

Motion Segmentation in the Presence of Outlying, Incomplete, or Corrupted Trajectories

Shankar Rao, *Member, IEEE*, Roberto Tron, *Student Member, IEEE*,
René Vidal, *Member, IEEE*, and Yi Ma, *Senior Member, IEEE*

Abstract—In this paper, we study the problem of segmenting tracked feature point trajectories of multiple moving objects in an image sequence. Using the affine camera model, this problem can be cast as the problem of segmenting samples drawn from multiple linear subspaces. In practice, due to limitations of the tracker, occlusions, and the presence of nonrigid objects in the scene, the obtained motion trajectories may contain grossly mistracked features, missing entries, or corrupted entries. In this paper, we develop a robust subspace separation scheme that deals with these practical issues in a unified mathematical framework. Our methods draw strong connections between lossy compression, rank minimization, and sparse representation. We test our methods extensively on the Hopkins155 motion segmentation database and other motion sequences with outliers and missing data. We compare the performance of our methods to state-of-the-art motion segmentation methods based on expectation-maximization and spectral clustering. For data without outliers or missing information, the results of our methods are on par with the state-of-the-art results and, in many cases, exceed them. In addition, our methods give surprisingly good performance in the presence of the three types of pathological trajectories mentioned above. All code and results are publicly available at <http://perception.csl.uiuc.edu/coding/motion/>.

Index Terms—Motion segmentation, subspace separation, lossy compression, incomplete data, error correction, sparse representation, matrix rank minimization.

1 INTRODUCTION

A fundamental problem in computer vision is to infer structures and movements of 3D objects from a video sequence. While classical multiple-view geometry typically deals with the situation where the scene is static, recently there has been growing interest in the analysis of dynamic scenes. Such scenes often contain multiple motions as there could be multiple objects moving independently in a scene, in addition to the motion of the camera. Thus, an important initial step in the analysis of video sequences is the *motion segmentation* problem. That is, given a set of feature points that are tracked through a sequence of video frames, one seeks to cluster the trajectories of those points according to the different motions these trajectories belong to.

In the literature, many different camera models have been proposed and studied, such as orthographic, paraperspective, affine, and perspective. Among these, the affine camera model (which includes orthographic and paraperspective) is arguably the most popular, due largely

to its generality and simplicity. Thus, in this paper, we assume the affine camera model and show how to develop a more robust solution to the motion segmentation problem.

1.1 Basic Formulation of Motion Segmentation

Under the affine camera model, a feature point in 3D space $(X, Y, Z) \in \mathbb{R}^3$ is related to its projection on the image plane $(x, y) \in \mathbb{R}^2$ by

$$\begin{bmatrix} x \\ y \end{bmatrix} = \underbrace{\mathbf{K} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}}_{\mathbf{A} \in \mathbb{R}^{2 \times 4}} \begin{bmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}, \quad (1)$$

where \mathbf{A} is the *affine motion matrix*, parameterized by the camera calibration matrix $\mathbf{K} \in \mathbb{R}^{2 \times 3}$ and the relative orientation of the image plane with respect to the world coordinates $(\mathbf{R}, \mathbf{t}) \in SE(3)$.

Suppose we are given trajectories of P tracked feature points of a rigid object $\{(x_{fp}, y_{fp})\}_{f=1 \dots P}^P$ from F 2D image frames taken by a moving camera. The linear constraints in (1) can be combined for multiple points across multiple frames so that the tracked feature points are related to their 3D coordinates $\{(X_p, Y_p, Z_p)\}_{p=1}^P$ by the matrix equation:

$$\underbrace{\begin{bmatrix} x_{11} & x_{12} & \cdots & x_{1P} \\ y_{11} & y_{12} & \cdots & y_{1P} \\ \vdots & \vdots & \ddots & \vdots \\ x_{F1} & x_{F2} & \cdots & x_{FP} \\ y_{F1} & y_{F2} & \cdots & y_{FP} \end{bmatrix}}_{\mathbf{Y} \in \mathbb{R}^{2F \times P}} = \underbrace{\begin{bmatrix} \mathbf{A}_1 \\ \vdots \\ \mathbf{A}_F \end{bmatrix}}_{\mathbf{A} \in \mathbb{R}^{2F \times 4}} \underbrace{\begin{bmatrix} X_1 & \cdots & X_P \\ Y_1 & \cdots & Y_P \\ Z_1 & \cdots & Z_P \\ 1 & \cdots & 1 \end{bmatrix}}_{\mathbf{X} \in \mathbb{R}^{4 \times P}}, \quad (2)$$

$$\mathbf{Y} = \mathbf{A}\mathbf{X},$$

- S. Rao is with HRL Laboratories, LLC, 3011 Malibu Canyon Rd., Malibu, CA 90265. E-mail: srrao@hrl.com.
- R. Tron and R. Vidal are with the Center for Imaging Science, Department of Biomedical Engineering, Johns Hopkins University, 301B Clark Hall, 3400 N. Charles St., Baltimore, MD 21218. E-mail: {tron, rvidal}@cis.jhu.edu.
- Y. Ma is with the Electrical and Computer Engineering Department, University of Illinois at Urbana-Champaign, Coordinated Science Laboratory, 1308 W. Main St., Urbana, IL 61801 and with Microsoft Research Asia, 5/F, Beijing Sigma Center, No. 49, Zhichun Road, Hai Dian District, Beijing, China 100190. E-mail: yima@illinois.edu.

Manuscript received 3 Dec. 2008; revised 30 May 2009; accepted 3 Sept. 2009; published online 18 Nov. 2009.

Recommended for acceptance by J. Kosecka.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-2008-12-0832.

Digital Object Identifier no. 10.1109/TPAMI.2009.191.

where A_f is the affine motion matrix at frame f . From this formulation, we see that

$$\text{rank}(Y) = \text{rank}(AX) \leq \min(\text{rank}(A), \text{rank}(X)) \leq 4. \quad (3)$$

Thus, the affine camera model postulates that trajectories of feature points from a single rigid motion will all lie in a linear subspace of \mathbb{R}^{2F} of dimension at most four.

A dynamic scene can contain multiple moving objects, in which case the affine camera model for a single rigid motion cannot be directly applied. Now let us assume that the given P trajectories correspond to N moving objects. In this case, the set of all trajectories will lie in a *union of N linear subspaces* in \mathbb{R}^{2F} (see, for instance, [27] for details), but we do not know which trajectories belong to which subspace. Thus, the problem of assigning each trajectory to its corresponding motion reduces to the problem of segmenting data drawn from multiple subspaces, which we refer to as *subspace separation*.

Problem 1 (Motion segmentation via subspace separation). *Given a set of trajectories of P feature points $Y = [y_1, y_2, \dots, y_P] \in \mathbb{R}^{2F \times P}$ from N rigidly moving objects in a dynamic scene, find a permutation Γ of the columns of the data matrix Y :*

$$Y\Gamma = [Y_1, Y_2, \dots, Y_N] \quad (4)$$

such that the columns of each submatrix Y_n , $n = 1, \dots, N$, are trajectories of a single motion.

1.2 Related Work on Motion Segmentation

In the literature, there are many approaches to motion segmentation that can roughly be grouped into three categories: factorization-based, algebraic, and statistical.

Factorization-based approaches [5], [12], [16], [17] attempt to directly factor Y according to (4). To make such approaches tractable, the motions must be independent of one another, i.e., the motion subspaces intersect only at the origin. However, for most dynamic scenes with a moving camera or containing articulated objects, the motions are at least partially dependent on each other. This has motivated the development of algorithms designed to deal with dependent motions.

Algebraic methods such as Generalized Principal Component Analysis (GPCA) [26] are designed as generic subspace separation algorithms that do not place any restriction on the relative orientations of the motion subspaces. For instance, they allow the subspaces to intersect into lower dimensional subspaces, and hence, they can deal with partially dependent motions. In principle, algebraic methods such as GPCA can be extended to deal with missing data [27] and outliers [32]. However, its complexity grows *exponentially* with respect to both the dimension of the ambient space and the number of motions in the scene, and so is not scalable in practice.

The statistical methods come in many flavors. Many formulate motion segmentation as a statistical clustering problem that is tackled with Expectation-Maximization (EM) or variations of it [14], [18], [23]. As such, they are iterative methods that require good initialization, and can potentially get stuck in suboptimal local minima. Other statistical methods use local information around each trajectory to create a pairwise similarity matrix that can then be segmented using *spectral clustering* techniques [10], [31], [33].

1.3 Robustness Issue

Many of the above approaches assume that all trajectories are good, with perhaps a moderate amount of noise. However, real motion data acquired by a tracker can be much more complicated:

1. A trajectory may correspond to certain nonrigid or random motions that do not obey the affine camera model (an *outlying trajectory*).
2. Some of the features may be missing in some frames, causing a trajectory to have some missing entries (an *incomplete trajectory*).
3. Even worse, some feature points may be mis-tracked (with the tracker unaware), causing a trajectory to have some entries with gross errors (a *corrupted trajectory*).

While some of the methods can be modified to be robust to *one* of such problems [10], [14], [27], [31], [32], to our knowledge there is no motion segmentation algorithm that can elegantly deal with all of these problems in a unified fashion.

1.4 Our Approach

In order to uniformly and effectively deal with clustering and robustness issues, we rely on Occam's Razor: *All other things being equal, the simplest solution is the best*. This means that when choosing among multiple viable segmentations for motion data, one should pick the segmentation that most simply explains the data. There are many empirical metrics that can be used to express the simplicity of data. One such measure is the *coding length*, which is the minimal number of bits needed to represent data. The coding length has been used effectively for data compression and model selection [1] as well as for segmentation [20]. In recent years, there has been increasing interest in findings representations for data that are *sparse*, i.e., having few nonzero entries. This interest has been mainly fueled by the discovery that when the sparsity is high enough, such representations can be efficiently computed using convex optimization [4], [8]. The sparse structure of data has also been shown to be highly robust and can be used to deal with incomplete and corrupted data [3].

In this paper, we propose a new motion segmentation scheme that draws heavily from the principles of both *data compression* and *sparse representation*. We show that the notion of coding length and sparsity are highly related, and by properly exploiting them, we are able to make motion segmentation robust to all three types of pathological trajectories listed above. In particular, we adapt the lossy compression-based agglomerative clustering algorithm from [20], referred to as *Agglomerative Lossy Compression (ALC)*, to the problem of motion segmentation. The algorithm is noniterative and requires only a single parameter. We will show how it can be naturally adapted to deal with outliers in our context. We supplement ALC with techniques from sparse representation, allowing our method to handle incomplete and corrupted trajectories even before the segmentation is obtained. To our knowledge, our paper is the first to apply sparse representation to the problem of motion segmentation.

1.4.1 Organization of This Paper

We first review our agglomerative algorithm (Section 2.1), then show how we apply the derived algorithm to motion segmentation (Section 2.2), and test the effectiveness of the algorithm on the publicly available Hopkins155 motion segmentation database (Section 2.3). We show that the new algorithm naturally handles outlying trajectories (Section 3.1), and can be extended to repair incomplete (Section 3.2) or corrupted trajectories (Section 3.3). Note our distinction between incomplete and corrupted trajectories: For incomplete trajectories, we know in which frames the features are missing; for corrupted ones, we do not have that knowledge. Our methods use the affine camera model assumption, so we make comparisons with similar methods, but not with perspective camera-based methods.¹ As most extant methods for motion segmentation assume that the number of motions is known, for fair comparison we also assume that the group count is given.

2 AGGLOMERATIVE LOSSY COMPRESSION (ALC)

In this section, we describe the subspace separation method that we use for motion segmentation. Section 2.1 reviews the principles of matrix rank minimization, data compression, and sparse representation behind ALC. Section 2.2 shows how ALC can be applied to the motion segmentation problem when the motion trajectories are complete and contain no outliers. Finally, Section 2.3 shows the results of the segmentation algorithm on the Hopkins155 database (which does not contain outliers).

2.1 Matrix Rank Minimization and Lossy Data Compression

According to the problem formulation (4), to a large extent the goal of subspace separation is to find a partition of the data matrix Y into submatrices $\{Y_n\}_{n=1}^N$ such that each Y_n spans a subspace of the lowest possible dimension. In other words, each Y_n as a matrix is maximally rank deficient. *Matrix rank minimization* (MRM) is itself a very challenging problem. The rank function is neither smooth nor convex, and it is notoriously difficult to minimize directly. Finding a matrix M^* that is maximally rank deficient among a convex set of matrices Ω is known to be NP-Hard [25]. Also, the rank function is highly unstable in the presence of noise.

Recent progress in compressed sensing has led to some groundbreaking work in rank minimization. In particular, it has been shown that when the matrix rank is low enough, minimizing the matrix rank over a convex domain is equivalent to minimizing the matrix nuclear norm² $\|M\|_*$, which can be solved efficiently by semidefinite programming [21]. If $\bar{\Omega}$ is a convex set of *symmetric positive semidefinite* matrices, one can find the minimum rank $M^* \in \bar{\Omega}$ by solving

$$M^* = \operatorname{argmin}_{M \in \bar{\Omega}} J_\delta(M) \doteq \log_2 \det \left(I + \frac{1}{\delta} M \right), \quad (5)$$

1. Refer to [22] for work on robust motion segmentation with a perspective camera model.

2. The nuclear norm of a matrix M is the sum of all of its singular values $\sum_i \sigma_i$.

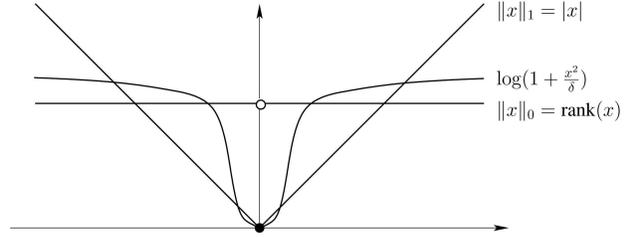


Fig. 1. Comparison of $J_\delta(x)$, $\operatorname{rank}(x) = \|x\|_0$, and the nuclear norm (1-norm) $\|x\|_1 = |x|$ in one dimension.

where the constant $\delta > 0$ is a small regularization parameter [11]. It is easy to see that the function J_δ is approximately the sum of the logarithm of the singular values (up to a scale). So, unlike the nuclear norm, which is convex, the function J_δ is no longer convex in \mathbb{M} , though it is a smooth surrogate. Nevertheless, $J_\delta(\mathbb{M})$ has the same global minimum as $\operatorname{rank}(\mathbb{M})$, as shown in Fig. 1 (for each singular value).

Here, we are not minimizing $\operatorname{rank}(Y_n)$ over a convex set. Recall that each Y_n is a submatrix of Y , and the set $\{Y_n\}_{n=1}^N$ forms a partition of Y . The number of partitions of the data matrix Y into $\{Y_n\}_{n=1}^N$ is combinatorial and this makes the space of all segmentations of Y a very complicated domain. Thus, technically subspace separation *cannot* be reduced to an instance of MRM over a convex domain. However, with slight modification to the function $J_\delta(\mathbb{M})$ in (5), we can see a clear connection between (5) and the principle of (lossy) *minimum description length* (MDL) [20]. Given a data matrix $Y_n \in \mathbb{R}^{D \times P_n}$, Ma et al. [20] proposed the following function for estimating the number of bits needed to code Y_n up to distortion ε^2 :

$$\begin{aligned} L(Y_n, \varepsilon) &\doteq \frac{D + P_n}{2} J_{\frac{\varepsilon^2}{D}} \left(\frac{1}{P_n} Y_n Y_n^T \right) \\ &= \frac{D + P_n}{2} \log_2 \det \left(I + \frac{D}{P_n \varepsilon^2} Y_n Y_n^T \right). \end{aligned} \quad (6)$$

This function is still a smooth surrogate for $\operatorname{rank}(Y_n)$ as it is obtained by scaling $J_\delta(\mathbb{M})$ by a constant term with $\mathbb{M} = \frac{1}{P_n} Y_n Y_n^T$ and $\delta = \frac{\varepsilon^2}{D}$. It was shown in [20] that as $\varepsilon \rightarrow 0$, (6) converges to the optimal rate distortion for a Gaussian source, and it is also a tight upper bound for the coding length of subspace-like data. Thus, (6) provides a reasonable estimate of the number of bits needed to code a set of samples drawn from either a Gaussian distribution or a linear subspace.

Now suppose the data matrix $Y \in \mathbb{R}^{D \times P}$ can be partitioned into disjoint subsets $Y = [Y_1 \dots Y_N]$ of corresponding sizes $P_1 + \dots + P_N = P$. If we encode each subset separately, the total number of bits required to encode Y up to distortion ε^2 is

$$L^s(\{Y_1, \dots, Y_N\}, \varepsilon) \doteq \sum_{n=1}^N L(Y_n, \varepsilon) - P_n \log_2 \frac{P_n}{P}. \quad (7)$$

The second term in this equation counts the number of bits needed to represent the membership of the P vectors in the N subsets (e.g., by Huffman coding). In [20], Ma et al. showed that for data drawn from a mixture of multiple (degenerate) Gaussians, the segmentation that minimizes

(7) is, in fact, the segmentation that partitions the samples into groups corresponding to different Gaussians in the mixture. Thus, by finding the global minimum of (7), we also find the “true” segmentation of the data. It is worth noticing that once the distortion parameter ε is fixed, the number of groups in the segmentation is automatically determined. This completely avoids the necessity of additional model selection criterion usually required with traditional segmentation methods.

The issue now is that finding a global minimum of (7) is a combinatorial problem. Nevertheless, an agglomerative algorithm, proposed in [20], has been shown to be very effective for minimizing (7). Algorithm 1, listed below, initially treats each sample as its own group, iteratively merging pairs of groups so that the resulting coding length is maximally reduced at each iteration. The algorithm terminates when it can no longer reduce the coding length. We refer to Algorithm 1 as *ALC*. See [20] for more details.

Algorithm 1. Agglomerative Lossy Compression.

```

1: Input:  $Y = [y_1, y_2, \dots, y_P] \in \mathbb{R}^{D \times P}$ ,  $\varepsilon \in \mathbb{R}$ 
2: Let  $\mathcal{S} = \{\{y_1\}, \dots, \{y_P\}\}$ 
3: done := false
4: while not done do
5:    $\{Y_i^*, Y_j^*\} := \operatorname{argmin}_{\{Y_i, Y_j\} \in \mathcal{S}} L^s(\{\{Y_i, Y_j\}\}, \varepsilon) - L^s(\{Y_i, Y_j\}, \varepsilon)$ 
6:   if  $L^s(\{\{Y_i^*, Y_j^*\}\}, \varepsilon) - L^s(\{Y_i^*, Y_j^*\}, \varepsilon) \geq 0$  then
7:     done := true
8:   else
9:      $\mathcal{S} := (\mathcal{S} \setminus \{Y_i^*, Y_j^*\}) \cup \{\{Y_i^*, Y_j^*\}\}$ 
10:  end if
11: end while
12: output:  $\mathcal{S}$ 
    
```

2.2 ALC and Motion Segmentation

In this section, we explore many of the practical issues with applying ALC to the motion segmentation problem. We first show how trajectories from many types of motion lie on low-dimensional linear subspaces. We propose a simple and effective method for choosing the distortion level ε , the single parameter required by ALC. We then describe how dimensionality reduction techniques can be used to improve both the convergence and tractability of our agglomerative approach. Finally, we discuss the computational complexity of the method and show how it can be improved.

2.2.1 Motions as Linear Subspaces

As seen in Section 1.1, under the affine camera model, trajectories from a single general rigid-body motion lie on a linear subspace of dimensionality four in the $2F$ -dimensional trajectory space. For motions along a line or within a plane, their corresponding trajectories lie on subspaces of dimensionality two or three, respectively. Thus, in a dynamic scene with multiple motions, trajectories from the different motions lie on multiple subspaces with possibly different dimensionalities. Because ALC is designed to cluster data from multiple subspaces of mixed dimensionalities, it should be highly effective for affine motion segmentation.

We can also use linear subspaces to model many kinds of nonrigid motion. In an articulated motion, the motion

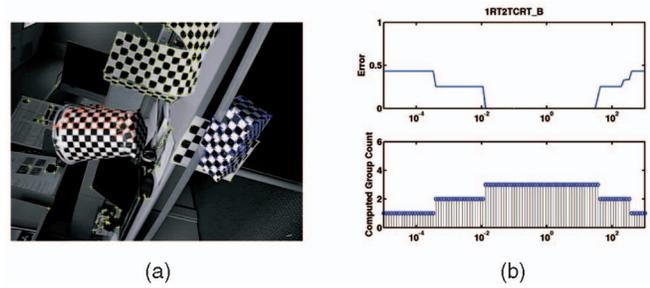


Fig. 2. (a) The “1RT2TCRT_B” sequence from the Hopkins155 database. (b) The misclassification rate (top) and estimated group count (bottom) as a function of ε .

consists of two or more “submotions” that are joined at a link. Each submotion is a rigid-body motion whose trajectories lie on a subspace of dimensionality at most four. However, the motions of the linked parts are dependent, and consequently, their subspaces have intersections of dimensionality one or two, depending on whether the link is a joint or an axis. Other nonrigid motions, such as facial expressions, can be approximated by a piecewise linear combination of a number of “key shapes.” As shown in [31], the trajectories of a nonrigid motion with K key shapes will lie on a linear subspace of dimensionality at most $3K + 1$. ALC should, in principle, be able to segment trajectories in scenes containing articulated and nonrigid motions. However, due to the greedy nature of ALC, the dependencies between submotions, and higher dimensional subspace embeddings, for scenes with these kinds of motions, it is possible for ALC to obtain suboptimal segmentations.

2.2.2 Choosing the Distortion Level ε

In principle, ε could be determined in some heuristic fashion from the statistics of the data, see, e.g., [17]. However, note that the distortion level ε is directly related to the number of motions N : The smaller ε is, the larger N is and vice versa. Since most extant motion segmentation algorithms require the number of motions as a parameter, in order to make a fair comparison with other methods, we assume that the number of motions is given, and use it to determine ε . Fig. 2 shows an example motion sequence. We run ALC on this sequence for several choices of ε . On the right, we plot the misclassification rate and estimated group count as a function of ε . We see that the correct segmentation is stable over a fairly large interval. Using this observation, we developed the following voting scheme:

1. For a given motion sequence, run the algorithm multiple times over a number of choices of ε .³
2. Discard any ε that does not give rise to a segmentation with the correct number of groups.⁴
3. With the remaining choices of ε , find all of the distinct segmentations that are produced.
4. Choose the ε that minimizes the coding length for the most segmentations, relative to the other choices of ε .

This scheme is quite simple, and by no means optimal, but as our experiments show, it works very well in practice.

3. Our experiments use 101 steps of ε in the interval $[10^{-5}, 10^3]$.

4. If none of the choices of ε produce the right number of groups, we select the ε that minimizes the “penalized” coding length proposed in [20].

2.2.3 Improving ALC with Dimensionality Reduction

ALC applies a greedy approach to make minimization of (7) computationally feasible. Due to this greedy approach, ALC can obtain a segmentation that does not globally minimize the coding length. In fact, precise theoretical conditions for ALC to converge to the minimum coding length segmentation are not yet known. Ma et al. demonstrated empirically that for data in high-dimensional spaces, suboptimal segmentations can be found if the samples do not adequately cover each subspace [20]. Thus, dimensionality reduction can potentially improve the results of ALC by making the subspaces more dense with samples.

Dimensionality reduction can also improve the computational tractability of subspace separation without adversely affecting the quality of the segmentation. This is because, with probability one, projection onto an arbitrary d -dimensional subspace preserves the multisubspace structure of data lying on subspaces with dimensionality strictly less than d . Thus, for segmenting affine motions, Vidal et al. [27] suggest projecting the trajectories onto a 5D subspace. However, as we have discussed, for more complicated scenes, such as scenes with articulated or nonrigid motion, five dimensions may not be sufficient.

A balance needs to be struck between expressiveness and sample density. One choice, recently proposed in the sparse representation community [9], is the *sparsity-preserving* dimension d_{sp} :

$$d_{sp} = \min d \text{ subject to } d \geq 2k \log(D/d), \quad (8)$$

where D is the dimensionality of the ambient space and k is the true low dimensionality of the data. It has been shown that, asymptotically, as $D \rightarrow \infty$, this d is the smallest dimensionality of projection such that the low-dimensional multisubspace structure of the data is preserved with high probability under a *random* projection. For our problem, using the affine camera model, the dimensionality of the motion subspaces is at most 4, so we can assume that $k = 4$ and obtain a conservative estimate for the dimensionality of the projection d . As our experimental results will show, this choice works well in practice.

In our experiments, we test ALC with dimensionalities of projection $d = 5$ (as suggested in [27]), and the sparsity-preserving d stated above.⁵ We refer to the two versions of the algorithm as ALC_5 and ALC_{sp} , respectively.

2.2.4 Algorithmic Improvements to ALC.

As discussed in [20], the computational complexity of a straightforward implementation of ALC is

$$O(P^3 + P^2D^3). \quad (9)$$

The first term in (9) corresponds to, for each of $O(P)$ iterations, searching a table of size $O(P) \times O(P)$ for the pair of groups that, when merged, maximally decrease the overall coding length. The second term in (9) corresponds to, for each of $O(P)$ iterations, the cost of updating $O(P)$ entries in the table via an $O(D^3)$ log-determinant computation. In practice, the running time of ALC is dominated by this second term. We have observed empirically that the vast majority of the time, one of the two groups to be merged contains only one sample. In this case, the log-determinant

5. In our implementations of ALC, we use Principal Component Analysis (PCA) for dimensionality reduction.

TABLE 1
Misclassification Rates (in Percent) for Sequences of Two and Three Motions in the Hopkins155 Database

| Checkerboard | MSL | LSA | ALC_5 | ALC_{sp} |
|---------------|-------|--------------|---------|--------------|
| Average | 4.46% | 2.57% | 2.56% | 1.49% |
| Median | 0.00% | 0.27% | 0.00% | 0.27% |
| Traffic | MSL | LSA | ALC_5 | ALC_{sp} |
| Average | 2.23% | 5.43% | 2.83% | 1.75% |
| Median | 0.00% | 1.48% | 0.30% | 1.51% |
| Articulated | MSL | LSA | ALC_5 | ALC_{sp} |
| Average | 7.23% | 4.10% | 6.90% | 10.70% |
| Median | 0.00% | 1.22% | 0.89% | 0.95% |
| All Sequences | MSL | LSA | ALC_5 | ALC_{sp} |
| Average | 4.14% | 3.45% | 3.03% | 2.40% |
| Median | 0.00% | 0.59% | 0.00% | 0.43% |

(a)

| Checkerboard | MSL | LSA | ALC_5 | ALC_{sp} |
|---------------|--------------|--------|--------------|--------------|
| Average | 10.38% | 5.80% | 6.78% | 5.00% |
| Median | 4.61% | 1.77% | 0.92% | 0.66% |
| Traffic | MSL | LSA | ALC_5 | ALC_{sp} |
| Average | 1.80% | 25.07% | 4.01% | 8.86% |
| Median | 0.00% | 23.79% | 1.35% | 0.51% |
| Articulated | MSL | LSA | ALC_5 | ALC_{sp} |
| Average | 2.71% | 7.25% | 7.25% | 21.08% |
| Median | 2.71% | 7.25% | 7.25% | 21.08% |
| All Sequences | MSL | LSA | ALC_5 | ALC_{sp} |
| Average | 8.23% | 9.73% | 6.26% | 6.69% |
| Median | 1.76% | 2.33% | 1.02% | 0.67% |

(b)

(a) Two-motion sequences and (b) three-motion sequences.

can be computed via a rank-1 update to the Cholesky factorization of a $D \times D$ matrix [29]. By doing so, the computational complexity of ALC becomes

$$O(P^3 + P^2D^2 + PD^3), \quad (10)$$

allowing the speed of ALC to scale more gracefully with the dimensionality of the data. We quantitatively demonstrate this decrease in the running time of ALC in the next section.

2.3 Results on the Hopkins155 Database

We now test the efficacy of ALC for motion segmentation, by applying the algorithm to the Hopkins155 database [24]. The Hopkins155 database consists of 155 motion sequences that can be categorized as checkerboard, traffic, or articulated. The motion sequences were obtained using an automatic tracker, and errors in tracking were manually corrected for each sequence. Thus, in this experiment, there is no attempt to deal with incomplete or corrupted trajectories. See [24] for more details on the Hopkins155 database.

We run ALC_5 and ALC_{sp} on the checkerboard, traffic, and articulated sequences using the voting scheme described earlier to determine ε . For each category of sequences, we compute the average and median misclassification rates, and the average computation times. We list these results in Tables 1 and 2 along with the reported results for Multistage Learning (MSL) [18] and Local Subspace Affinity (LSA) [31]⁶

6. For LSA, we report the results for the version that projects the data onto a $4N$ -dimensional space.

TABLE 2
Performance over Entire Hopkins155 Database

| Checkerboard | MSL | LSA | ALC ₅ | ALC _{sp} |
|---------------|--------------|--------------|------------------|-------------------|
| Average | 5.94% | 3.38% | 3.61% | 2.37% |
| Median | 0.00% | 0.57% | 0.00% | 0.31% |
| Traffic | MSL | LSA | ALC ₅ | ALC _{sp} |
| Average | 2.15% | 9.05% | 3.05% | 3.06% |
| Median | 0.00% | 1.96% | 0.92% | 1.35% |
| Articulated | MSL | LSA | ALC ₅ | ALC _{sp} |
| Average | 6.53% | 4.58% | 6.95% | 12.30% |
| Median | 0.00% | 1.22% | 0.89% | 0.95% |
| All Sequences | MSL | LSA | ALC ₅ | ALC _{sp} |
| Average | 5.06% | 4.87% | 3.76% | 3.37% |
| Median | 0.00% | 0.90% | 0.26% | 0.49% |

(a)

| Method | MSL | LSA | ALC ₅ | ALC _{sp} |
|---------------|---------|---------|---------------------|---------------------|
| Checkerboard | 17h 40m | 10.423s | 12m 6s (6m 5s) | 24m 4s (7m 12s) |
| Traffic | 12h 42m | 8.433s | 8m 42s (4m 15s) | 17m 19s (4m 56s) |
| Articulated | 7h 35m | 3.551s | 4m 51s (2m 16s) | 10m 43s (2m 40s) |
| All Sequences | 15h 36m | 9.474s | 10m 32s (5m 15s) | 21m 3s (6m 11s) |

(b)

(a) Misclassification rates (in percent) and (b) average computation times. Results in parentheses for ALC use the rank-1 Cholesky update discussed in Section 2.2.

on the same database. Fig. 3 gives two histograms of the misclassification rates over the sequences with two and three motions, respectively. There are several other algorithms that have been tested on the Hopkins155 database (GPCA, RANSAC, etc.), but we list these two algorithms because they have, to date, the best reported misclassification rates in many categories of sequences.

As these results show, ALC performs well compared to the state of the art. It has the best overall misclassification rate as well as for the checkerboard sequences. In categories where ALC is not the best, its performance is still competitive. As expected, the performance of ALC for the articulated sequences is not as good, primarily because, in the Hopkins155 database, many of the scenes with articulated motions are, in fact, scenes of human motion with only a few tracked features.

In terms of computation time, we see that the algorithms fall into three categories: The spectral method LSA runs on the order of seconds, our agglomerative methods run on the order of minutes, and the iterative method MSL runs on the order of hours. Keep in mind that our methods are run for 101 different choices of the parameter ε . Also, by using the rank-1 Cholesky update, both ALC₅ and ALC_{sp} run two to four times faster on each sequence. Finally, with regard to the projection dimension, our results indicate that, overall, ALC_{sp} performs better than ALC₅.

3 ROBUST SUBSPACE SEPARATION

In this section, we show how to make subspace separation robust to the three kinds of pathologies discussed earlier. In

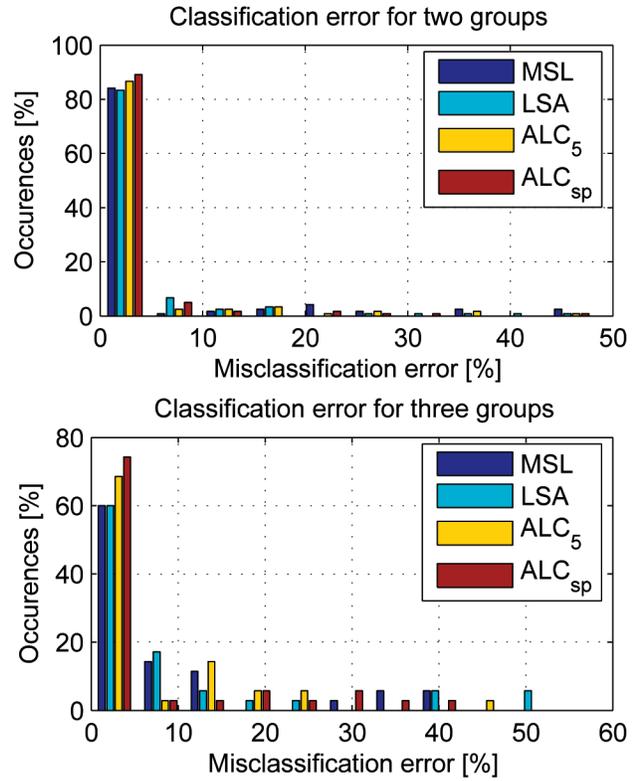


Fig. 3. Misclassification rate histograms for various algorithms on the Hopkins155 database.

particular, we show that ALC naturally deals with outliers, and by harnessing the low-dimensional subspace structure of the data, we can repair incomplete and corrupted samples prior to subspace separation.

3.1 Outlying Trajectories

Dynamic scenes often contain trajectories that do not correspond to any of the motion models in the scene. Such trajectories can arise from motions not well described by the affine camera model, such as the motion of nonrigid objects. These kinds of trajectories have been referred to as “sample outliers” by [7], suggesting that no subset of the trajectory corresponds to any affine motion model. Fortunately, ALC deals with such sample outliers in an elegant fashion. In [20], it was observed that in low-dimensional spaces, a sufficient number of outliers will cover the entire space, and so, the algorithm tends to group all outliers into a single group. Such a group can be easily detected because the number of bits per vector in that group will be very large relative to other groups. However, in higher dimensional spaces, such as in our motion segmentation problem, it would require an enormous number of outliers to fill the space. If outliers are thinly scattered in the ambient space, they will be most efficiently encoded when each outlier is in its own group. Such small groups are also easily detectable.

3.1.1 Experiments with Simulated Outliers

In these experiments, we compare the robustness to sample outliers of ALC⁷ and LSA [31], a spectral clustering-based

7. For this simulation, we use ALC₅, the version of ALC that projects the data onto a 5D space.

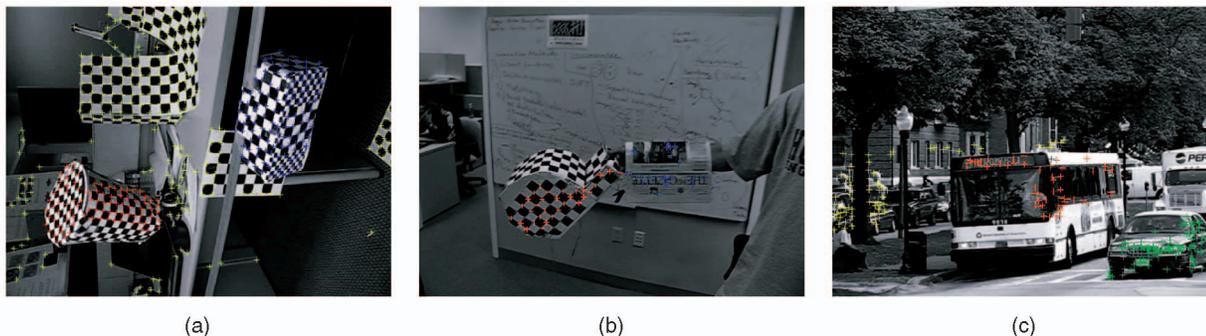


Fig. 4. Example image frames from three motion sequences from the Hopkins155 database [24]. (a) 1R2RC sequence. (b) Arm sequence. (c) cars10 sequence.

motion segmentation algorithm that is reasonably robust to outliers. We choose three representative sequences from the Hopkins155 database for simulation: “1R2RC” (checkerboard), “arm” (articulation), and “cars10” (traffic) (see Fig. 4). We add between 0 and 25 percent outlying trajectories to the data set of a given motion sequence. Outlying trajectories were generated by choosing a random initial point in the first frame and then selecting a random increment between successive frames. Each increment is generated by taking the difference between the coordinates of a randomly chosen point in two randomly chosen consecutive frames. In this way, the outlying trajectories will qualitatively have the same statistical properties as the other trajectories, but will not obey any particular motion model. We then input these data sets into LSA and ALC, respectively, and compute the misclassification rate and outlier detection rate for both algorithms.⁸ For each experiment, we run 100 trials with different randomly generated outlying trajectories. Table 3 shows the average misclassification rates and outlier detection rates for each experiment. As the results show, ALC can easily detect outliers without hindering motion segmentation, whereas for LSA, the outliers tend to interfere with the classification of valid trajectories. Hence, for subsequent experiments in this paper, we will not compare our methods with LSA.

3.1.2 Experiments with Real Outliers

We apply ALC to four motion sequences with real outlying trajectories, as shown in Fig. 5. For each sequence, trajectories were obtained with an automatic tracker, and the ground-truth segmentation was manually determined. A trajectory was termed an inlier if it is correctly tracked in all frames, and an outlier if it is incorrectly tracked in all frames.⁹ Information about the number of motions, samples for each group, and the number of outliers in each sequence are listed in Table 4.

The Misclassification and Outlier Detection rates are listed in Table 5. As these results show, ALC_{sp} is able to detect and remove real outliers without substantially

8. In ALC, a trajectory is labeled an outlier if it belongs to a group with less than five samples. In our implementation of LSA, a trajectory is labeled as an outlier if its distance from the nearest motion subspace is greater than a predetermined threshold.

9. For this experiment, trajectories with partial corruption were removed from the data set. This is because trajectories with partial corruption still retain a valid class label. Thus, it is better to deal with such trajectories as incomplete or corrupted, which we will discuss in Sections 3.2 and 3.3.

affecting segmentation of inliers, while ALC_5 is not. The one exception is the “carsbus3” sequence, where ALC_5 seems to outperform ALC_{sp} . However, qualitatively examining the segmentation results, we see that ALC_{sp} achieves its 9.74 percent misclassification rate by falsely grouping some outliers with features from the car in the foreground. However, these trajectories are fairly close to the car, so it could be argued that they are, in fact, noisy or corrupted trajectories rather than outliers. On the other hand, ALC_5 gets its low misclassification rate of 1.62 percent by falsely rejecting most of the trajectories from that same car as outliers. This experiment suggests that to reliably segment motion data in the presence of outliers, the data should be projected into a space with more than just five dimensions.

3.2 Incomplete Trajectories

In practice, due to occlusions or limitations of the tracker, some features may be missing in some image frames and lead to incomplete trajectories in Y . There are many methods in the computer vision literature for filling in the missing entries of a matrix of motion trajectories [14], [15], [19]. These methods typically assume that the data matrix is low rank. For a matrix with low column rank, the problem of completing missing data can, in fact, be cast as a rank minimization problem:

$$\hat{Y} = \underset{X}{\operatorname{argmin}} \operatorname{rank}(X) \quad \text{subject to} \quad \mathcal{M}(X) = \mathcal{M}(Y), \quad (11)$$

where $\mathcal{M}(\cdot)$ is a mask that matches given entries in Y . As we mentioned earlier, rank minimization is a difficult problem and most of the methods in computer vision mentioned above rely on an iterative alternative minimization scheme. There has been a significant breakthrough in the compressed sensing literature that shows that the above problem can be solved correctly and efficiently by semi-definite programming when the rank is low enough. In fact, a very sharp bound is derived for how many entries are needed for an exact completion of the matrix [2].¹⁰

However, these powerful tools for entry completion run into serious problems when the columns of the data matrix are from multiple subspaces. Data drawn from a union of subspaces can potentially be full rank—the matrix \hat{Y} is often overcomplete. As such, the problem becomes extremely

10. According to this new result, rather surprisingly the percentage of entries needed for an exact completion goes to zero as the dimension goes to infinity, whereas for the iterative schemes, such as Power Factorization [15], the conventional rule of thumb is that one needs about at least 20-30 percent entries for a good chance of success.

TABLE 3

Misclassification and Outlier Detection Rates for LSA and ALC as a Function of the Outlier Percentage (from 0 to 25 Percent) for Three Motion Sequences in Fig. 4

| | 1R2RC [%] | | arm [%] | | cars10 [%] | |
|-----|-----------|-------------|---------|-------------|------------|-------------|
| [%] | LSA | ALC | LSA | ALC | LSA | ALC |
| 0 | 2.40 | 1.09 | 22.08 | 0.00 | 16.84 | 1.34 |
| 7 | 6.91 | 1.29 | 24.17 | 0.13 | 31.97 | 0.40 |
| 15 | 3.09 | 1.31 | 15.38 | 0.06 | 26.43 | 0.19 |
| 25 | 2.69 | 1.16 | 10.25 | 0.04 | 24.59 | 0.17 |

(a)

| | 1R2RC [%] | | arm [%] | | cars10 [%] | |
|-----|-----------|-------|---------|-----|------------|-------|
| [%] | LSA | ALC | LSA | ALC | LSA | ALC |
| 0 | 98.04 | 100 | 77.9 | 100 | 86.87 | 100 |
| 7 | 94.75 | 99.99 | 92.79 | 100 | 96.82 | 99.70 |
| 15 | 98.04 | 99.98 | 91.34 | 100 | 98.84 | 99.81 |
| 25 | 98.20 | 99.97 | 95.56 | 100 | 98.76 | 99.83 |

(b)

(a) Misclassification rates and (b) outlier detection rates.

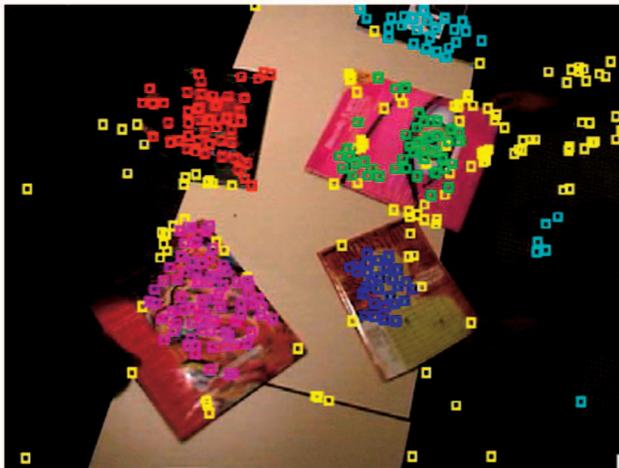
TABLE 4

Information about the Four Motion Sequences in Fig. 5 Containing Real Outlying Trajectories (Numbers of Motions, Samples for Each Group, and Outliers)

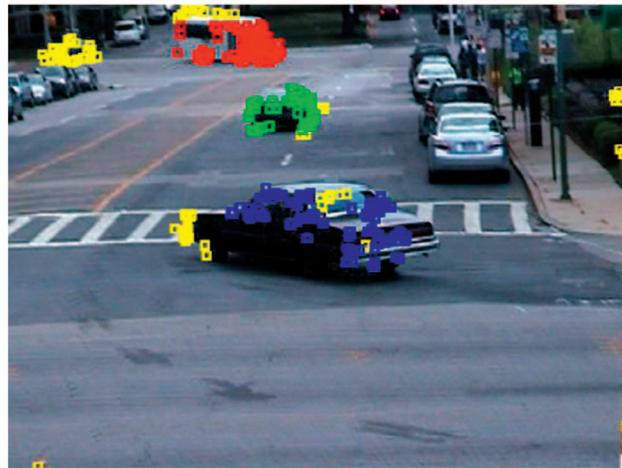
| Sequence | # motions | # samples | # outliers |
|-------------|-----------|--------------------|------------|
| books | 5 | 45, 41, 28, 71, 30 | 127 |
| carsbus3 | 3 | 85, 45, 89 | 89 |
| carsTurning | 4 | 51, 114, 52, 517 | 43 |
| nrbooks3 | 3 | 129, 168, 91 | 35 |

underdetermined as there is, in general, no unique solution for the values of the missing entries as a linear combination of the known entries. However, by harnessing the low-dimensional multiple-subspace structure of the data set, it is actually possible to *complete* these trajectories *prior* to subspace separation.

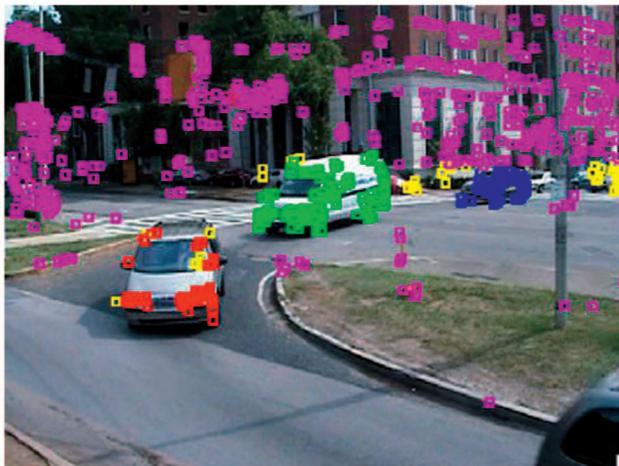
The key observation is that samples drawn from a low-dimensional linear subspace are *self-expressive*, meaning that a sample can be expressed in terms of a few other samples from the same linear subspace. More precisely, if the given sample is $\mathbf{y} \in \mathbb{R}^D$ and $\mathbf{Y} \in \mathbb{R}^{D \times P}$ is the data matrix whose



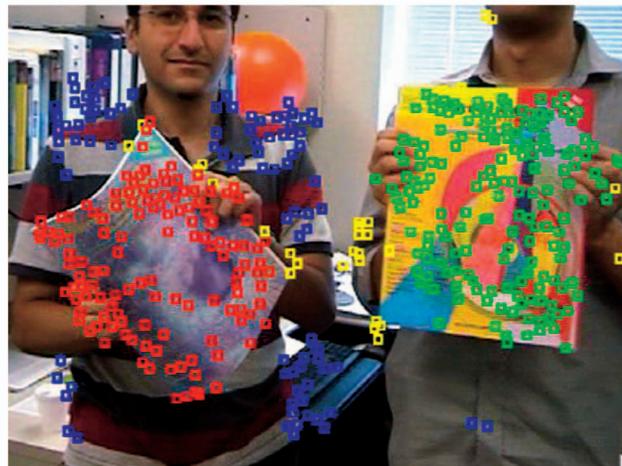
(a)



(b)



(c)



(d)

Fig. 5. Example image frames from four motion sequences containing real outlying trajectories. Feature points from outlying trajectories are labeled as yellow squares. (a) Books sequence. (b) carsbus3 sequence. (c) carsTurning sequence. (d) nrbooks3 sequence.

TABLE 5

Misclassification and Outlier Detection Rates for ALC_5 and ALC_{sp} on Four Motion Sequences with Real Outlying Trajectories

| | books [%] | | carsbus3 [%] | | carsTurning [%] | | nrbooks3 [%] | |
|------------------------|-----------|-------------|--------------|------------|-----------------|-------------|--------------|-------------|
| | ALC_5 | ALC_{sp} | ALC_5 | ALC_{sp} | ALC_5 | ALC_{sp} | ALC_5 | ALC_{sp} |
| Misclassification Rate | 7.89 | 2.05 | 1.62 | 9.74 | 15.44 | 0.26 | 11.11 | 0.47 |
| Outlier Detection Rate | 98.25 | 99.42 | 76.95 | 100 | 75.16 | 97.04 | 27.66 | 98.58 |

columns are all of the *other* samples in the data set, then there exists a coefficient vector $c \in \mathbb{R}^P$ that satisfies

$$y = Yc. \quad (12)$$

As the number of samples P is usually much greater than the dimension of the ambient space D , (12) is a highly underdetermined system of linear equations, and so, in general, c is not unique. In fact, any D vectors in the set that span \mathbb{R}^D can serve as a basis for representing y . However, since y lies in a low-dimensional linear subspace, it can be represented as a linear combination of only a few vectors from the same subspace. Hence, its coefficient vector should have only a few nonzero entries corresponding to vectors from the same subspace. Thus, what we seek is the *sparsest* c :

$$c^* = \underset{c}{\operatorname{argmin}} \|c\|_0 \text{ subject to } y = Yc, \quad (13)$$

where $\|\cdot\|_0$ is the " ℓ^0 norm" equal to the number of nonzero entries in the vector. The sparsest coefficient vector c^* is unique when $\|c^*\|_0 < D/2$. In the general case, ℓ^0 minimization, like MRM, is known to be NP-Hard.¹¹ Fortunately, due to the findings of Donoho [8], it is known that if c^* is sufficiently sparse (i.e., $\|c^*\|_0 \lesssim \lfloor \frac{D+1}{3} \rfloor$), then the ℓ^0 minimization in (13) is equivalent to the following ℓ^1 minimization:

$$c^* = \underset{c}{\operatorname{argmin}} \|c\|_1 \text{ subject to } y = Yc, \quad (14)$$

which is essentially a linear program.

We apply these results to the problem of dealing with incomplete data. Suppose we have a sample $y \in \mathbb{R}^D$ with missing entries $\{y_i\}_{i \in I}$, $I \subset \{1, \dots, D\}$ and a data set $Y \in \mathbb{R}^{D \times P}$ with *no* missing entries. The idea is to use the available entries in y and the corresponding rows in Y to complete the vector. Let $\hat{y} \in \mathbb{R}^{D-|I|}$ and $\hat{Y} \in \mathbb{R}^{(D-|I|) \times P}$ be y and Y with the rows indexed by I removed, respectively. By removing these rows, we are essentially projecting the data onto the $(D - |I|)$ -dimensional subspace orthogonal to $\operatorname{span}(\{e_i : i \in I\})$, where e_i is the i th vector in the canonical basis for \mathbb{R}^D . This is licit because, as long as the dimension of each subspace is strictly less than $d = (D - |I|)$, an arbitrary d -dimensional projection preserves the structural relationships between the subspaces with probability one. Thus, if we solve the linear program¹²

$$c^* = \underset{c}{\operatorname{argmin}} \|c\|_1 \text{ subject to } \hat{y} = \hat{Y}c, \quad (15)$$

11. In fact, when MRM is applied to a set of *diagonal* matrices, it reduces to ℓ^0 minimization.

12. As suggested in [30], one can deal with noisy data by replacing the equality constraint in (15) with $\|\hat{y} - \hat{Y}c\|_2 \leq \epsilon$. Though no longer a linear program, the problem can still be solved efficiently via semidefinite programming.

then the completed vector y^* can be recovered as

$$y^* = Yc^*. \quad (16)$$

3.2.1 Experiments with Simulated Missing Data

We now test the accuracy of our ℓ^1 -based method for entry completion. In each trial, we randomly select a trajectory y_p from the data set for a given sequence, and remove $1 \leq m \leq D - 1 = 2F - 1$ of its entries. We then apply (15) and (16) to recover the missing entries.¹³ In order to simulate many trajectories with missing entries in the data set, we perform five different experiments. In each experiment, we use a subset Y_c containing between 20 and 100 percent of the remaining trajectories to complete y_p .

We also compare the performance of our method with Power Factorization [15], an iterative technique that has been applied to incomplete motion data [27]. Note that the two approaches work under different operating conditions. Our ℓ^1 -based approach uses a set of complete vectors to fill in the missing entries of incomplete vectors, *one* vector at a time. Power Factorization fills in the entries of *all* incomplete vectors simultaneously, but relies on a low-rank representation of the whole matrix. For fair comparison, we embed the trajectories in a data matrix Y , and then randomly remove m entries from y_p , as well as each column of $Y \setminus Y_c$. We then apply Power Factorization to Y to fill in its missing entries, subject to a rank constraint of $r = 4N$, where N is the number of motions in the scene.

Fig. 6 shows the results for 200 trials. For each method and each sequence, we plot the average per-entry error of the recovered trajectory \hat{y}_p with respect to the ground truth versus the percentage of missing entries in each incomplete trajectory. The different colored plots are for the experiments with varying percentage of the data set used for completion. For all motion sequences, our method is able to reconstruct trajectories to within subpixel accuracy even with over 80 percent of the entries missing. The performance of both methods remains consistent even when the entries are completed with small subsets of the remaining data. This suggests that both methods can work well even if a large number of trajectories have missing features. However, as these simulations show, our method clearly outperforms Power Factorization, obtaining lower per-entry error and converging for a larger percentage of missing entries. This is because our method takes advantage of the multiple-subspace structure in the data, while Power Factorization does not.

13. For all of our experiments that use ℓ^1 -minimization, we use the freely available CVX toolbox for MATLAB [13].

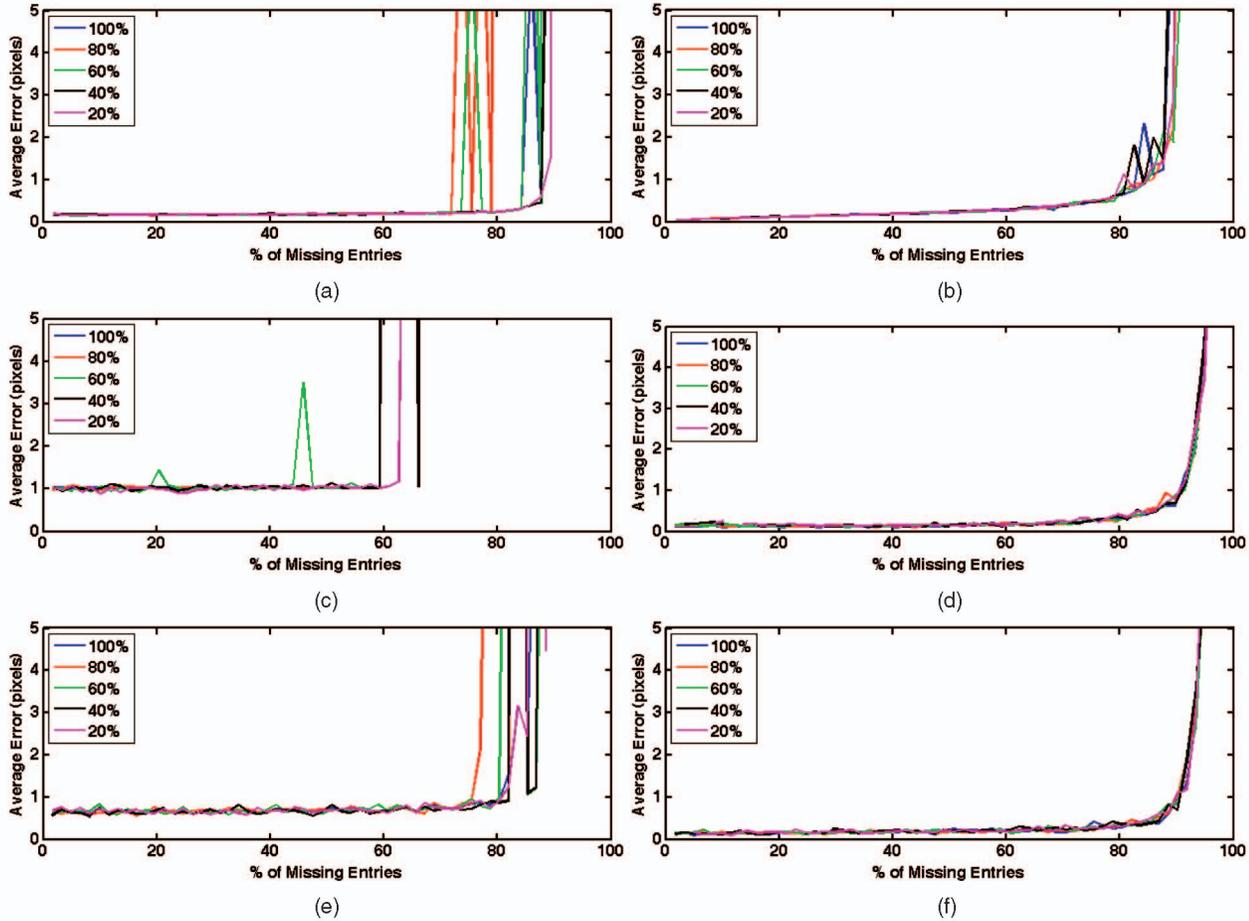


Fig. 6. Errors on the recovered trajectories using our ℓ^1 -based trajectory completion for the sequences: “1R2RC,” “arm,” and “cars10.” The different colored plots are for experiments with varying percentage of the data set used for completion. (a) 1R2RC—PF. (b) 1R2RC— ℓ^1 . (c) arm—PF. (d) arm— ℓ^1 . (e) cars10—PF. (f) cars10— ℓ^1 .

3.2.2 Experiments with Real Missing Data

We now test our robust subspace separation method on real motion sequences with incomplete or corrupted trajectories. We first use the three motion sequences, as shown in Fig. 7. These sequences are taken from [27] and are similar to the checkerboard sequences in Hopkins155. Each sequence contains three different motions and was split into three new sequences containing only trajectories from the first and second groups, first and third groups, and second and third groups, respectively. Thus, in total, we have 12 motion sequences: nine with two motions and three with three motions. For these sequences, between 4 and 35 percent of the entries in the data matrix of trajectories are corrupted. These entries were manually located and labeled.

To see how ℓ^1 -based entry completion affects the quality of segmentation, we remove the entries of trajectories that were marked as corrupted so that we may treat them as missing entries. We apply our ℓ^1 -based entry completion method to this data, and input the completed data into ALC_5 and ALC_{sp} , respectively. For comparison, we also use Power Factorization and Robust PCA [7] to complete the data before segmentation. The misclassification rate for each sequence is listed in Table 6a. Our ℓ^1 -based approach performs competitively with both Power Factorization and Robust PCA. The average performance of $\ell^1 + \text{ALC}_5$ is skewed by its misclassification rate for the “oc2R3RCRT” sequence. This is likely an artifact of the method we use to

choose ε . Note that while both Robust PCA and Power Factorization work well when combined with ALC_5 , their performance degrades significantly when combined with ALC_{sp} . Thus, alternative minimization techniques like Power Factorization and Robust PCA tend to work well only when the dimensionality of projection is small.

We also test our Power Factorization, Robust PCA, and our ℓ^1 -based approach on the four motion sequences in Fig. 5. In this experiment, we remove the outlying trajectories from each sequence and instead use the partially corrupted trajectories. Each trajectory has between 0 and 75 percent of its entries missing. The number and location of missing entries for each trajectory were manually determined. These sequences contain many corrupted trajectories, and so, it is possible that an incomplete trajectory cannot be satisfactorily completed, and will likely be classified as an outlier. Thus, to get a sense of how well the entries of incomplete trajectories are filled in, we compute both the misclassification rate *and* the outlier detection rate for each sequence. The results are listed in Tables 6b and 6c. For all four sequences, our ℓ^1 -based approach in conjunction with ALC_{sp} can effectively deal with incomplete trajectories, treating the fewest as possible as outliers. For the cases where $\text{RPCA} + \text{ALC}$ or $\text{PF} + \text{ALC}$ achieves low misclassification rates, note that the outlier detection rate is also low. This suggests that these iterative methods were unable to recover the missing entries of the incomplete trajectories, and so, such trajectories are incorrectly rejected as outliers.

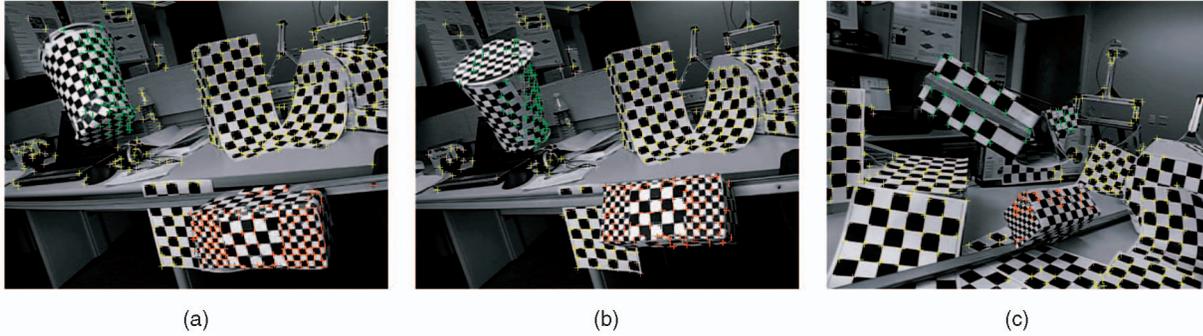


Fig. 7. Example frames from three motion sequences with incomplete or corrupted trajectories. Sequences taken from [27]. (a) oc1R2RC. (b) oc1R2RCT. (c) oc2R3RCRT.

3.3 Corrupted Trajectories

Corrupted entries can be present in a trajectory when the tracker unknowingly loses track of feature points.¹⁴ Such entries are gross errors that could have arbitrary magnitude. One could treat corrupted trajectories as sample outliers.¹⁵ However, in a corrupted trajectory, a portion of the entries still corresponds to a motion in the scene; hence, it seems wasteful to simply discard such information.

Repairing a vector with corrupted entries is a much more difficult problem than the entry completion problem in Section 3.2 because both the number and location of the corrupted entries in the vector are *not known*. Once again, by taking advantage of the low-dimensional multisubspace structure of the data set, we can both detect and repair vectors with corrupted entries *prior* to subspace separation.

A corrupted vector \hat{y} can be modeled as

$$\hat{y} = y + e, \quad (17)$$

where y is the uncorrupted vector and $e \in \mathbb{R}^D$ is a vector that contains all of the gross errors. We assume that there are only a few gross errors, so e will only have a few nonzero entries and thus be sparse.¹⁶ As long as there are enough uncorrupted vectors in the data set, we can express y as a linear combination of the other vectors in the data set, as shown in Section 3.2. If $Y \in \mathbb{R}^{P \times D}$ is a matrix whose columns are the other vectors in the data set, then (17) becomes

$$\hat{y} = Yc + e = [Y \ I] \begin{bmatrix} c \\ e \end{bmatrix} \doteq Bw. \quad (18)$$

We would like both the coefficient vector c and the error vector e to be sparse.¹⁷ If the true c and e are sufficiently sparse, we can simultaneously find the sparsest c and e by solving the linear program:

$$w^* = \underset{w}{\operatorname{argmin}} \|w\|_1 \quad \text{subject to} \quad \hat{y} = Bw. \quad (19)$$

The convex optimization problem in (19) has been successfully applied to robust face recognition in the

14. These kinds of trajectories are called “intrasample outliers” in [7].

15. Indeed, if a data set with some corrupted trajectories is input to ALC, the algorithm will classify those trajectories as outliers as the gross errors will greatly increase the coding length of their ground truth motion group.

16. We realize that, in practice, trajectories may be corrupted by a large number of gross errors. However, it is unlikely that *any* method can repair such trajectories and so it is the best to treat them as sample outliers.

17. The columns of Y should be scaled to have unit ℓ^2 norm to ensure that no vector is preferred in the sparse representation of w .

presence of occlusion [30], and is provably optimal for certain types of corruption [28]. Once w^* is computed, we decompose it into $w^* = [c^* \ e^*]^T$, where $c^* \in \mathbb{R}^P$ is the recovered coefficient vector and $e^* \in \mathbb{R}^D$ is the recovered error vector. The repaired vector y^* is simply

$$y^* = Yc^*. \quad (20)$$

TABLE 6
Comparison of Power Factorization and Robust PCA with Our ℓ^1 -Based Approach for Real Motion Sequences with Incomplete Data

| [%] | PF | | RPCA | | ℓ^1 | |
|---------------|------------------|-------------------|------------------|-------------------|------------------|-------------------|
| | ALC ₅ | ALC _{Sp} | ALC ₅ | ALC _{Sp} | ALC ₅ | ALC _{Sp} |
| oc1R2RC | 0.15 | 6.25 | 0.15 | 2.28 | 0.15 | 0.15 |
| oc1R2RC_g12 | 8.79 | 14.01 | 0.00 | 8.79 | 0.00 | 0.00 |
| oc1R2RC_g13 | 0.00 | 1.02 | 3.67 | 7.76 | 0.00 | 0.00 |
| oc1R2RC_g23 | 0.19 | 1.75 | 0.19 | 1.55 | 0.19 | 0.19 |
| oc1R2RCT | 1.82 | 3.45 | 2.00 | 6.91 | 0.91 | 1.27 |
| oc1R2RCT_g12 | 0.00 | 15.15 | 0.43 | 14.28 | 0.87 | 0.87 |
| oc1R2RCT_g13 | 5.16 | 2.58 | 0.94 | 7.75 | 0.70 | 1.41 |
| oc1R2RCT_g23 | 3.39 | 3.61 | 0.45 | 2.48 | 0.00 | 0.90 |
| oc2R3RCRT | 2.36 | 21.20 | 2.36 | 27.62 | 42.40 | 2.57 |
| oc2R3RCRT_g12 | 0.00 | 34.57 | 0.00 | 41.36 | 0.00 | 1.23 |
| oc2R3RCRT_g13 | 0.51 | 16.62 | 0.51 | 21.74 | 0.51 | 4.35 |
| oc2R3RCRT_g23 | 0.26 | 9.45 | 0.00 | 22.83 | 0.00 | 2.36 |
| Average | 1.89 | 10.81 | 0.89 | 13.78 | 3.81 | 1.28 |
| Median | 0.39 | 7.85 | 0.44 | 8.27 | 0.17 | 1.07 |

(a)

| [%] | PF | | RPCA | | ℓ^1 | |
|-------------|------------------|-------------------|------------------|-------------------|------------------|-------------------|
| | ALC ₅ | ALC _{Sp} | ALC ₅ | ALC _{Sp} | ALC ₅ | ALC _{Sp} |
| books | 0.00 | 1.88 | 0.13 | 0.00 | 0.75 | 0.00 |
| carsbus3 | 0.00 | 0.00 | 15.60 | 0.00 | 0.00 | 0.00 |
| carsTurning | 15.03 | 1.44 | 0.00 | 0.85 | 16.07 | 0.00 |
| nrbooks3 | 10.05 | 0.00 | 5.52 | 0.00 | 5.19 | 0.00 |

(b)

| [%] | PF | | RPCA | | ℓ^1 | |
|-------------|------------------|-------------------|------------------|-------------------|------------------|-------------------|
| | ALC ₅ | ALC _{Sp} | ALC ₅ | ALC _{Sp} | ALC ₅ | ALC _{Sp} |
| books | 57.06 | 63.09 | 40.65 | 44.03 | 76.84 | 91.34 |
| carsbus3 | 74.79 | 93.59 | 77.68 | 75.22 | 74.36 | 100.00 |
| carsTurning | 78.06 | 76.00 | 74.47 | 94.04 | 65.29 | 98.35 |
| nrbooks3 | 52.51 | 76.18 | 37.50 | 75.29 | 62.89 | 87.52 |

(c)

(a) Misclassification rates for the 12 sequences in Fig. 7, (b) misclassification rates for the four sequences in Fig. 5, and (c) outlier detection rates for the four sequences in Fig. 5.

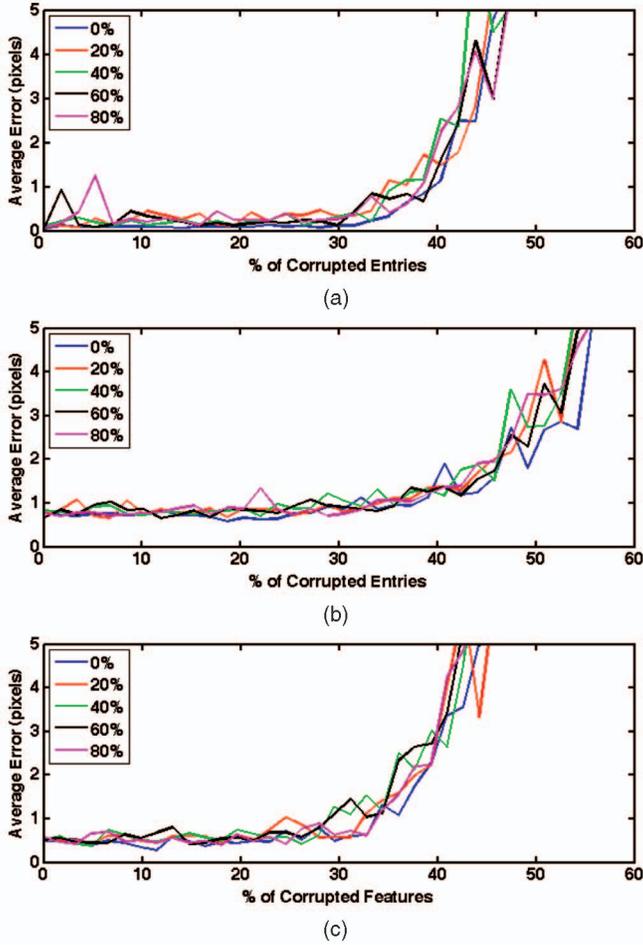


Fig. 8. Results for our ℓ^1 -based detection and repair of corrupted trajectories for the sequences: (a) “1R2RC,” (b) “arm,” and (c) “cars10.” The different colors represent experiments with varying percentage of corrupted trajectories in the data set.

The error vector e^* also provides useful information. The nonzero entries of e^* are precisely the gross errors in \hat{y} .

3.3.1 Experiments with Simulated Corrupted Data

We now test the limits of our ℓ^1 -based method for repair of corrupted trajectories. For each trial in the experiments, we randomly select a trajectory y_p from the given data set, and randomly select and corrupt between 1 and $D - 1 = 2F - 1$ entries in the vector. To corrupt the selected entries, we replace them with random values drawn from a distribution that is uniform in the pixel coordinate space. We then apply (19) and (20) to both detect the locations of corrupted entries, as well as repair them. In each experiment, we run 200 trials and average the errors. We perform five experiments of this type, each with a portion (from 0 to 80 percent) of the remaining data set Y being corrupted in the same way as y_p . The results of these experiments are shown in Fig. 8c. For each sequence, we plot the average per-entry error of the repaired vector with respect to the ground truth versus the percentage of corrupted entries in each vector. The different colors represent experiments with varying portions of corrupted Y . As shown in Fig. 8, this method is able to reconstruct vectors to within subpixel accuracy even with roughly one-third of the entries corrupted. This is in line

TABLE 7
Comparison of Robust PCA with Our ℓ^1 -Based Approach for Real Motion Sequences with Corrupted Data

| [%] | RPCA | | ℓ^1 | |
|---------------|------------------|-------------------|------------------|-------------------|
| | ALC ₅ | ALC _{sp} | ALC ₅ | ALC _{sp} |
| oc1R2RC | 1.68 | 0.15 | 0.15 | 0.15 |
| oc1R2RC_g12 | 0.00 | 2.61 | 0.00 | 0.00 |
| oc1R2RC_g13 | 0.00 | 0.20 | 0.00 | 0.00 |
| oc1R2RC_g23 | 0.19 | 0.00 | 0.19 | 0.00 |
| oc1R2RCT | 8.36 | 1.64 | 0.91 | 1.45 |
| oc1R2RCT_g12 | 0.43 | 0.00 | 0.00 | 0.43 |
| oc1R2RCT_g13 | 0.47 | 1.88 | 0.23 | 1.64 |
| oc1R2RCT_g23 | 0.19 | 0.00 | 0.00 | 1.35 |
| oc2R3RCRT | 42.61 | 7.49 | 41.97 | 9.64 |
| oc2R3RCRT_g12 | 0.62 | 0.62 | 0.62 | 0.00 |
| oc2R3RCRT_g13 | 3.83 | 9.97 | 2.81 | 8.95 |
| oc2R3RCRT_g23 | 6.56 | 9.97 | 2.89 | 12.60 |
| Average | 5.66 | 3.01 | 4.15 | 3.02 |
| Median | 1.15 | 1.61 | 0.21 | 0.89 |

(a)

| [%] | RPCA | | ℓ^1 | |
|-------------|------------------|-------------------|------------------|-------------------|
| | ALC ₅ | ALC _{sp} | ALC ₅ | ALC _{sp} |
| books | 21.58 | 3.08 | 2.78 | 0.00 |
| carsbus3 | 0.00 | 0.00 | 0.00 | 0.00 |
| carsTurning | 14.02 | 0.00 | 17.95 | 3.67 |
| nrbooks3 | 3.25 | 0.00 | 6.45 | 0.00 |

(b)

| [%] | RPCA | | ℓ^1 | |
|-------------|------------------|-------------------|------------------|-------------------|
| | ALC ₅ | ALC _{sp} | ALC ₅ | ALC _{sp} |
| books | 30.16 | 36.60 | 85.19 | 83.95 |
| carsbus3 | 62.13 | 95.32 | 78.83 | 100.00 |
| carsTurning | 95.32 | 76.66 | 72.82 | 97.72 |
| nrbooks3 | 6.20 | 65.29 | 66.73 | 82.26 |

(c)

(a) Misclassification rates for the 12 real sequences in Fig. 7, (b) misclassification rates for the four real sequences in Fig. 5, and (c) outlier detection rates for the four real sequences in Fig. 5

with the bound $\|e^*\|_0 < \lfloor \frac{D+1}{3} \rfloor$ given by [8]. We also see that the performance remains consistent even if 80 percent of the entire data set is corrupted.

3.3.2 Experiments with Real Corrupted Data

We test our ability to repair corrupted trajectories and observe the effects of the repair on segmentation. We apply our ℓ^1 -based repair and detection method to the raw motion sequences in Fig. 7, and then input the repaired data to ALC₅ and ALC_{sp}, respectively. For comparison, we also use Robust PCA to complete the data before segmentation. The misclassification rate for each sequence is listed in Table 7a. Both Robust PCA and our ℓ^1 -based approach can repair corrupted trajectories to achieve reasonable segmentations.

We also test our ℓ^1 -based approach for error corrections on the four motion sequences in Fig. 5. In this experiment, each trajectory has between 0 and 25 percent of its entries corrupted. The misclassification and outlier detection rates for each sequence are listed in Table 7b. For these more realistic sequences, we see that our ℓ^1 -based approach can still effectively deal with corrupted trajectories, treating the

fewest possible as outliers. For the cases where RPCA+ALC achieves good misclassification rates, note that the outlier detection is also low, meaning that Robust PCA was unable to detect and correct the errors in the corrupted trajectories. Finally, in both of these experiments, we note that both methods tend to work better when combined with ALC_{sp} .

4 CONCLUSIONS AND FUTURE WORK

In this paper, we have investigated the problem of motion segmentation from the perspective of robust subspace separation. We have shown that the key for correct segmentation, data completion, and error correction is to correctly harness the intrinsic low-dimensional, sparse structures within such data. This has made the proper choice of measures for sparsity and compactness the central issue. We have shown that in our context, both the (lossy) coding length and 1-norm are good surrogates for the matrix rank and vector sparsity, respectively. Not only is the use of these measures theoretically well founded, but we have also demonstrated with extensive simulations and experiments that they indeed lead to algorithms with superior performance for segmenting motion trajectories despite outliers, incomplete data, and random errors. The proposed techniques and algorithms are, in fact, generic to subspace separation, and can conceivably be used in other application domains with little modification.

This paper provides strong, encouraging empirical evidence for people to work on many exciting open theoretical problems. We have explored several schemes for improving both the speed and convergence of the coding-length-based agglomerative algorithm. In the algorithm, the coding length is used as a "distance" measure between pairs of subsets. It is worth investigating if such a measure exhibits *locality-sensitive hashing* properties [6] as other norms so that more principled speedup algorithms can be derived.

We have seen that, typically, the agglomerative algorithm converges to the correct motion segmentation for a wide range of choice of ϵ . There is still a lack of proof for under what conditions the agglomerative algorithm is expected to converge to the segmentation with *globally* minimum coding length. Experiments in this paper and simulations in [20] seem to indicate that there is a phase transition between convergence and divergence of the agglomerative algorithm, similar to the phase transition for ℓ^1 - ℓ^0 equivalence observed in [9]. This remains an open problem that we will investigate in the future.

Although the problem of completing a low-rank matrix has recently been solved [2], the problem of completing a matrix with columns from *multiple subspaces* remains a widely open problem. In this paper, we have seen surprisingly good performance with the ℓ^1 -minimization. However, there is no proof yet whether this is the best one can do for this problem nor is there a clear characterization for the amount of entries needed.

Empirically, we have observed that the sparse coefficients c computed in our method are indicative of the membership of motion trajectories involved. This suggests that the sparse coefficients can be used as a measure of *similarity* for the trajectories' membership. Hence, one could potentially

use graphical cuts or spectral clustering methods for segmenting the trajectories. It would be interesting to find out if such an approach could lead to more competitive clustering results than other similarity measures, such as the Local Subspace Affinity [31], or results even better than the methods proposed in this paper.

ACKNOWLEDGMENTS

This work is partially supported by US National Science Foundation (NSF) grants CRS-EHS-0509151, NSF CCF-TF-0514955, ONR YIP N00014-05-1-0633, NSF IIS 07-03756, NSF CAREER IIS-0447739, NSF EHS-0509101, and ONR N00014-05-10836, and by contract JHU APL-934652.

REFERENCES

- [1] A. Barron, J. Rissanen, and B. Yu, "The Minimum Description Length Principle in Coding and Modeling," *IEEE Trans. Information Theory*, vol. 44, no. 6, pp. 2743-2760, Oct. 1998.
- [2] E. Candes and B. Recht, "Exact Matrix Completion via Convex Optimization," *Foundations of Computational Math.*, 2009.
- [3] E. Candes, M. Rudelson, R. Vershynin, and T. Tao, "Error Correction via Linear Programming," *Proc. IEEE Symp. Foundations of Computer Science*, pp. 295-308, 2005.
- [4] E. Candes and T. Tao, "Decoding by Linear Programming," *IEEE Trans. Information Theory*, vol. 51, no. 12, pp. 4203-4215, Dec. 2005.
- [5] J. Costeira and T. Kanade, "A Multibody Factorization Method for Independently Moving Objects," *Int'l J. Computer Vision*, vol. 29, no. 3, pp. 159-179, 1998.
- [6] M. Datar, N. Immorlica, P. Indyk, and V. Mirrokni, "Locality-Sensitive Hashing Scheme Based of p -Stable Distributions," *Proc. ACM Symp. Computational Geometry*, 2004.
- [7] F. De la Torre and M.J. Black, "Robust Principal Component Analysis for Computer Vision," *Proc. Int'l Conf. Computer Vision*, pp. 362-369, 2001.
- [8] D.L. Donoho, "For Most Large Underdetermined Systems of Linear Equations the Minimal ℓ^1 -Norm Solution Is Also the Sparsest Solution," *Comm. Pure and Applied Math.*, vol. 59, no. 6, pp. 797-829, Mar. 2006.
- [9] D.L. Donoho and J. Tanner, "Counting Faces of Randomly Projected Polytopes When the Projection Radically Lowers Dimension," *J. Am. Math. Soc.*, vol. 22, no. 1, pp. 1-53, 2009.
- [10] Z. Fan, J. Zhou, and Y. Wu, "Multibody Grouping by Inference of Multiple Subspaces from High Dimensional Data Using Oriented-Frames," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 28, no. 1, pp. 90-105, Jan. 2006.
- [11] M. Fazel, H. Hindi, and S. Boyd, "Log-Det Heuristic for Matrix Rank Minimization with Applications to Hankel and Euclidean Distance Matrices," *Proc. Am. Control Conf.*, pp. 2156-2162, June 2003.
- [12] C. Gear, "Multibody Grouping from Motion Images," *Int'l J. Computer Vision*, vol. 29, no. 2, pp. 133-150, 1998.
- [13] M. Grant and S. Boyd, "CVX: MATLAB Software for Disciplined Convex Programming [Web Page and Software]," <http://www.stanford.edu/~boyd/cvx/>, Nov. 2007.
- [14] A. Gruber and Y. Weiss, "Multibody Factorization with Uncertainty and Missing Data Using the EM Algorithm," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, vol. 1, pp. 769-775, 2004.
- [15] R. Hartley and F. Schaffalitzky, "PowerFactorization: An Approach to Affine Reconstruction with Missing and Uncertain Data," *Proc. Australia-Japan Advanced Workshop Computer Vision*, 2003.
- [16] N. Ichimura, "Motion Segmentation Based on Factorization and Discriminant Criterion," *Proc. IEEE Int'l Conf. Computer Vision*, pp. 600-605, 1999.
- [17] K. Kanatani, "Motion Segmentation by Subspace Separation and Model Selection," *Proc. IEEE Int'l Conf. Computer Vision*, vol. 2, pp. 586-591, 2001.
- [18] K. Kanatani and Y. Sugaya, "Multi-Stage Optimization for Multi-Body Motion Segmentation," *Proc. Australia-Japan Advanced Workshop Computer Vision*, 2003.

- [19] Q. Ke and T. Kanade, "Robust ℓ^1 -Norm Factorization in the Presence of Outliers and Missing Data by Alternative Convex Programming," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 739-746, 2005.
- [20] Y. Ma, H. Derksen, W. Hong, and J. Wright, "Segmentation of Multivariate Mixed Data via Lossy Coding and Compression," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 29, no. 9, pp. 1546-1562, Sept. 2007.
- [21] B. Recht, M. Fazel, and P.A. Parillo, "Guaranteed Minimum-Rank Solutions of Linear Matrix Equations via Nuclear Norm Minimization," *SIAM Rev.*, 2007.
- [22] K. Schindler, D. Suter, and H. Wang, "A Model-Selection Framework for Multibody Structure-and-Motion of Image Sequences," *Int'l J. Computer Vision*, vol. 79, no. 2, pp. 159-177, Aug. 2008.
- [23] P. Torr, "Geometric Motion Segmentation and Model Selection," *Philosophical Trans. Royal Soc. of London*, vol. 356, pp. 1321-1340, 1998.
- [24] R. Tron and R. Vidal, "A Benchmark for the Comparison of 3D Motion Segmentation Algorithms," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 1-8, 2007.
- [25] L. Vandenberghe and S. Boyd, "Semidefinite Programming," *SIAM Rev.*, vol. 38, no. 1, pp. 49-95, Mar. 1996.
- [26] R. Vidal, Y. Ma, and S. Sastry, "Generalized Principal Component Analysis (GPCA)," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 27, no. 12, pp. 1945-1959, Dec. 2005.
- [27] R. Vidal, R. Tron, and R. Hartley, "Multiframe Motion Segmentation with Missing Data Using PowerFactorization and GPCA," *Int'l J. Computer Vision*, vol. 79, no. 1, pp. 85-105, 2007.
- [28] J. Wright and Y. Ma, "Dense Error Correction via ℓ^1 Minimization," *IEEE Trans. Information Theory*, 2008.
- [29] J. Wright, Y. Ma, Y. Tao, Z. Lin, and H.-Y. Shum, "Classification via Minimum Incremental Coding Length," *SIAM J. Imaging Sciences*, vol. 2, no. 2, pp. 367-395, 2009.
- [30] J. Wright, A.Y. Yang, A. Ganesh, S.S. Sastry, and Y. Ma, "Robust Face Recognition via Sparse Representation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 31, no. 2, pp. 210-227, Feb. 2009.
- [31] J. Yan and M. Pollefeys, "A General Framework for Motion Segmentation: Independent, Articulated, Rigid, Non-Rigid, Degenerate and Non-Degenerate," *Proc. European Conf. Computer Vision*, pp. 94-106, 2006.
- [32] A.Y. Yang, S. Rao, and Y. Ma, "Robust Statistical Estimation and Segmentation of Multiple Subspaces," *Proc. IEEE Computer Vision and Pattern Recognition Workshop 25 Years of RANSAC*, 2006.
- [33] L. Zelnik-Manor and M. Irani, "Degeneracies, Dependencies and Their Implications in Multi-Body and Multi-Sequence Factorization," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, vol. 2, pp. 287-293, 2003.



Shankar Rao received the BS degree in electrical engineering and computer science from the University of California, Berkeley, in 2001, and the MS degree in electrical and computer engineering, the MS degree in applied mathematics, and the PhD degree in electrical and computer engineering from the University of Illinois at Urbana-Champaign in 2004, 2006, and 2009, respectively. He is currently a research staff member at HRL Laboratories, LLC. His

research interests include lossy coding-based clustering, robust subspace segmentation, manifold learning, and segmentation of images, motion, and videos. He is a member of the IEEE.



Roberto Tron received the BSc and MSc degrees (highest honors) in telecommunication engineering from the Politecnico di Torino in 2004 and 2007, respectively, the Diplôme d'Ingénieur from the Eurecom Institute, and the DEA degree from the Université de Nice Sophia-Antipolis in 2006. He is currently working toward the PhD degree in the Department of Electrical and Computer Engineering at Johns Hopkins University. His research interests include motion segmentation and distributed algorithms on camera sensor networks. He is a student member of the IEEE.



René Vidal received the BSc degree (highest honors) in electrical engineering from the Universidad Católica de Chile in 1997, and the MSc and PhD degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 2000 and 2003, respectively. In 2004, he joined the Department of Biomedical Engineering and the Center for Imaging Science at Johns Hopkins University as an assistant professor. He was a coeditor (with Anders Heyden and Yi Ma) of the book *Dynamical Vision* and has coauthored more than 100 articles in biomedical imaging, computer vision, machine learning, hybrid systems, robotics, and vision-based control. He is an associate editor of the *Journal of Mathematical Imaging and Vision*, and a member of the program committees of all of the major computer vision conferences. He is the recipient of the 2005 US National Science Foundation (NSF) CAREER Award, the 2004 Best Paper Award Honorable Mention at the European Conference on Computer Vision, the 2004 Sakrison Memorial Prize, the 2003 Eli Jury Award, and the 1997 Award of the School of Engineering of the Universidad Católica de Chile to the best graduating student of the school. He is a member of the IEEE.



Yi Ma received the two bachelor's degrees in automation and applied mathematics from Tsinghua University, Beijing, China, in 1995, and the MS degree in electrical engineering and computer science in 1997, the MA degree in mathematics in 2000, and the PhD degree in electrical engineering and computer science in 2000 from the University of California, Berkeley. Since 2000, he has been on the faculty of the Electrical and Computer Engineering Department of the University of Illinois at Urbana-Champaign, where he now holds the rank of associate professor. Currently, he is also the research manager for the Visual Computing Group of Microsoft Research Asia in Beijing. His main research areas are in systems theory and computer vision. He was the recipient of the David Marr Best Paper Prize at the International Conference on Computer Vision in 1999 and Honorable Mention for the Longuet-Higgins Best Paper Award at the European Conference on Computer Vision in 2004. He received the CAREER Award from the US National Science Foundation in 2004 and the Young Investigator Program Award from the US Office of Naval Research in 2005. He is an associate editor of the *IEEE Transactions on Pattern Analysis and Machine Intelligence*. He is a senior member of the IEEE and a member of the ACM.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.