

Oracle Based Active Set Algorithm for Scalable Elastic Net Subspace Clustering

Chong You[†] Chun-Guang Li^{*} Daniel P. Robinson[‡] René Vidal[†]

[†]Center for Imaging Science, Johns Hopkins University

^{*}SICE, Beijing University of Posts and Telecommunications

[‡]Applied Mathematics and Statistics, Johns Hopkins University

Abstract

State-of-the-art subspace clustering methods are based on expressing each data point as a linear combination of other data points while regularizing the matrix of coefficients with ℓ_1 , ℓ_2 or nuclear norms. ℓ_1 regularization is guaranteed to give a subspace-preserving affinity (i.e., there are no connections between points from different subspaces) under broad theoretical conditions, but the clusters may not be connected. ℓ_2 and nuclear norm regularization often improve connectivity, but give a subspace-preserving affinity only for independent subspaces. Mixed ℓ_1 , ℓ_2 and nuclear norm regularizations offer a balance between the subspace-preserving and connectedness properties, but this comes at the cost of increased computational complexity. This paper studies the geometry of the elastic net regularizer (a mixture of the ℓ_1 and ℓ_2 norms) and uses it to derive a provably correct and scalable active set method for finding the optimal coefficients. Our geometric analysis also provides a theoretical justification and a geometric interpretation for the balance between the connectedness (due to ℓ_2 regularization) and subspace-preserving (due to ℓ_1 regularization) properties for elastic net subspace clustering. Our experiments show that the proposed active set method not only achieves state-of-the-art clustering performance, but also efficiently handles large-scale datasets.

1. Introduction

In many computer vision applications, including image representation and compression [19], motion segmentation [6, 41, 34], temporal video segmentation [39], and face clustering [18], high-dimensional datasets can be well approximated by a union of low-dimensional subspaces. In this case, the problem of clustering a high-dimensional dataset into multiple classes or categories reduces to the problem of assigning each data point to its own subspace and recovering the underlying low-dimensional structure of the data, a problem known in the literature as *subspace clustering* [37].

Prior Work. Over the past decade, the subspace clustering

problem has received a lot of attention in the literature and many methods have been developed. Among them, spectral clustering based methods have become extremely popular [4, 48, 9, 10, 26, 25, 13, 38, 28, 7, 17, 32, 23, 36] (see [37] for details). These methods usually divide the problem into two steps: a) learning an affinity matrix that characterizes whether two points are likely to lie in the same subspace, and b) applying spectral clustering to this affinity. Arguably, the first step is the most important, as the success of spectral clustering depends on having an appropriate affinity matrix.

State-of-the-art methods for constructing the affinity matrix are based on the *self-expressiveness model* [9]. Under this model, each data point \mathbf{x}_j is expressed as a linear combination of all other data points, i.e., $\mathbf{x}_j = \sum_{i \neq j} \mathbf{x}_i c_{ij} + \mathbf{e}_j$, where the coefficient c_{ij} is used to define an affinity between points i and j , and the vector \mathbf{e}_j captures deviations from the self-expressive model. The coefficients are typically found by solving an optimization problem of the form

$$\min_{\mathbf{c}_j, \mathbf{e}_j} r(\mathbf{c}_j) + \gamma \cdot h(\mathbf{e}_j) \quad \text{s.t.} \quad \mathbf{x}_j = X\mathbf{c}_j + \mathbf{e}_j, c_{jj} = 0, \quad (1)$$

where $X = [\mathbf{x}_1, \dots, \mathbf{x}_N]$ is the data matrix, $\mathbf{c}_j = [c_{1j}, \dots, c_{Nj}]^\top$ is the vector of coefficients, $r(\cdot)$ is a properly chosen regularizer on the coefficients, $h(\cdot)$ is a properly chosen regularizer on the noise or corruption, and $\gamma > 0$ is a parameter that balances these two regularizers.

The main difference among state-of-the-art methods lies in the choice of the regularizer $r(\cdot)$. The sparse subspace clustering (SSC) method [9] searches for a sparse representation using $r(\cdot) = \|\cdot\|_1$. While under broad theoretical conditions (see [10, 35, 47]) the representation produced by SSC is guaranteed to be *subspace preserving* (i.e., $c_{ij} \neq 0$ only if \mathbf{x}_i and \mathbf{x}_j are in the same subspace), the affinity matrix may lack *connectedness* [29] (i.e., data points from the same subspace may not form a connected component of the affinity graph due to the sparseness of the connections, which may cause over-segmentation). Other recently proposed sparsity based methods, such as orthogonal matching pursuit (OMP) [7, 46] and nearest subspace neighbor (NSN) [32], also suffer from the same connectivity issue.

As an alternative, the least squares regression (LSR)

method [28] uses the regularizer $r(\cdot) = \frac{1}{2} \|\cdot\|_2^2$. One benefit of LSR is that the representation matrix is generally dense, which alleviates the connectivity issue of sparsity based methods. However, the representation is known to be subspace preserving only when the subspaces are independent,¹ which significantly limits its applicability. Nuclear norm regularization based methods, such as low rank representation (LRR) [26] and low rank subspace clustering (LRSC) [13, 38], also suffer from the same limitation [44].

To bridge the gap between the subspace preserving and connectedness properties, [44, 31, 12] propose to use mixed norms. For example, the low rank sparse subspace clustering (LRSSC) method [44], which uses a mixed ℓ_1 and nuclear norm regularizer, is shown to give a subspace preserving representation under conditions which are similar to but stronger than those of SSC. However, the justification for the improvements in connectivity given by LRSSC is merely experimental. Likewise, [31, 12] propose to use a mixed ℓ_1 and ℓ_2 norm given by

$$r(\mathbf{c}) = \lambda \|\mathbf{c}\|_1 + \frac{1-\lambda}{2} \|\mathbf{c}\|_2^2, \quad (2)$$

where $\lambda \in [0, 1]$ controls the trade-off between the two regularizers. However, [31, 12] do not provide a theoretical justification for the benefits of the method. Other subspace clustering regularizers studied in [27] and [21] use the trace lasso [15] and the k -support norm [1], respectively. However, no theoretical justification is provided in [27, 21] for the benefit of their methods.

Another issue with the aforementioned methods [44, 31, 12, 27, 21] is that they do not provide efficient algorithms to deal with large-scale datasets. To address this issue, [5] proposes to find the representation of X by a few anchor points that are sampled from X and then perform spectral clustering on the anchor graph. In [33] the authors propose to cluster a small subset of the original data and then classify the rest of the data based on the learned groups. However, both of these strategies are suboptimal in that they sacrifice clustering accuracy for computational efficiency.

Paper Contributions. In this paper, we exploit a mixture of ℓ_1 and ℓ_2 norms to balance the subspace preserving and connectedness properties. Specifically, we use $r(\cdot)$ as in (2) and $h(\mathbf{e}) = \frac{1}{2} \|\mathbf{e}\|_2^2$. The method is thus a combination of SSC and LSR and reduces to each of them when $\lambda = 1$ and $\lambda = 0$, respectively. In the statistics literature, the optimization program using this regularization is called *Elastic Net* and is used for variable selection in regression problems [49]. Thus we refer to this method as the Elastic Net Subspace Clustering (EnSC).

This work makes the following contributions:

1. We propose an efficient and provably correct active-set based algorithm for solving the elastic net prob-

lem. The proposed algorithm exploits the fact that the nonzero entries of the elastic net solution fall into an *oracle region*, which we use to define and efficiently update an active set. The proposed update rule leads to an iterative algorithm which is shown to converge to the optimal solution in a finite number of iterations.

2. We provide theoretical conditions under which the affinity generated by EnSC is subspace preserving, as well as a clear geometric interpretation for the balance between the subspace-preserving and connectedness properties. Our conditions depend on a *local* characterization of the distribution of the data, which improves over prior *global* characterizations.
3. We present experiments on computer vision datasets that demonstrate the superiority of our method in terms of both clustering accuracy and scalability.

2. Elastic Net: Geometry and a New Algorithm

In this section, we study the elastic net optimization problem, and present a new active-set based optimization algorithm for solving it. Consider the objective function

$$f(\mathbf{c}; \mathbf{b}, A) := \lambda \|\mathbf{c}\|_1 + \frac{1-\lambda}{2} \|\mathbf{c}\|_2^2 + \frac{\gamma}{2} \|\mathbf{b} - A\mathbf{c}\|_2^2, \quad (3)$$

where $\mathbf{b} \in \mathbb{R}^D$, $A = [\mathbf{a}_1, \dots, \mathbf{a}_N] \in \mathbb{R}^{D \times N}$, $\gamma > 0$, and $\lambda \in [0, 1]$ (the reader is referred to [45] for a study of the case $\lambda = 1$). Without loss of generality, we assume that \mathbf{b} and $\{\mathbf{a}_j\}_{j=1}^N$ are normalized to be of unit ℓ_2 norm in our analysis. The elastic net model then computes

$$\mathbf{c}^*(\mathbf{b}, A) := \arg \min_{\mathbf{c}} f(\mathbf{c}; \mathbf{b}, A). \quad (4)$$

We note that $\mathbf{c}^*(\mathbf{b}, A)$ is unique since $f(\mathbf{c}; \mathbf{b}, A)$ is a strongly convex function; we use the notation \mathbf{c}^* in place of $\mathbf{c}^*(\mathbf{b}, A)$ when the meaning is clear.

In the next two sections, we present a geometric analysis of the elastic net solution, and use this analysis to design an active-set algorithm for efficiently solving (4).

2.1. Geometric structure of the elastic net solution

We first introduce the concept of an oracle point.

Definition 2.1 (Oracle Point). *The oracle point associated with the optimization problem (4) is defined to be*

$$\delta(\mathbf{b}, A) := \gamma \cdot (\mathbf{b} - A\mathbf{c}^*(\mathbf{b}, A)). \quad (5)$$

When there is no risk of confusion, we omit the dependency of the oracle point on \mathbf{b} and A and write $\delta(\mathbf{b}, A)$ as δ .

Notice that the oracle point is unique since \mathbf{c}^* is unique, and that the oracle point cannot be computed until the optimal solution \mathbf{c}^* has been computed. The next result gives a critical relationship involving the oracle point that is exploited by our active-set method (see [45] for the proof).

¹Subspaces $\{\mathcal{S}_\kappa\}$ are independent if $\dim(\sum_{\kappa} \mathcal{S}_\kappa) = \sum_{\kappa} \dim(\mathcal{S}_\kappa)$.




Figure 1. Illustration of the structure of the solution \mathbf{c}^* for a data matrix A containing 100 randomly generated points in \mathbb{R}^2 , which are shown as blue dots in the x - y plane. The z direction shows the magnitude for each coefficient c_j^* . The red dot represents the oracle point $\delta(\mathbf{b}, A)$, with its direction denoted by the red dashed line. The value for γ is fixed at 50, but the value for λ varies as depicted.

Theorem 2.1. *The solution \mathbf{c}^* to problem (4) satisfies*

$$(1 - \lambda)\mathbf{c}^* = \mathcal{T}_\lambda(A^\top \delta), \quad (6)$$

where $\mathcal{T}_\lambda(\cdot)$ is the soft-thresholding operator (applied componentwise to $A^\top \delta$) defined as $\mathcal{T}_\lambda(v) = \text{sgn}(v)(|v| - \lambda)$ if $|v| > \lambda$ and 0 otherwise.

Theorem 2.1 shows that if the oracle point δ is known, the solution \mathbf{c}^* can be written out directly. Moreover, it follows from (5) and (6) that $\delta = \mathbf{0}$ if and only if $\mathbf{b} = \mathbf{0}$.

In Figure 1, we depict a two dimensional example of the solution to the elastic net problem (4) for different values of the tradeoff parameter λ . As expected, the solution \mathbf{c}^* becomes denser as λ decreases. Moreover, as predicted by Theorem 2.1, the magnitude of the coefficient c_j^* is a decaying function of the angle between the corresponding dictionary atom \mathbf{a}_j and the oracle point δ (shown in red). If \mathbf{a}_j is far enough from δ such that $|\langle \mathbf{a}_j, \delta \rangle| \leq \lambda$ holds true, then the corresponding coefficient c_j^* is zero. We call the region containing the nonzero coefficients the *oracle region*. We can formally define the oracle region by using the quantity $\mu(\cdot, \cdot)$ to denote the coherence of two vectors, i.e.,

$$\mu(\mathbf{v}, \mathbf{w}) := \frac{|\langle \mathbf{v}, \mathbf{w} \rangle|}{\|\mathbf{v}\|_2 \|\mathbf{w}\|_2}. \quad (7)$$

Definition 2.2 (Oracle Region). *The oracle region associated with the optimization problem (4) is defined as*

$$\Delta(\mathbf{b}, A) := \left\{ \mathbf{v} \in \mathbb{R}^D : \|\mathbf{v}\|_2 = 1, \mu(\mathbf{v}, \delta) > \frac{\lambda}{\|\delta\|_2} \right\}. \quad (8)$$

The oracle region is composed of an antipodal pair of spherical caps of the unit ball of \mathbb{R}^D that are located at the symmetric locations $\pm \delta / \|\delta\|_2$, both with an angular radius of $\theta = \arccos(\lambda / \|\delta\|_2)$ (see Figure 2). From the definition of the oracle region and Theorem 2.1, it follows that $c_j^* \neq 0$ if and only if $\mathbf{a}_j \in \Delta(\mathbf{b}, A)$. In other words, the support of the solution \mathbf{c}^* are those vectors \mathbf{a}_j in the oracle region.

The oracle region also captures the behavior of the solution when columns from the matrix A are removed or new columns are added. This provides the key insight into designing an active-set method for solving the optimization.




Figure 2. The oracle region $\Delta(\mathbf{b}, A)$ is illustrated in red. Note that the size of the oracle region increases as the quantity $\lambda / \|\delta\|_2$ decreases, and vice versa.

Proposition 2.1. *For any $\mathbf{b} \in \mathbb{R}^D$, $A \in \mathbb{R}^{D \times N}$ and $A' \in \mathbb{R}^{D \times N'}$, if no column of A' is contained in $\Delta(\mathbf{b}, A)$, then $\mathbf{c}^*(\mathbf{b}, [A, A']) = [\mathbf{c}^*(\mathbf{b}, A)^\top, \mathbf{0}_{N' \times 1}^\top]^\top$.*

The interpretation for Proposition 2.1 is that the solution $\mathbf{c}^*(\mathbf{b}, A)$ does not change (modulo padding with additional zeros) when new columns are added to the dictionary A , as long as the new columns are not inside the oracle region $\Delta(\mathbf{b}, A)$. From another perspective, $\mathbf{c}^*(\mathbf{b}, [A, A'])$ does not change if one removes columns from the dictionary $[A, A']$ that are not in the oracle region $\Delta(\mathbf{b}, [A, A'])$.

Proposition 2.2. *For any $\mathbf{b} \in \mathbb{R}^D$, $A \in \mathbb{R}^{D \times N}$ and $A' \in \mathbb{R}^{D \times N'}$, denote $\mathbf{c}^*(\mathbf{b}, [A, A']) = [\mathbf{c}_A^\top, \mathbf{c}_{A'}^\top]^\top$. If any column of A' lies within $\Delta(\mathbf{b}, A)$, then $\mathbf{c}_{A'} \neq \mathbf{0}$.*

This result means that the solution to the elastic net problem will certainly be changed by adding new columns that lie within the oracle region to the dictionary.

In the next section, we describe an efficient algorithm for solving the elastic net problem (4) that is based on the geometric structure and concentration behavior of the solution.

2.2. A new active-set algorithm

Although the elastic net optimization problem [49] has been recently introduced for subspace clustering in [11, 12, 31], such prior work does not provide an efficient algorithm that can handle large-scale datasets. In fact, such prior work




Figure 3. Conceptual illustration of the ORGEN algorithm. All the dots on the unit circle illustrate the dictionary A . (a) active set T_k at step k , illustrated by red dots. (b) The oracle region $\Delta(\mathbf{b}, A_{T_k})$ illustrated by red arcs. (c) The new active set T_{k+1} illustrated in green, which is the set of indices of points that are in $\Delta(\mathbf{b}, A_{T_k})$.

solves the elastic net problem using existing algorithms that require calculations involving the full data matrix A (e.g., the accelerated proximal gradient (APG) [2] is used in [11] and the linearized alternating direction method (LADM) [24] is used in [31]). Here, we propose to solve the elastic net problem (4) with an active-set algorithm that is more efficient than both APG and LADM, and can handle large-scale datasets. We call our new method (see Algorithm 1) ORacle Guided Elastic Net solver, or ORGEN for short.

The basic idea behind ORGEN is to solve a sequence of reduced-scale subproblems defined by an active set that is itself determined from the oracle region. Let T_k be the active set at iteration k . Then, the next active set T_{k+1} is selected to contain the indices of columns that are in the oracle region $\Delta(\mathbf{b}, A_{T_k})$, where A_{T_k} denotes the submatrix of A with columns indexed by T_k . We use Figure 3 for a conceptual illustration. In Figure 3(a) we show the columns of A that correspond to the active set T_k by labeling the corresponding columns of A_{T_k} in red. The oracle region $\Delta(\mathbf{b}, A_{T_k})$ is the union of the red arcs in Figure 3(b). Notice that at the bottom left there is one red dot that is not in $\Delta(\mathbf{b}, A_{T_k})$ and thus must not be included in T_{k+1} , and two blue dots that are not in T_k but lie in the oracle region $\Delta(\mathbf{b}, A_{T_k})$ and thus must be included in T_{k+1} . In Figure 3(c) we illustrate T_{k+1} by green dots. This iterative procedure is terminated once T_{k+1} does not contain any new points, i.e., when $T_{k+1} \subseteq T_k$, at which time T_{k+1} is the support for $\mathbf{c}^*(\mathbf{b}, A)$.

The next lemma helps explain why ORGEN converges.

Lemma 2.1. *In Algorithm 1, if $T_{k+1} \not\subseteq T_k$, then*

$$f(\mathbf{c}^*(\mathbf{b}, A_{T_{k+1}}); \mathbf{b}, A_{T_{k+1}}) < f(\mathbf{c}^*(\mathbf{b}, A_{T_k}); \mathbf{b}, A_{T_k}).$$

The following convergence result holds for ORGEN.

Theorem 2.2. *Algorithm 1 converges to the optimal solution $\mathbf{c}^*(\mathbf{b}, A)$ in a finite number of iterations.*

The result follows from Lemma 2.1, because it implies that an active set can never be repeated. Since there are only finitely many distinct active sets, the algorithm must eventually terminate with $T_{k+1} \subseteq T_k$. The remaining part of

Algorithm 1 ORacle Guided Elastic Net (ORGEN) solver

Input: $A = [\mathbf{a}_1, \dots, \mathbf{a}_N] \in \mathbb{R}^{D \times N}$, $\mathbf{b} \in \mathbb{R}^D$, λ and γ .

- 1: Initialize the support set T_0 and set $k \leftarrow 0$.
- 2: **loop**
- 3: Compute $\mathbf{c}^*(\mathbf{b}, A_{T_k})$ as in (4) using any solver.
- 4: Compute $\delta(\mathbf{b}, A_{T_k})$ from $\mathbf{c}^*(\mathbf{b}, A_{T_k})$ as in (5).
- 5: Active set update: $T_{k+1} \leftarrow \{j : \mathbf{a}_j \in \Delta(\mathbf{b}, A_{T_k})\}$.
- 6: If $T_{k+1} \subseteq T_k$, terminate; otherwise set $k \leftarrow k + 1$.
- 7: **end loop**

Output: A vector \mathbf{c} such that $\mathbf{c}_{T_k} = \mathbf{c}^*(\mathbf{b}, A_{T_k})$ and zeros otherwise. Its support is T_{k+1} .

the proof establishes that if $T_{k+1} \subseteq T_k$, then $\mathbf{c}^*(\mathbf{b}, A_{T_{k+1}})$ gives the nonzero entries of the solution.

ORGEN solves large-scale problems by solving a sequence of reduced-size problems in step 3 of Algorithm 1. If the active set T_k is small, then step 3 is a small-scale problem that can be efficiently solved. However, there is no procedure in Algorithm 1 that explicitly controls the size of T_k . To address this concern, we propose an alternative to step 5 in which only a small number of new points—the ones most correlated with δ —are added. Specifically,

$$5' : T_{k+1} = \{j \in T_k : \mathbf{a}_j \in \Delta(\mathbf{b}, A_{T_k})\} \cup S_k, \quad (9)$$

where S_k holds the indices of the largest n entries in $\{|\mathbf{a}_j^\top \delta(\mathbf{b}, A_{T_k})| : j \notin T_k, \mathbf{a}_j \in \Delta(\mathbf{b}, A_{T_k})\}$; ideally, n should be chosen so that the size of T_{k+1} is bounded by a predetermined value N_{\max} that represents the maximum size subproblem that can be handled in step 3. If N_{\max} is chosen large enough that the second set in the union in (9) is non-empty, then our convergence result still holds.

Initialization. We suggest the following procedure for computing the initial active set T_0 . First, compute the solution to (4) with $\lambda = 0$, which has a closed form solution and can be computed efficiently if the ambient dimension D of the data is not too big. Then, the l largest entries (in absolute value) of the solution for some pre-specified value l are added to T_0 . Our experiments suggest that this strategy promotes fast convergence of Algorithm 1.

3. Elastic Net Subspace Clustering (EnSC)

Although the elastic net has been recently introduced for subspace clustering in [31, 12], these works do not provide conditions under which the affinity is guaranteed to be subspace preserving or potential improvements in connectivity. In this section, we give conditions for the affinity to be subspace preserving and for the balance between the subspace-preserving and connectedness properties. To the best of our knowledge, this is the first time that such theoretical guarantees have been established.

We first formally define the subspace clustering problem.

Problem 3.1 (Subspace Clustering). Let $X \in \mathbb{R}^{D \times N}$ be a real-valued matrix whose columns are drawn from a union of n subspaces of \mathbb{R}^D , say $\bigcup_{\ell=1}^n \mathcal{S}_\ell$, where the dimension d_ℓ of the ℓ -th subspace satisfies $d_\ell < D$ for $\ell = 1, \dots, n$. The goal of subspace clustering is to segment the columns of X into their representative subspaces.

Let $X = [\mathbf{x}_1, \dots, \mathbf{x}_N]$, where each \mathbf{x}_j is assumed to be of unit norm. Using the same notation as for (4), the proposed EnSC computes $\mathbf{c}^*(\mathbf{x}_j, X_{-j})$ for each $\{\mathbf{x}_j\}_{j=1}^N$, i.e.,

$$\mathbf{c}^*(\mathbf{x}_j, X_{-j}) = \arg \min_{\mathbf{c}} f(\mathbf{c}; \mathbf{x}_j, X_{-j}), \quad (10)$$

where X_{-j} is X with the j -th column removed. In this section, we focus on a given vector, say \mathbf{x}_j . We suppose that $\mathbf{x}_j \in \mathcal{S}_\ell$ for some ℓ , and use X_{-j}^ℓ to denote the submatrix of X with columns from \mathcal{S}_ℓ except that \mathbf{x}_j is removed. Since our goal is to use the entries of $\mathbf{c}^*(\mathbf{x}_j, X_{-j})$ to construct an affinity graph in which only points in the same subspace are connected, we desire the nonzero entries of $\mathbf{c}^*(\mathbf{x}_j, X_{-j})$ to be a subset of the columns X_{-j}^ℓ so that no connections are built between points from different subspaces. If this is the case, we say that such a solution $\mathbf{c}^*(\mathbf{x}_j, X_{-j})$ is *subspace preserving*. On the other hand, we also want the nonzero entries of $\mathbf{c}^*(\mathbf{x}_j, X_{-j})$ to be as dense as possible in X_{-j}^ℓ so that within each cluster the affinity graph is well-connected². To some extent, these are conflicting goals: if the connections are few, it is more likely that the solution is subspace preserving, but the affinity graph of each cluster is not well connected. Conversely, as one builds more connections, it is more likely that some of them will be false, but the connectivity is improved.

In the next two sections, we give a geometric interpretation of the tradeoff between the subspace preserving and connectedness properties, and provide sufficient conditions for a representation to be subspace preserving.

3.1. Subspace-preserving vs. connected solutions

Our analysis is built upon the optimization problem $\min_{\mathbf{c}} f(\mathbf{c}; \mathbf{x}_j, X_{-j}^\ell)$. Note that its solution is trivially subspace preserving since the dictionary X_{-j}^ℓ is contained in \mathcal{S}_ℓ . We then treat all points from other subspaces as newly added columns to X_{-j}^ℓ and apply Propositions 2.1 and 2.2. We get the following geometric result.

Lemma 3.1. Suppose that $\mathbf{x}_j \in \mathcal{S}_\ell$. Then, the vector $\mathbf{c}^*(\mathbf{x}_j, X_{-j})$ is subspace preserving if and only if $\mathbf{x}_k \notin \Delta(\mathbf{x}_j, X_{-j}^\ell)$ for all $\mathbf{x}_k \notin \mathcal{S}_\ell$.

We illustrate the geometry implied by Lemma 3.1 in Figure 4, where we assume \mathcal{S}_ℓ is a two dimensional subspace




Figure 4. The structure of the solution for an example in \mathbb{R}^3 associated with a point \mathbf{x}_j (not shown) that lies in the 2-dimensional subspace \mathcal{S}_ℓ . The blue dots illustrate the columns of X_{-j}^ℓ , the union of the two red regions is the oracle region $\Delta(\mathbf{x}_j, X_{-j}^\ell)$, and the green points are vectors from other subspaces.

in \mathbb{R}^3 . The dictionary X_{-j}^ℓ is represented by the blue dots in the plane and the oracle region $\Delta(\mathbf{x}_j, X_{-j}^\ell)$ is denoted as the two red circles. The green dots are all other points in the dictionary. Lemma 3.1 says that $\mathbf{c}^*(\mathbf{x}_j, X_{-j})$ is subspace preserving if and only if all green dots lie outside of the red region.

To ensure that a solution is subspace preserving one desires a small oracle region, while to ensure connectedness one desires a large oracle region. These facts again highlight the trade-off between these two properties. Recall that the elastic net balances ℓ_1 regularization (promotes sparse solutions) and ℓ_2 regularization (promotes dense solutions). Thus, one should expect that the oracle region will decrease in size as λ is increased from 0 towards 1. Theorem 3.1 formalizes this claim, but first we need the following definition that characterizes the distribution of the data in X_{-j}^ℓ .

Definition 3.1 (inradius). The inradius of a convex body \mathcal{P} is the radius $r(\mathcal{P})$ of the largest ℓ_2 ball inscribed in \mathcal{P} .

To understand the next result, we comment that the size of the oracle region $\Delta(\mathbf{x}_j, X_{-j}^\ell)$ is controlled by the quantity $\lambda / \|\delta(\mathbf{x}_j, X_{-j}^\ell)\|_2$ as depicted in Figure 2.

Theorem 3.1. If $\mathbf{x}_j \in \mathcal{S}_\ell$, then

$$\frac{\lambda}{\|\delta(\mathbf{x}_j, X_{-j}^\ell)\|_2} \geq \frac{r_j^2}{r_j + \frac{1-\lambda}{\lambda}}, \quad (11)$$

where r_j is the inradius of the convex hull of the symmetrized points in X_{-j}^ℓ , i.e.,

$$r_j := r(\text{conv}\{\pm \mathbf{x}_k : \mathbf{x}_k \in \mathcal{S}_\ell \text{ and } k \neq j\}). \quad (12)$$

We define the right-hand-side of (11) to be zero when $\lambda = 0$.

The above theorem allows us to determine an upper bound for the size of the oracle region. This follows since a lower bound on the size of $\lambda / \|\delta(\mathbf{x}_j, X_{-j}^\ell)\|_2$ implies an

²In fact, even when each cluster is well-connected, further improving connectivity within clusters is still beneficial since it enhances the ability of the subsequent step of spectral clustering in correcting any erroneous connections in the affinity graph [42, 40].

upper bound on the size of the oracle region (see (8) and Figure 2). Also notice that the right hand side of (11) is in the range $[0, r_j)$ and is monotonically increasing with λ . Thus, it provides an upper bound on the area of the oracle region, which decreases as λ increases. This highlights that the trade-off between the subspace-preserving and connectedness properties is controlled by λ .

Remark 3.1. *It would be nice if $\lambda/\|\delta(\mathbf{x}_j, X_{-j}^\ell)\|_2$ was increasing as a function of λ (we already know that its lower bound given in Theorem 3.1 is increasing in λ). However, one can show using the data $\mathbf{x}_j = [0.22, 0.72, 0.66]^\top$,*

$$X_{-j}^\ell = \begin{bmatrix} -0.55 & -0.82 & -0.05 & 0.22 \\ 0.22 & 0.57 & 0.84 & 0.78 \\ -0.80 & 0.00 & 0.55 & 0.58 \end{bmatrix}, \quad (13)$$

and parameter choice $\gamma = 10$, that $\lambda/\|\delta\|$ (with $\lambda = 0.88$) is larger than $\lambda/\|\delta\|$ (with $\lambda = 0.95$).

3.2. Conditions for a subspace-preserving solution

A sufficient condition for a solution to be subspace preserving is obtained by combining the geometry in Lemma 3.1 with the bound on the size of the oracle region implied by Theorem 3.1.

Theorem 3.2. *Let $\mathbf{x}_j \in \mathcal{S}_\ell$, $\delta_j = \delta(\mathbf{x}_j, X_{-j}^\ell)$ be the oracle point, and r_j be the inradius characterization of X_{-j}^ℓ as given by (12). Then, $\mathbf{c}^*(\mathbf{x}_j, X_{-j})$ is subspace preserving if*

$$\max_{k:\mathbf{x}_k \notin \mathcal{S}_\ell} \mu(\mathbf{x}_k, \delta_j) \leq \frac{r_j^2}{r_j + \frac{1-\lambda}{\lambda}}. \quad (14)$$

Notice that in Theorem 3.2 the quantity δ_j is determined from X_{-j}^ℓ and that it lies within the subspace \mathcal{S}_ℓ by definition of $\delta(\mathbf{x}_j, X_{-j}^\ell)$. Thus the left-hand-side of (14) characterizes the separation between the oracle point—which is in \mathcal{S}_ℓ —and the set of points outside of \mathcal{S}_ℓ . On the right-hand-side, r_j characterizes the distribution of points in X_{-j}^ℓ . In particular, r_j is large when points are well spread within \mathcal{S}_ℓ and not skewed toward any direction. Finally, note that the right-hand-side of (14) is an increasing function of λ , showing that the solution is more likely to be subspace preserving if more weight is placed on the ℓ_1 regularizer relative to the ℓ_2 regularizer.

Theorem 3.2 has a close relationship to the sufficient condition for SSC to give a subspace preserving solution (the case $\lambda = 1$) [35]. Specifically, [35] shows that if $\max_{k:\mathbf{x}_k \notin \mathcal{S}_\ell} \mu(\mathbf{x}_k, \delta_j) < r_j$, then SSC gives a subspace preserving solution. We can observe that condition (14) approaches the condition for SSC as $\lambda \rightarrow 1$.

The result stated in Theorem 3.2 is a special case of the following more general result.

Theorem 3.3. *Let $\mathbf{x}_j \in \mathcal{S}_\ell$, $\delta_j = \delta(\mathbf{x}_j, X_{-j}^\ell)$ be the oracle point, and $\kappa_j = \max_{k \neq j, \mathbf{x}_k \in \mathcal{S}_\ell} \mu(\mathbf{x}_k, \delta_j)$ be the coherence of δ_j with its nearest neighbor in X_{-j}^ℓ . Then, the solution $\mathbf{c}^*(\mathbf{x}_j, X_{-j})$ is subspace preserving if*

$$\max_{k:\mathbf{x}_k \notin \mathcal{S}_\ell} \mu(\mathbf{x}_k, \delta_j) \leq \frac{\kappa_j^2}{\kappa_j + \frac{1-\lambda}{\lambda}}. \quad (15)$$

The only difference between this result and that in Theorem 3.2 is that κ_j is used instead of r_j for characterizing the distribution of points in X_{-j}^ℓ . We show in [45] that $r_j \leq \kappa_j$, which makes Theorem 3.3 more general than Theorem 3.2. Geometrically, r_j is large if the subspace \mathcal{S}_ℓ is well-covered by X_{-j}^ℓ , while κ_j is large if the neighborhood of the oracle point closest to δ_j is well-covered, i.e., there is a point in X_{-j}^ℓ that is close to δ_j . Thus, while the condition in Theorem 3.2 requires each subspace to have global coverage by the data, the condition in Theorem 3.3 allows the data to be biased, and only requires a local region to be well-covered. In addition, condition (15) can be checked when the membership of the data points is known. This advantage allows us to check the tightness of the condition (15), which is studied in more details in [45]. In contrast, condition (14) and previous work on SSC [35, 43] use the inradius r_j , which is generally NP-hard to calculate [35, 44].

4. Experiments

4.1. ORGEN on synthetic data

We conducted synthetic experiments to illustrate the computational efficiency of the proposed algorithm ORGEN. Three popular solvers are exploited: the regularized feature sign search (RFSS) is an active set type method [20]; the LASSO version of the LARS algorithm [8] that is implemented in the sparse modeling software (SPAMS); and the gradient projection for sparse reconstruction (GPSR) algorithm proposed in [14]. These three solvers are used to solve the subproblem in step 3 of ORGEN, resulting in three implementations of ORGEN. We also used the three solvers as stand-alone solvers for comparison purposes.

In all experiments, the vector \mathbf{b} and columns of A are all generated independently and uniformly at random on the unit sphere of \mathbb{R}^{100} . The results are averages over 50 trials.

In the first experiment, we test the scaling behavior of ORGEN by varying N ; the results are shown in Figure 5(a). We can see that our active-set scheme improves the computational efficiency of all three solvers. Moreover, as N grows, the improvement becomes more significant.

Next, we test the performance of ORGEN for various values of the parameter λ that controls the tradeoff between the subspace preserving and connectedness properties; the running times and sparsity level are shown in Figures 5(b) and 5(c), respectively. The performance of SPAMS is not

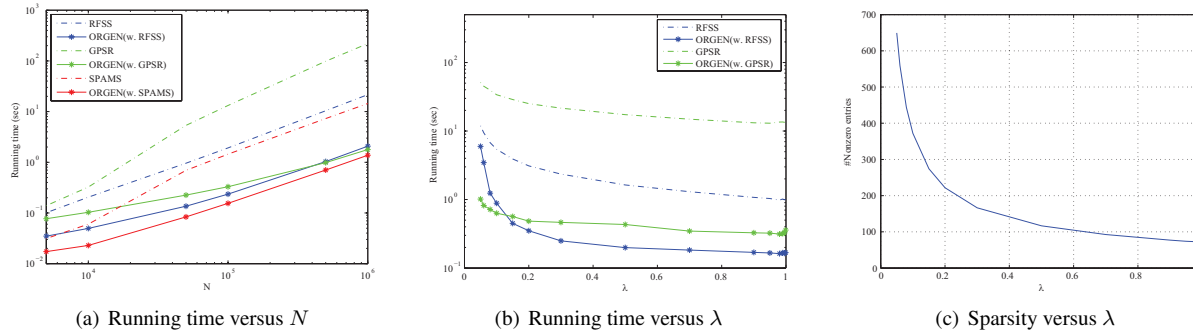


Figure 5. Performance with varying N and λ : (a) $\lambda = 0.9$ and $N \in [5000, 10^6]$; and (b, c) $N = 100,000$ and $\lambda \in [0.05, 0.999]$.

Table 1. Dataset information.

	N (#data)	D (ambient dim.)	n (#groups)
Coil-100	7,200	1024	100
PIE	11,554	1024	68
MNIST	70,000	500	10
CovType	581,012	54	7

reported since it performs poorly even for moderately small values of λ . For all methods, the computational efficiency decreases as λ becomes smaller. For the two versions of ORGEN, this is expected since the solution becomes denser as λ becomes smaller (see Figure 5(c)). Thus the active sets become larger, which leads directly to larger and more time consuming subproblems in step 3.

4.2. EnSC on real data

In this section, we use ORGEN to solve the optimization problems arising in EnSC, where each subproblem in step 3 is solved using the RFSS method. To compute the coefficient vectors $c^*(x_j, X_{-j})$, the parameter λ is set to be the same for all j , while the parameter γ is set as $\gamma = \alpha\gamma_0$ where $\alpha > 1$ is a hyperparameter and γ_0 is the smallest value of γ such that $c^*(x_j, X_{-j})$ is nonzero. The algorithm is run for at most 2 iterations, as we observe that this is sufficient for the purpose of subspace clustering and that subsequent iterations do not boost performance. We measure clustering performance by clustering accuracy, which is calculated as the best matching rate between the label predicted by the algorithm and that of the ground truth.

Datasets. We test our method on the four datasets presented in Table 1. The Coil-100 dataset [30] contains 7,200 grayscale images of 100 different objects. Each object has 72 images taken at pose intervals of 5 degrees, with the images being of size 32×32 . The PIE dataset [16] contains images of the faces of 68 people taken under 13 different poses, 43 different illuminations, and 4 different expressions. In the experiments, we use the five near frontal poses and all images under different illuminations and expressions. Each image is manually cropped and normalized to 32×32 pixels. The MNIST dataset [22] contains 70,000 images of

handwritten digits 0–9. For each image, we extract a feature vector of dimension 3,472 via the scattering convolution network [3], and then project to dimension 500 using PCA. Finally, the Covtype database³ has been collected to predict forest cover type from 54 cartographic variables.

Methods. We compare our method with several state-of-the-art subspace clustering methods that may be categorized into three groups. The first group contains TSC [17], OMP [7], NSN [32], and SSC [9]. TSC is a variant of the k -nearest neighbors method, OMP and NSN are two sparse greedy methods, and SSC is a convex optimization method. These algorithms build sparse affinity matrices and are computationally efficient, and therefore can perform large-scale clustering. For TSC and NSN we use the code provided by the respective authors. We note that the code may not be optimized for computational efficiency considerations. For OMP, we use our implementation, which has been optimized for subspace clustering. For SSC we use the SPAMS solver described in the previous section.

The second group consists of LRSC and SSC (with a different solver). We use the code provided by their respective authors, which uses the Alternating Direction Method of Multipliers (ADMM) to solve the optimization problems. To distinguish the two versions of SSC, we refer to this one as SSC-ADMM and to the previous one as SSC-SPAMS.

The final group consists of ENSC [31] and KMP [21], and are the closest in spirit to our method. Our method and ENSC both balance the ℓ_1 and ℓ_2 regularizations, but ENSC uses $h(e) = \|e\|_1$ to penalize the noise (see (1)) and the linearized alternating direction method to minimize their objective. In KMP, the k -support norm is used to blend the ℓ_1 and ℓ_2 regularizers. We implemented ENSC and KMP according to the descriptions in their original papers.

Results. To the best of our knowledge, a comparison of all these methods on large scale datasets has not been reported in prior work. Thus, we run all experiments and tune the parameters for each method to give the best clustering accuracy. The results are reported in Table 2.

³<http://archive.ics.uci.edu/ml/datasets/Covertype>

Table 2. Performance of different clustering algorithms. The running time includes the time for computing the affinity matrix and for performing spectral clustering. The sparsity is the number of nonzero coefficients in each representation c_j averaged over $j = 1, \dots, N$. The value “M” means that the memory limit of 16GB was exceeded, and the value “T” means that the time limit of seven days was reached.

	TSC	OMP	NSN	SSC-SPAMS	SSC-ADMM	LRSC	ENSC	KMP	EnSC-ORGEN
<i>Clustering accuracy (%)</i>									
Coil-100	61.32	33.64	50.32	53.75	57.10	55.76	51.11	61.97	69.24
PIE	22.15	11.28	35.02	39.05	41.94	46.65	21.40	16.55	52.98
MNIST	85.00	46.84	85.82	92.46	M	M	M	M	93.79
CovType	35.45	48.76	38.04	T	M	M	M	M	53.52
<i>Running time (min.)</i>									
Coil-100	2	2	11	16	127	3	8	63	3
PIE	3	8	25	67	412	12	25	361	13
MNIST	30	24	298	1350	-	-	-	-	28
CovType	999	783	3572	-	-	-	-	-	1452
<i>Sparsity</i>									
Coil-100	4	2	18	7.0	5.4	7199	7199	20.9	6.3
PIE	8	22	17	20.4	28.5	11553	11553	30.0	82.4
MNIST	8	10	12	25.4	-	-	-	-	26.6
CovType	20	15	10	-	-	-	-	-	34.9

We see that our proposed method achieves the best clustering performance on every dataset. Our method is also among the most efficient in terms of computing time. The methods SSC-ADMM, ENSC, LRSC and KMP cannot handle large-scale data because they perform calculations over the full data matrix and put the entire kernel matrix $X^T X$ in memory, which is infeasible for large datasets. The method of SSC-SPAMS uses an active set method that can deal with massive data, however, it is computationally much less efficient than our solver ORGEN.

For understanding the advantages of our method, in Table 2 we report the sparsity of the representation coefficients, which is the number of nonzero entries in c_j averaged over all $j = 1, \dots, N$. For TSC, OMP and NSN, the sparsity is directly provided as a parameter of the algorithms. For SSC and our method EnSC-ORGEN, the sparsity is indirectly controlled by the parameters of the models. We can see that our method usually gives more nonzero entries than the sparsity based methods of TSC, OMP, NSN, and SSC. This shows the benefit of our method: while the number of correct connections built by OMP, NSN and SSC are in general upper-bounded by the dimension of the subspace, our method does not have this limit and is capable of constructing more correct connections and producing well-connected affinity graphs. On the other hand, the affinity graph of LRSC is dense, so although each cluster is self-connected, there are abundant wrong connections. This highlights the advantage of our method, which is flexible in controlling the number of nonzero entries by adjusting the trade-off parameter λ . Our results illustrate that this trade-off improves clustering accuracy.

Finally, ENSC and KMP are two representatives of other methods that also exploit the trade-off between ℓ_1 and ℓ_2

regularizations. A drawback of both works is that the solvers for their optimization problems are not as effective as our ORGEN algorithm, as they cannot deal with large datasets due to memory requirements. Moreover, we observe that their algorithms converge to modest accuracy in a few iterations but can be very slow in giving a high precision solution. This may explain why their clustering accuracy is not as good as that of EnSC-ORGEN. Especially, we see that ENSC gives dense solutions although the true solution is expected to be sparser, and this is explained by the fact that the solution paths of their solver are dense solutions.

5. Conclusion

We investigated elastic net regularization (i.e., a mixture of the ℓ_1 and ℓ_2 norms) for scalable and provable subspace clustering. Specifically, we presented an active set algorithm that efficiently solves the elastic net regularization subproblem by capitalizing on the geometric structure of the elastic net solution. We then gave theoretical justifications—based on a geometric interpretation for the trade-off between the subspace preserving and connectedness properties—for the correctness of subspace clustering via the elastic net. Extensive experiments verified that that our proposed active set method achieves state-of-the-art clustering accuracy and can handle large-scale datasets.

Acknowledgments. C. You, D. P. Robinson and R. Vidal are supported by the National Science Foundation under grant 1447822. C.-G. Li is partially supported by National Natural Science Foundation of China under grants 61273217 and 61511130081, and the 111 project under grant B08004. The authors thank Ben Haeffele for insightful comments on the design of the ORGEN algorithm.

References

- [1] A. Argyriou, R. Foygel, and N. Srebro. Sparse prediction with the k -support norm. In *Neural Information Processing Systems*, pages 1466–1474, 2012. 2
- [2] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, Mar 2009. 4
- [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] G. Chen and G. Lerman. Spectral curvature clustering (SCC). *International Journal of Computer Vision*, 81(3):317–330, 2009. 1
- [5] X. Chen and D. Cai. Large scale spectral clustering with landmark-based representation. In *AAAI Conference on Artificial Intelligence*, 2011. 2
- [6] J. Costeira and T. Kanade. A multibody factorization method for independently moving objects. *International Journal of Computer Vision*, 29(3):159–179, 1998. 1
- [7] 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. 1, 7
- [8] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *Annals of Statistics*, 32(2):407–499, 2004. 6
- [9] E. Elhamifar and R. Vidal. Sparse subspace clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2009. 1, 7
- [10] 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
- [11] Y. Fang, R. Wang, and B. Dai. Graph-oriented learning via automatic group sparsity for data analysis. In *IEEE International Conference on Data Mining*, pages 251–259, 2012. 3, 4
- [12] Y. Fang, R. Wang, B. Dai, and X. Wu. Graph-based learning via auto-grouped sparse regularization and kernelized extension. *IEEE Transactions on Knowledge and Data Engineering*, 27(1):142–154, 2015. 2, 3, 4
- [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. 1, 2
- [14] M. Figueiredo, R. Nowak, and S. Wright. Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems. *IEEE Journal of Selected Topics in Signal Processing*, 1(4):586–597, 2007. 6
- [15] E. Grave, G. Obozinski, and F. Bach. Trace lasso: a trace norm regularization for correlated designs. In *Neural Information Processing Systems*, 2011. 2
- [16] R. Gross, I. Matthews, J. F. Cohn, T. Kanade, and S. Baker. Multi-PIE. *Image Vision Computing*, 28(5):807–813, 2010. 7
- [17] R. Heckel and H. Bölcskei. Robust subspace clustering via thresholding. *CoRR*, abs/1307.4891, 2013. 1, 7
- [18] J. Ho, M. H. Yang, J. Lim, K. Lee, and D. Kriegman. Clustering appearances of objects under varying illumination conditions. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2003. 1
- [19] W. Hong, J. Wright, K. Huang, and Y. Ma. Multi-scale hybrid linear models for lossy image representation. *IEEE Transactions on Image Processing*, 15(12):3655–3671, 2006. 1
- [20] B. Jin, D. Lorenz, and S. Schiffler. Elastic-net regularization: error estimates and active set methods. *Inverse Problems*, 25(11), 2009. 6
- [21] 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, 7
- [22] 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
- [23] C.-G. Li and R. Vidal. Structured sparse subspace clustering: A unified optimization framework. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 277–286, 2015. 1
- [24] Z. Lin, R. Liu, and Z. Su. Linearized alternating direction method with adaptive penalty for low rank representation. In *Neural Information Processing Systems*, 2011. 4
- [25] 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
- [26] G. Liu, Z. Lin, and Y. Yu. Robust subspace segmentation by low-rank representation. In *International Conference on Machine Learning*, 2010. 1, 2
- [27] C. Lu, Z. Lin, and S. Yan. Correlation adaptive subspace segmentation by trace lasso. In *IEEE International Conference on Computer Vision*, 2013. 2
- [28] 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, 2
- [29] B. Nasihatkon and R. Hartley. Graph connectivity in sparse subspace clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2011. 1
- [30] S. A. Nene, S. K. Nayar, and H. Murase. Columbia object image library (COIL-100). *Technical Report CUCS-006-96*, 1996. 7
- [31] Y. Panagakis and C. Kotropoulos. Elastic net subspace clustering applied to pop/rock music structure analysis. *Pattern Recognition Letters*, 38:46–53, 2014. 2, 3, 4, 7
- [32] D. Park, C. Caramanis, and S. Sanghavi. Greedy subspace clustering. In *Neural Information Processing Systems*, 2014. 1, 7
- [33] X. Peng, L. Zhang, and Z. Yi. Scalable sparse subspace clustering. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 430–437, 2013. 2
- [34] 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

- [35] M. Soltanolkotabi and E. J. Candès. A geometric analysis of subspace clustering with outliers. *Annals of Statistics*, 2013. [1](#), [6](#)
- [36] M. Tsakiris and R. Vidal. Filtrated spectral algebraic subspace clustering. In *ICCV Workshop on Robust Subspace Learning and Computer Vision*, pages 28–36, 2015. [1](#)
- [37] R. Vidal. Subspace clustering. *IEEE Signal Processing Magazine*, 28(3):52–68, March 2011. [1](#)
- [38] R. Vidal and P. Favaro. Low rank subspace clustering (LRSC). *Pattern Recognition Letters*, 43:47–61, 2014. [1](#), [2](#)
- [39] R. Vidal, Y. Ma, and S. Sastry. Generalized Principal Component Analysis (GPCA). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(12):1–15, 2005. [1](#)
- [40] R. Vidal, Y. Ma, and S. Sastry. *Generalized Principal Component Analysis*. Springer Verlag, 2016. [5](#)
- [41] 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](#)
- [42] U. von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17, 2007. [5](#)
- [43] Y.-X. Wang and H. Xu. Noisy sparse subspace clustering. In *International Conference on Machine learning*, 2013. [6](#)
- [44] Y.-X. Wang, H. Xu, and C. Leng. Provable subspace clustering: When LRR meets SSC. In *Neural Information Processing Systems*, 2013. [2](#), [6](#)
- [45] C. You, C.-G. Li, D. Robinson, and R. Vidal. Oracle based active set algorithm for scalable elastic net subspace clustering. *Arxiv*, 2016. [2](#), [6](#)
- [46] C. You, D. Robinson, and R. Vidal. Scalable sparse subspace clustering by orthogonal matching pursuit. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. [1](#)
- [47] C. You and R. Vidal. Geometric conditions for subspace-sparse recovery. In *International Conference on Machine Learning*, pages 1585–1593, 2015. [1](#)
- [48] 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](#)
- [49] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society, Series B*, 67:301–320, 2005. [2](#), [3](#)