

An Online Algorithm for Separating Sparse and Low-Dimensional Signal Sequences From Their Sum

Han Guo, Chenlu Qiu, and Namrata Vaswani

Abstract—This paper designs and extensively evaluates an online algorithm, called practical recursive projected compressive sensing (Prac-ReProCS), for recovering a time sequence of sparse vectors S_t and a time sequence of dense vectors L_t from their sum, $M_t := S_t + L_t$, when the L_t 's lie in a slowly changing low-dimensional subspace of the full space. A key application where this problem occurs is in real-time video layering where the goal is to separate a video sequence into a slowly changing background sequence and a sparse foreground sequence that consists of one or more moving regions/objects on-the-fly. Prac-ReProCS is a practical modification of its theoretical counterpart which was analyzed in our recent work. Extension to the undersampled case is also developed. Extensive experimental comparisons demonstrating the advantage of the approach for both simulated and real videos, over existing batch and recursive methods, are shown.

Index Terms—Online robust PCA, recursive sparse recovery, large but structured noise, compressed sensing.

I. INTRODUCTION

THIS paper designs and evaluates a practical algorithm for recovering a time sequence of sparse vectors S_t and a time sequence of dense vectors L_t from their sum, $M_t := S_t + L_t$, when the L_t 's lie in a slowly changing low-dimensional subspace of \mathbf{R}^n . The magnitude of the entries of L_t could be larger, roughly equal or smaller than that of the nonzero entries of S_t . The extension to the undersampled case, $M_t := AS_t + BL_t$, is also developed. The above problem can be interpreted as one of online/recursive sparse recovery from potentially large but structured noise. In this case, S_t is the quantity of interest and L_t is the potentially large but structured low-dimensional noise. Alternatively it can be posed as a recursive/online robust principal components analysis (PCA) problem. In this case L_t , or in fact, the subspace in which it lies, is the quantity of interest while S_t is the outlier.

A key application where the above problem occurs is in video layering where the goal is to separate a slowly changing background from moving foreground objects/regions [4], [5].

Manuscript received September 03, 2013; revised February 07, 2014; accepted May 20, 2014. Date of publication June 18, 2014; date of current version July 18, 2014. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Martin Haardt. This work was partially supported by NSF grants CCF-0917015, CCF-1117125 and IIS-1117509. A portion of this work was presented at Allerton 2010, Allerton 2011, and ICASSP 2014.

The authors are with the Department of Electrical and Computer Engineering, Iowa State University, Ames, IA 50010 USA (e-mail: hanguo@iastate.edu; chenlu@iastate.edu; namrata@iastate.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2014.2331612

The foreground layer, e.g. moving people/objects, is of interest in applications such as automatic video surveillance, tracking moving objects, or video conferencing. The background sequence is of interest in applications such as background editing (video editing applications). In most static camera videos, the background images do not change much over time and hence the background image sequence is well modeled as lying in a fixed or slowly-changing low-dimensional subspace of \mathbf{R}^n [5], [6]. Moreover the changes are typically global, e.g. due to lighting variations, and hence modeling it as a dense image sequence is valid too [5]. The foreground layer usually consists of one or more moving objects/persons/regions that move in a correlated fashion, i.e. it is a sparse image sequence that often changes in a correlated fashion over time. Other applications where the above problem occurs include solving the video layering problem from compressive video measurements, e.g. those acquired using a single-pixel camera; online detection of brain activation patterns from full or undersampled functional MRI (fMRI) sequences (the “active” part of the brain forms the sparse image, while the rest of the brain which does not change much over time forms the low-dimensional part); or sensor networks based detection and tracking of abnormal events such as forest fires or oil spills. The single pixel imaging and undersampled fMRI applications are examples of the compressive case, $M_t = AS_t + BL_t$ with $B = A$.

Related Work: Most high dimensional data often approximately lie in a lower dimensional subspace. Principal components’ analysis (PCA) is a widely used dimension reduction technique that finds a small number of orthogonal basis vectors (principal components), along which most of the variability of the dataset lies. For a given dimension, r , PCA finds the r -dimensional subspace that minimizes the mean squared error between data vectors and their projections into this subspace [7]. It is well known that PCA is very sensitive to outliers. Computing the PCs in the presence of outliers is called robust PCA. Solving the robust PCA problem recursively as more data comes in is referred to as online or recursive robust PCA. “Outlier” is a loosely defined term that usually refers to any corruption that is not small compared to the true signal (or data vector) and that occurs only occasionally. As suggested in [8], an outlier can be nicely modeled as a sparse vector.

In the last few decades, there has been a large amount of work on robust PCA, e.g. [4], [9]–[12], and recursive robust PCA e.g. [13]–[15]. In most of these works, either the locations of the missing/corrupted data points are assumed known [13] (not a practical assumption); or they first detect the corrupted data points and then replace their values using nearby values [14]; or weight each data point in proportion to its reliability (thus soft-detecting and down-weighting the likely outliers) [4], [15];

or just remove the entire outlier vector [11], [12]. Detecting or soft-detecting outliers (S_t) as in [4], [14], [15] is easy when the outlier magnitude is large, but not when it is of the same order or smaller than that of the L_t 's.

In a series of recent works [5], [16], a new and elegant solution to robust PCA called Principal Components' Pursuit (PCP) has been proposed, that does not require a two step outlier location detection/correction process and also does not throw out the entire vector. It redefines batch robust PCA as a problem of separating a low rank matrix, $\mathcal{L}_t := [L_1, \dots, L_t]$, from a sparse matrix, $\mathcal{S}_t := [S_1, \dots, S_t]$, using the measurement matrix, $\mathcal{M}_t := [M_1, \dots, M_t] = \mathcal{L}_t + \mathcal{S}_t$. Other recent works that also study batch algorithms for recovering a sparse S_t and a low-rank \mathcal{L}_t from $\mathcal{M}_t := \mathcal{L}_t + \mathcal{S}_t$ or from undersampled measurements include [17]–[26]. It was shown in [5] that by solving PCP:

$$\min_{\mathcal{L}, \mathcal{S}} \|\mathcal{L}\|_* + \lambda \|\mathcal{S}\|_1 \text{ subject to } \mathcal{L} + \mathcal{S} = \mathcal{M}_t \quad (1)$$

one can recover \mathcal{L}_t and \mathcal{S}_t exactly, provided that (a) \mathcal{L}_t is “dense”; (b) any element of the matrix \mathcal{S}_t is nonzero w.p. ρ , and zero w.p. $1 - \rho$, independent of all others (in particular, this means that the support sets of the different S_t 's are independent over time); and (c) the rank of \mathcal{L}_t and the support size of \mathcal{S}_t are small enough. Here $\|A\|_*$ is the nuclear norm of a matrix A (sum of singular values of A) while $\|A\|_1$ is the ℓ_1 norm of A seen as a long vector.

Notice that most applications described above require an online solution. A batch solution would need a long delay; and would also be much slower and more memory-intensive than a recursive solution. Moreover, the assumption that the foreground support is independent over time is not usually valid. To address these issues, in the conference versions of this work [1], [2], we introduced a novel recursive solution called Recursive Projected Compressive Sensing (ReProCS). In recent work [27]–[29], we have obtained performance guarantees for ReProCS. Under mild assumptions (denseness, slow enough subspace change of L_t and “some” support change at least every h frames of S_t), we showed that, with high probability (w.h.p.), ReProCS can exactly recover the support set of S_t at all times; and the reconstruction errors of both S_t and L_t are upper bounded by a time invariant and small value. The work of [27], [28] contains a partial result while [29] is a complete correctness result.

Contributions: The contributions of this work are as follows. (1) We design a practically usable modification of the ReProCS algorithm studied in [27]–[29]. By “practically usable”, we mean that (a) it requires fewer parameters and we develop simple heuristics to set these parameters without any model knowledge; (b) it exploits practically motivated assumptions and we demonstrate that these assumptions are valid for real video data. While denseness and gradual support change are also used in earlier works—[5], [16] and [30] respectively—*slow subspace change* is the key new (and valid) assumption introduced in ReProCS. (2) We show via extensive simulation and real video experiments that practical-ReProCS is more robust to correlated support change of S_t than PCP and other existing work. Also, it is also able to recover small magnitude sparse vectors significantly better than other existing recursive as well as batch algorithms. (3) We also develop a

compressive practical-ReProCS algorithm that can recover S_t from $M_t := AS_t + BL_t$. In this case A and B can be fat, square or tall.

More Related Work: Other very recent work on recursive/online robust PCA includes [31]–[35].

Some other related work includes work that uses structured sparsity models, e.g. [36]. For our problem, if it is known that the sparse vector consists of one or a few connected regions, these ideas could be incorporated into our algorithm as well. On the other hand, the advantage of the current approach that only uses sparsity is that it works both for the case of a few connected regions as well as for the case of multiple small sized moving objects, e.g. see the airport video results at http://www.ece.iastate.edu/~chenlu/ReProCS/Video_ReProCS.htm.

Paper Organization: We give the precise problem definition and assumptions in Section II. The practical ReProCS algorithm is developed in Section III. The algorithm for the compressive measurements' case is developed in Section IV. In Section V, we demonstrate using real videos that the key assumptions used by our algorithm are true in practice. Experimental comparisons on simulated and real data are shown in Section VI. Conclusions and future work are discussed in Section VII.

A. Notation

For a set $T \subseteq \{1, 2, \dots, n\}$, we use $|T|$ to denote its cardinality; and we use T^c to denote its complement, i.e. $T^c := \{i \in \{1, 2, \dots, n\} : i \notin T\}$. The symbols \cup, \cap, \setminus denote set union set intersection and set difference respectively (recall $T_1 \setminus T_2 := T_1 \cap T_2^c$). For a vector v , v_i denotes the i th entry of v and v_T denotes a vector consisting of the entries of v indexed by T . We use $\|v\|_p$ to denote the ℓ_p norm of v . The support of v , $\text{supp}(v)$, is the set of indices at which v is nonzero, $\text{supp}(v) := \{i : v_i \neq 0\}$. We say that v is s -sparse if $|\text{supp}(v)| \leq s$.

For a matrix B , B' denotes its transpose, and B^\dagger denotes its pseudo-inverse. For a matrix with linearly independent columns, $B^\dagger = (B'B)^{-1}B'$. The notation $[\]$ denotes an empty matrix. We use I to denote an identity matrix. For an $m \times n$ matrix B and an index set $T \subseteq \{1, 2, \dots, n\}$, B_T is the sub-matrix of B containing columns with indices in the set T . Notice that $B_T = BI_T$. We use $B \setminus B_T$ to denote B_{T^c} . Given another matrix B_2 of size $m \times n_2$, $[B \ B_2]$ constructs a new matrix by concatenating matrices B and B_2 in horizontal direction. Thus, $[(B \setminus B_T) \ B_2] = [B_{T^c} \ B_2]$. We use the notation $B \stackrel{SVD}{=} U\Sigma V'$ to denote the singular value decomposition (SVD) of B with the diagonal entries of Σ being arranged in non-decreasing order.

The interval notation $[t_1, t_2] := \{t_1, t_1 + 1, \dots, t_2\}$ and similarly the matrix $[L_{t_1}, \dots, L_{t_2}] := [L_{t_1}, L_{t_1+1}, \dots, L_{t_2}]$

Definition 1.1: The s -restricted isometry constant (RIC) [37], δ_s , for an $n \times m$ matrix Ψ is the smallest real number satisfying $(1 - \delta_s)\|x\|_2^2 \leq \|\Psi_T x\|_2^2 \leq (1 + \delta_s)\|x\|_2^2$ for all sets T with $|T| \leq s$ and all real vectors x of length $|T|$.

Definition 1.2: For a matrix M ,

- $\text{range}(M)$ denotes the subspace spanned by the columns of M .
- M is a *basis matrix* if $M'M = I$.
- The notation $Q = \text{basis}(\text{range}(M))$, or $Q = \text{basis}(M)$ for short, means that Q is a *basis matrix* for $\text{range}(M)$ i.e. Q satisfies $Q'Q = I$ and $\text{range}(Q) = \text{range}(M)$.

Definition 1.3:

- The $b\%$ left singular values' set of a matrix M is the smallest set of indices of its singular values that contains at least $b\%$ of the total singular values' energy. In other words, if $M \stackrel{SVD}{=} U\Sigma V'$, it is the smallest set T such that $\sum_{i \in T} (\Sigma)_{i,i}^2 \geq \frac{b}{100} \sum_{i=1}^n (\Sigma)_{i,i}^2$.
- The corresponding matrix of left singular vectors, U_T , is referred to as the $b\%$ left singular vectors' matrix.
- The notation $[Q, \Sigma] = \text{approx-basis}(M, b\%)$ means that Q is the $b\%$ left singular vectors' matrix for M and Σ is the diagonal matrix with diagonal entries equal to the $b\%$ left singular values' set.
- The notation $Q = \text{approx-basis}(M, r)$ means that Q contains the left singular vectors of M corresponding to its r largest singular values. This also sometimes referred to as: Q contains the r top singular vectors of M .

II. PROBLEM DEFINITION AND ASSUMPTIONS

The measurement vector at time t , M_t , is an n dimensional vector which can be decomposed as

$$M_t := S_t + L_t. \quad (2)$$

Let T_t denote the support set of S_t , i.e.,

$$T_t := \text{supp}(S_t) = \{i : (S_t)_i \neq 0\}.$$

We assume that S_t and L_t satisfy the assumptions given below in the next three subsections. Suppose that an initial training sequence which does not contain the sparse components is available, i.e. we are given $\mathcal{M}_{\text{train}} = [M_t; 1 \leq t \leq t_{\text{train}}]$ with $M_t = L_t$. This is used to get an initial estimate of the subspace in which the L_t 's lie.¹ At each $t > t_{\text{train}}$, the goal is to recursively estimate S_t and L_t and the subspace in which L_t lies. By "recursively" we mean: use $\hat{S}_{t-1}, \hat{L}_{t-1}$ and the previous subspace estimate to estimate S_t and L_t .

The magnitude of the entries of L_t may be small, of the same order, or large compared to that of the nonzero entries of S_t . In applications where S_t is the signal of interest, the case when $\|L_t\|_2$ is of the same order or larger than $\|S_t\|_2$ is the difficult case.

A key application where the above problem occurs is in separating a video sequence into background and foreground layers. Let Im_t denote the image at time t , F_t denote the foreground image at t and B_t the background image at t , all arranged as 1-D vectors. Then, the image sequence satisfies

$$(\text{Im}_t)_i = \begin{cases} (F_t)_i & \text{if } i \in \text{supp}(F_t) \\ (B_t)_i & \text{if } i \notin \text{supp}(F_t). \end{cases} \quad (3)$$

In fMRI, F_t is the sparse active region image while B_t is the background brain image. In both cases, it is fair to assume that an initial background-only training sequence is available. For video this means there are no moving objects/regions in the foreground. For fMRI, this means some frames are captured without providing any stimulus to the subject.

¹If an initial sequence without S_t 's is not available, one can use a batch robust PCA algorithm to get the initial subspace estimate as long as the initial sequence satisfies its required assumptions.

Let μ denote the empirical mean of the training background images. If we let $L_t := B_t - \mu, M_t := \text{Im}_t - \mu, T_t := \text{supp}(F_t)$, and

$$(S_t)_{T_t} := (F_t - B_t)_{T_t}, (S_t)_{T_t^c} := 0,$$

then, clearly, $M_t = S_t + L_t$. Once we get the estimates \hat{L}_t, \hat{S}_t , we can also recover the foreground and background as

$$\hat{B}_t = \hat{L}_t + \mu, \hat{T}_t = \text{supp}(\hat{S}_t), (\hat{F}_t)_{\hat{T}_t} = (\text{Im}_t)_{\hat{T}_t}, (\hat{F}_t)_{\hat{T}_t^c} = 0.$$

A. Slowly Changing Low-Dimensional Subspace Change

We assume that for τ large enough, any τ length subsequence of the L_t 's lies in a subspace of \mathbf{R}^n of dimension less than $\min(\tau, n)$, and usually much less than $\min(\tau, n)$. In other words, for τ large enough, $\max_t \text{rank}([L_{t-\tau+1}, \dots, L_t]) \ll \min(\tau, n)$. Also, this subspace is either fixed or changes slowly over time.

One way to model this is as follows [27]. Let $L_t = P_t a_t$ where P_t is an $n \times r_t$ basis matrix with $r_t \ll n$ that is piecewise constant with time, i.e. $P_t = P_{(j)}$ for all $t \in [t_j, t_{j+1})$ and $P_{(j)}$ changes as

$$P_{(j)} = [(P_{(j-1)} R_j \setminus P_{(j), \text{old}}), P_{(j), \text{new}}]$$

where $P_{(j), \text{new}}$ and $P_{(j), \text{old}}$ are basis matrices of size $n \times c_{j, \text{new}}$ and $n \times c_{j, \text{old}}$ respectively with $P'_{(j), \text{new}} P_{(j-1)} = 0$ and R_j is a rotation matrix. Moreover, (a) $0 \leq \sum_{i=1}^j (c_{i, \text{new}} - c_{i, \text{old}}) \leq c_{\text{dif}}$; (b) $0 \leq c_{j, \text{new}} \leq c_{\text{max}} < r_0$; (c) $(t_{j+1} - t_j) \gg r_0 + c_{\text{dif}}$; and (d) there are a total of J change times with $J \ll (n - r_0 - c_{\text{dif}})/c_{\text{max}}$.

Clearly, (a) implies that $r_t \leq r_0 + c_{\text{dif}} := r_{\text{max}}$ and (d) implies that $r_{\text{max}} + J c_{\text{max}} \ll n$. This, along with (b) and (c), helps to ensure that for any $\tau > r_{\text{max}} + c_{\text{max}}$, $r^{t, \tau} := \max_t \text{rank}([L_{t-\tau+1}, \dots, L_t]) < \min(\tau, n)$, and for $\tau \gg r_{\text{max}} + c_{\text{max}}$, $r^{t, \tau} \ll \min(\tau, n)^2$.

By slow subspace change, we mean that: for $t \in [t_j, t_{j+1})$, $\|(I - P_{(j-1)} P'_{(j-1)}) L_t\|_2$ is initially small and increases gradually. In particular, we assume that, for $t \in [t_j, t_j + \alpha)$,

$$\left\| (I - P_{(j-1)} P'_{(j-1)}) L_t \right\|_2 \leq \gamma_{\text{new}} \ll \min(\|L_t\|_2, \|S_t\|_2)$$

and increases gradually after $t_j + \alpha$. One model for "increases gradually" is as given in [27, Sec III-B]. Nothing in this paper requires the specific model and hence we do not repeat it here.

²To address a reviewer comment, we explain this in detail here. Notice first that (c) implies that $(t_{j+1} - t_j) \gg r_{\text{max}}$. Also, (b) implies that $\text{rank}([L_{t_j}, \dots, L_{t_{j+k}-1}]) \leq r_{\text{max}} + (k-1)c_{\text{max}}$. First consider the case when both $t - \tau + 1$ and t lie in $[t_j, t_{j+1} - 1]$. In this case, $r^{t, \tau} \leq r_{\text{max}}$ for any τ . Thus for any $t_{j+1} - t_j > \tau \gg r_{\text{max}}$, $r^{t, \tau} \ll \min(\tau, n)$. Next consider the case when $t - \tau + 1 \in [t_j, t_{j+1} - 1]$ and $t \in [t_{j+1}, t_{j+2} - 1]$. In this case, $r^{t, \tau} \leq r_{\text{max}} + c_{\text{max}}$. Thus, for any $t_{j+2} - t_j > \tau \gg r_{\text{max}} + c_{\text{max}}$, $r^{t, \tau} \ll \min(\tau, n)$. Finally consider the case when $t - \tau + 1 \in [t_j, t_{j+1} - 1]$ and $t \in [t_{j+k+1}, t_{j+k+2} - 1]$ for a $0 < k < J - 1$. In this case, τ can be rewritten as $\tau = (t_{j+k+1} - t_{j+1}) + \tau_1 + \tau_2$ with $\tau_1 := t_{j+1} - (t - \tau + 1)$ and $\tau_2 := t - (t_{j+k+1} - 1)$. Clearly, $r^{t, \tau} \leq (r_{\text{max}} + (k-1)c_{\text{max}}) + \min(\tau_1, c_{\text{max}}) + \min(\tau_2, c_{\text{max}}) < k r_{\text{max}} + \min(\tau_1, c_{\text{max}}) + \min(\tau_2, c_{\text{max}}) \ll (t_{j+k+2} - t_{j+1}) + \min(\tau_1, c_{\text{max}}) + \min(\tau_2, c_{\text{max}}) \leq (t_{j+k+2} - t_{j+1}) + \tau_1 + \tau_2 = \tau$. Moreover, $r^{t, \tau} \leq r_{\text{max}} + (k+1)c_{\text{max}} \leq r_{\text{max}} + J c_{\text{max}} \ll n$. Thus, in this case again for any τ , $r^{t, \tau} \ll \min(\tau, n)$.

The above piecewise constant subspace change model is a simplified model for what typically happens in practice. In most cases, P_t changes a little at each t in such a way that the low-dimensional assumption approximately holds. If we try to model this, it would result in a nonstationary model that is difficult to precisely define or to verify (it would require multiple video sequences of the same type to verify).³

Since background images typically change only a little over time (except in case of a camera viewpoint change or a scene change), it is valid to model the mean-subtracted background image sequence as lying in a slowly changing low-dimensional subspace. We verify this assumption in Section V.

B. Denseness

To state the denseness assumption, we first need to define the denseness coefficient. This is a simplification of the one introduced in our earlier work [27].

Definition 2.1 (Denseness Coefficient): For a matrix or a vector B , define

$$\kappa_s(B) = \kappa_s(\text{range}(B)) := \max_{|T| \leq s} \|I_T' \text{basis}(B)\|_2 \quad (4)$$

where $\|\cdot\|_2$ is the vector or matrix 2-norm. Recall that $\text{basis}(B)$ is short for $\text{basis}(\text{range}(B))$. Similarly $\kappa_s(B)$ is short for $\kappa_s(\text{range}(B))$. Notice that $\kappa_s(B)$ is a property of the subspace $\text{range}(B)$. Note also that $\kappa_s(B)$ is a non-decreasing function of s and of $\text{rank}(B)$.

We assume that the subspace spanned by the L_t 's is dense, i.e.

$$\kappa_{2s}(P_{(j)}) = \kappa_{2s}([L_{t_j}, \dots, L_{t_{j+1}-1}]) \leq \kappa_*$$

for a κ_* significantly smaller than one. Moreover, a similar assumption holds for $P_{(j),\text{new}}$ with a tighter bound: $\kappa_{2s}(P_{(j),\text{new}}) \leq \kappa_{\text{new}} < \kappa_*$. This assumption is similar to one of the denseness assumptions used in [5], [38]. In [5], a bound is assumed on $\kappa_1(U)$ and $\kappa_1(V)$ where U and V are the matrices containing the left and right singular vectors of the entire matrix, $[L_1, L_2 \dots L_t]$; and a tighter bound is assumed on $\max_{i,j} |(UV^T)_{i,j}|$. In our notation, $U = [P_{(0)}, P_{(1),\text{new}}, \dots, P_{(J),\text{new}}]$.

The following lemma, proved in [27], relates the RIC of $I - PP'$, when P is a *basis matrix*, to the denseness coefficient for $\text{range}(P)$. Notice that $I - PP'$ is an $n \times n$ matrix that has rank $(n - \text{rank}(P))$ and so it cannot be inverted.

Lemma 2.2: For a *basis matrix*, P ,

$$\delta_s(I - PP') = \kappa_s(P)^2.$$

Thus, the denseness assumption implies that the RIC of the matrix $(I - P_{(j)}P'_{(j)})$ is small. Using any of the RIC based sparse recovery results, e.g. [39], this ensures that for $t \in [t_j, t_{j+1})$, s -sparse vectors S_t are recoverable from $(I - P_{(j)}P'_{(j)})M_t = (I - P_{(j)}P'_{(j)})S_t$ by ℓ_1 minimization.

Very often, the background images primarily change due to lighting changes (in case of indoor sequences) or due to moving waters or moving leaves (in case of many outdoor sequences) [5], [27]. All of these result in global changes and hence it is

³With letting a_t be a zero mean random variable with a covariance matrix that is constant for sub-intervals within $[t_j, t_{j+1})$, the above model is a piecewise wide sense stationary approximation to the nonstationary model.

valid to assume that the subspace spanned by the background image sequences is dense.

C. Small Support Size, Some Support Change, Small Support Change Assumption on S_t

Let the sets of support additions and removals be

$$\Delta_t := T_t \setminus T_{t-1}, \quad \Delta_{e,t} := T_{t-1} \setminus T_t.$$

(1) We assume that

$$|T_t| + \min(|T_t|, |\Delta_t| + |\Delta_{e,t}|) \leq s + s_\Delta \text{ where } s_\Delta \ll s.$$

In particular, this implies that we either need $|T_t| \leq s$ and $|\Delta_t| + |\Delta_{e,t}| \leq s_\Delta$ (S_t is sparse with support size at most s , and its support changes slowly) or, in cases when the change $|\Delta_t| + |\Delta_{e,t}|$ is large, we need $|T_t| \leq 0.5(s + s_\Delta)$ (need a tighter bound on the support size).

(2) We also assume that there is *some* support change every few frames, i.e. at least once every h frames, $|\Delta_t| > s_{\Delta, \text{min}}$. Practically, this is needed to ensure that at least some of the background behind the foreground is visible so that the changes to the background subspace can be estimated.

In the video application, foreground images typically consist of one or more moving objects/people/regions and hence are sparse. Also, typically the objects are not static, i.e. there is some support change at least every few frames. On the other hand, since the objects usually do not move very fast, slow support change is also valid most of the time. The time when the support change is almost comparable to the support size is usually when the object is entering or leaving the image, but these are the exactly the times when the object's support size is itself small (being smaller than $0.5(s + s_\Delta)$ is a valid). We show some verification of these assumptions in Section V.

III. PRAC-REPROCS: PRACTICAL REPROCS

We first develop a practical algorithm based on the basic ReProCS idea from our earlier work [27]. Then we discuss how the sparse recovery and support estimation steps can be improved. The complete algorithm is summarized in Algorithm 1. Finally we discuss an alternate subspace update procedure in Section III-D.

A. Basic Algorithm

We use $\hat{S}_t, \hat{T}_t, \hat{L}_t$ to denote estimates of S_t , its support, T_t , and L_t respectively; and we use \hat{P}_t to denote the *basis matrix* for the estimated subspace of L_t at time t . Also, let

$$\Phi_t := (I - \hat{P}_{t-1}\hat{P}'_{t-1}). \quad (5)$$

Given the initial training sequence which does not contain the sparse components, $\mathcal{M}_{\text{train}} = [L_1, L_2, \dots, L_{t_{\text{train}}}]$ we compute \hat{P}_0 as an approximate basis for $\mathcal{M}_{\text{train}}$, i.e. $\hat{P}_0 = \text{approx-basis}(\mathcal{M}_{\text{train}}, b\%)$. Let $\hat{r} = \text{rank}(\hat{P}_0)$. We need to compute an approximate basis because for real data, the L_t 's are only approximately low-dimensional. We use $b\% = 95\%$ or $b\% = 99.99\%$ depending on whether the low-rank part is approximately low-rank or almost exactly low-rank. After this, at each time t , ReProCS involves 4 steps:

(a) Perpendicular Projection; (b) Sparse Recovery (recover T_t and S_t); (c) Recover L_t ; (d) Subspace Update (update \hat{P}_t).

Perpendicular Projection. In the first step, at time t , we project the measurement vector, M_t , into the space orthogonal to $\text{range}(\hat{P}_{t-1})$ to get the projected measurement vector,

$$y_t := \Phi_t M_t. \quad (6)$$

Sparse Recovery (Recover T_t and S_t). With the above projection, y_t can be rewritten as

$$y_t = \Phi_t S_t + \beta_t \text{ where } \beta_t := \Phi_t L_t. \quad (7)$$

Because of the slow subspace change assumption, projecting orthogonal to $\text{range}(\hat{P}_{t-1})$ nullifies most of the contribution of L_t and hence β_t can be interpreted as small “noise”. We explain this in detail in Appendix A.

Thus, the problem of recovering S_t from y_t becomes a traditional noisy sparse recovery/CS problem. Notice that, since the $n \times n$ projection matrix, Φ_t , has rank $n - \text{rank}(\hat{P}_{t-1})$, therefore y_t has only this many “effective” measurements, even though its length is n . To recover S_t from y_t , one can use ℓ_1 minimization [39], [40], or any of the greedy or iterative thresholding algorithms from literature. In this work we use ℓ_1 minimization: we solve

$$\min_x \|x\|_1 \text{ s.t. } \|y_t - \Phi_t x\|_2 \leq \xi \quad (8)$$

and denote its solution by $\hat{S}_{t,cs}$. By the denseness assumption, P_{t-1} is dense. Since \hat{P}_{t-1} approximates it, this is true for \hat{P}_{t-1} as well [27, Lemma 6.6]. Thus, by Lemma 2.2, the RIC of Φ_t is small enough. Using [39, Theorem 1], this and the fact that β_t is small ensures that S_t can be accurately recovered from y_t . The constraint ξ used in the minimization should equal $\|\beta_t\|_2$ or its upper bound. Since β_t is unknown we set $\xi = \|\hat{\beta}_t\|_2$ where $\hat{\beta}_t := \Phi_t \hat{L}_{t-1}$.

By thresholding on $\hat{S}_{t,cs}$ to get an estimate of its support followed by computing a least squares (LS) estimate of S_t on the estimated support and setting it to zero everywhere else, we can get a more accurate estimate, \hat{S}_t , as suggested in [41]. We discuss better support estimation and its parameter setting in Section III-C.

Recover L_t . The estimate \hat{S}_t is used to estimate L_t as $\hat{L}_t = M_t - \hat{S}_t$. Thus, if S_t is recovered accurately, so will L_t .

Subspace Update (Update \hat{P}_t). Within a short delay after every subspace change time, one needs to update the subspace estimate, \hat{P}_t . To do this in a provably reliable fashion, we introduced the projection PCA (p-PCA) algorithm in [27]. The algorithm studied there used knowledge of the subspace change times t_j and of the number of new directions $c_{j,new}$. Let $\hat{P}_{(j-1)}$ denote the final estimate of a basis for the span of $P_{(j-1)}$. It is assumed that the delay between change times is large enough so that $\hat{P}_{(j-1)}$ is an accurate estimate. At $t = t_j + \alpha - 1$, p-PCA gets the first estimate of the new directions, $\hat{P}_{(j),new,1}$, by projecting the last α \hat{L}_t 's perpendicular to $\hat{P}_{(j-1)}$ followed by computing the $c_{j,new}$ top left singular vectors of the projected data matrix. It then updates the subspace estimate as $\hat{P}_t = [\hat{P}_{(j-1)}, \hat{P}_{(j),new,1}]$. The same procedure is repeated at every $t = t_j + k\alpha - 1$ for $k = 2, 3, \dots, K$ and each time we update the subspace as $\hat{P}_t = [\hat{P}_{(j-1)}, \hat{P}_{(j),new,k}]$. Here K is chosen so that the subspace estimation error decays down to a small enough value within K p-PCA steps.

In this paper, we design a practical version of p-PCA which does not need knowledge of t_j or $c_{j,new}$. This is summarized in Algorithm 1. The key idea is as follows. We let $\hat{\sigma}_{\min}$ be the \hat{r} th largest singular value of the training dataset. This serves as the noise threshold for approximately low rank data. We split projection PCA into two phases: “detect subspace change” and “p-PCA”. We are in the detect phase when the previous subspace has been accurately estimated. Denote the basis matrix for this subspace by $\hat{P}_{(j-1)}$. We detect the subspace change as follows. Every α frames, we project the last α \hat{L}_t 's perpendicular to $\hat{P}_{(j-1)}$ and compute the SVD of the resulting matrix. If there are any singular values above $\hat{\sigma}_{\min}$, this means that the subspace has changed. At this point, we enter the “p-PCA” phase. In this phase, we repeat the K p-PCA steps described above with the following change: we estimate $c_{j,new}$ as the number of singular values above $\hat{\sigma}_{\min}$, but clipped at $\lceil \alpha/3 \rceil$ (i.e. if the number is more than $\lceil \alpha/3 \rceil$ then we clip it to $\lceil \alpha/3 \rceil$). We stop either when the stopping criterion given in step 4biv is achieved ($k \geq K_{\min}$ and the projection of L_t along $\hat{P}_{new,k}$ is not too different from that along $\hat{P}_{new,k}$) or when $k \geq K_{\max}$.

For the above algorithm, with theoretically motivated choices of algorithm parameters, under the assumptions from Section II, it is possible to show that, w.h.p., the support of S_t is exactly recovered, the subspace of L_t 's is accurately recovered within a finite delay of the change time. We provide a brief overview of the proof from [27], [29] in Appendix A that helps explain why the above approach works.

Remark 3.1: The p-PCA algorithm only allows addition of new directions. If the goal is to estimate the span of $[L_1, \dots, L_t]$, then this is what is needed. If the goal is sparse recovery, then one can get a smaller rank estimate of \hat{P}_t by also including a step to delete the span of the removed directions, $P_{(j),old}$. This will result in more “effective” measurements available for the sparse recovery step and hence possibly in improved performance. The simplest way to do this is to do one simple PCA step every some frames. In our experiments, this did not help much though. A provably accurate solution is described in [27, Sec VII].

Remark 3.2: The p-PCA algorithm works on small batches of α frames. This can be made fully recursive if we compute the SVD of $(I - \hat{P}_{(j-1)} \hat{P}'_{(j-1)})[\hat{L}_{t-\alpha+1}, \dots, \hat{L}_t]$ using the incremental SVD (inc-SVD) procedure summarized in Algorithm 2 [13] for one frame at a time. As explained in [13] and references therein, we can get the left singular vectors and singular values of any matrix $M = [M_1, M_2, \dots, M_\alpha]$ recursively by starting with $\hat{P} = []$, $\hat{\Sigma} = []$ and calling $[\hat{P}, \hat{\Sigma}] = \text{inc-SVD}(\hat{P}, \hat{\Sigma}, M_i)$ for every column i or for short batches of columns of size of α/k . Since we use $\alpha = 20$ which is a small value, the use of incremental SVD does not speed up the algorithm in practice and hence we do not report results using it.

B. Exploiting Slow Support Change When Valid

In [27], [28], we always used ℓ_1 minimization followed by thresholding and LS for sparse recovery. However if slow support change holds, one can replace simple ℓ_1 minimization by modified-CS [30] which requires fewer measurements for exact/accurate recovery as long as the previous support estimate, \hat{T}_{t-1} , is an accurate enough predictor of the current support, T_t . In our application, \hat{T}_{t-1} is likely to contain a

Algorithm 1 Practical ReProCS-pPCA

Input: M_t ; **Output:** $\hat{T}_t, \hat{S}_t, \hat{L}_t$; **Parameters:** $q, b, \alpha, K_{\min}, K_{\max}$. We used $\alpha = 20, K_{\min} = 3, K_{\max} = 10$ in all experiments (α needs to only be large compared to c_{\max}); we used $b = 95$ for approximately low-rank data (all real videos and the lake video with simulated foreground) and used $b = 99.99$ for almost exactly low rank data (simulated data); we used $q = 1$ whenever $\|S_t\|_2$ was of the same order or larger than $\|L_t\|_2$ (all real videos and the lake video) and used $q = 0.25$ when it was much smaller (simulated data with small magnitude S_t).

Initialization

- $[\hat{P}_0, \hat{\Sigma}_0] \leftarrow \text{approx-basis}(\frac{1}{\sqrt{t_{\text{train}}}}[M_1, \dots, M_{t_{\text{train}}}], b\%)$.
- Set $\hat{r} \leftarrow \text{rank}(\hat{P}_0)$, $\hat{\sigma}_{\min} \leftarrow ((\hat{\Sigma}_0)_{\hat{r}, \hat{r}})$, $\hat{t}_0 = t_{\text{train}}$, flag = detect
- Initialize $\hat{P}_{(t_{\text{train}})} \leftarrow \hat{P}_0$ and $\hat{T}_t \leftarrow [\cdot]$.

For $t > t_{\text{train}}$ **do**

- 1) Perpendicular Projection: compute $y_t \leftarrow \Phi_t M_t$ with $\Phi_t \leftarrow I - \hat{P}_{t-1} \hat{P}'_{t-1}$
 - 2) Sparse Recovery (Recover S_t and T_t)
 - a) If $\frac{|\hat{T}_{t-2} \cap \hat{T}_{t-1}|}{|\hat{T}_{t-2}|} < 0.5$
 - i) Compute $\hat{S}_{t,cs}$ as the solution of (8) with $\xi = \|\Phi_t \hat{L}_{t-1}\|_2$.
 - ii) $\hat{T}_t \leftarrow \text{Thresh}(\hat{S}_{t,cs}, \omega)$ with $\omega = q\sqrt{\|M_t\|^2/n}$. Here $T \leftarrow \text{Thresh}(x, \omega)$ means that $T = \{i : |(x)_i| \geq \omega\}$.
Else
 - i) Compute $\hat{S}_{t,cs}$ as the solution of (9) with $T = \hat{T}_{t-1}$, $\lambda = \frac{|\hat{T}_{t-2} \setminus \hat{T}_{t-1}|}{|\hat{T}_{t-1}|}$, $\xi = \|\Phi_t \hat{L}_{t-1}\|_2$.
 - ii) $\hat{T}_{\text{add}} \leftarrow \text{Prune}(\hat{S}_{t,cs}, 1.4|\hat{T}_{t-1}|)$. Here $T \leftarrow \text{Prune}(x, k)$ returns indices of the k largest magnitude elements of x .
 - iii) $\hat{S}_{t,add} \leftarrow \text{LS}(y_t, \Phi_t, \hat{T}_{\text{add}})$. Here $\hat{x} \leftarrow \text{LS}(y, A, T)$ means that $\hat{x}_T = (A_T' A_T)^{-1} A_T' y$ and $\hat{x}_{T^c} = 0$.
 - iv) $\hat{T}_t \leftarrow \text{Thresh}(\hat{S}_{t,add}, \omega)$ with $\omega = q\sqrt{\|M_t\|^2/n}$.
 - b) $\hat{S}_t \leftarrow \text{LS}(y_t, \Phi_t, \hat{T}_t)$
 - 3) Estimate L_t : $\hat{L}_t \leftarrow M_t - \hat{S}_t$
 - 4) Update \hat{P}_t : projection PCA
 - a) If flag = detect and $\text{mod}(t - \hat{t}_j + 1, \alpha) = 0$, (here $\text{mod}(t, \alpha)$ is the remainder when t is divided by α)
 - i) compute the SVD of $\frac{1}{\sqrt{\alpha}}(I - \hat{P}_{(j-1)} \hat{P}'_{(j-1)})[\hat{L}_{t-\alpha+1}, \dots, \hat{L}_t]$ and check if any singular values are above $\hat{\sigma}_{\min}$
 - ii) if the above number is more than zero then set flag \leftarrow pPCA, increment $j \leftarrow j + 1$, set $\hat{t}_j \leftarrow t - \alpha + 1$, reset $k \leftarrow 1$
Else $\hat{P}_t \leftarrow \hat{P}_{t-1}$.
 - b) If flag = pPCA and $\text{mod}(t - \hat{t}_j + 1, \alpha) = 0$,
 - i) compute the SVD of $\frac{1}{\sqrt{\alpha}}(I - \hat{P}_{(j-1)} \hat{P}'_{(j-1)})[\hat{L}_{t-\alpha+1}, \dots, \hat{L}_t]$,
 - ii) let $\hat{P}_{j,\text{new},k}$ retain all its left singular vectors with singular values above $\hat{\sigma}_{\min}$ or all $\alpha/3$ top left singular vectors whichever is smaller,
 - iii) update $\hat{P}_t \leftarrow [\hat{P}_{(j-1)} \hat{P}_{j,\text{new},k}]$, increment $k \leftarrow k + 1$
 - iv) If $k \geq K_{\min}$ and $\frac{\|\sum_{i=\alpha+1}^t (\hat{P}_{j,\text{new},i-1} \hat{P}'_{j,\text{new},i-1} - \hat{P}_{j,\text{new},i} \hat{P}'_{j,\text{new},i}) L_t\|_2}{\|\sum_{i=\alpha+1}^t \hat{P}_{j,\text{new},i-1} \hat{P}'_{j,\text{new},i-1} L_t\|_2} < 0.01$ for $i = k - 2, k - 1, k$; or $k = K_{\max}$,
then $K \leftarrow k$, $\hat{P}_{(j)} \leftarrow [\hat{P}_{(j-1)} \hat{P}_{j,\text{new},K}]$ and reset flag \leftarrow detect.
Else $\hat{P}_t \leftarrow \hat{P}_{t-1}$.
-

significant number of extras and in this case, a better idea is to solve the following weighted ℓ_1 problem [42], [43]

$$\min_x \lambda \|x_T\|_1 + \|x_{T^c}\|_1 \text{ s.t. } \|y_t - \Phi_t x\|_2 \leq \xi, T := \hat{T}_{t-1} \quad (9)$$

with $\lambda < 1$ (modified-CS solves the above with $\lambda = 0$). Denote its solution by $\hat{S}_{t,cs}$. One way to pick λ is to let it be proportional to the estimate of the percentage of extras in \hat{T}_{t-1} . If slow support change does not hold, the previous support estimate is not a good predictor of the current support. In this case, doing the above is a bad idea and one should instead solve simple ℓ_1 , i.e. solve (9) with $\lambda = 1$. As explained in [43], if the support estimate contains at least 50% correct entries, then weighted ℓ_1 is better than simple ℓ_1 . We use the above criteria with true values replaced by estimates. Thus, if $\frac{|\hat{T}_{t-2} \cap \hat{T}_{t-1}|}{|\hat{T}_{t-2}|} > 0.5$, then we solve (9) with $\lambda = \frac{|\hat{T}_{t-2} \setminus \hat{T}_{t-1}|}{|\hat{T}_{t-1}|}$, else we solve it with $\lambda = 1$.

C. Improved Support Estimation

A simple way to estimate the support is by thresholding the solution of (9). This can be improved by using the Add-LS-Del procedure for support and signal value estimation [30]. We proceed as follows. First we compute the set $\hat{T}_{t,add}$ by thresholding on $\hat{S}_{t,cs}$ in order to retain its k largest magnitude entries. We then compute a LS estimate of S_t on $\hat{T}_{t,add}$ while setting it to

zero everywhere else. As explained earlier, because of the LS step, $\hat{S}_{t,add}$ is a less biased estimate of S_t than $\hat{S}_{t,cs}$. We let $k = 1.4|\hat{T}_{t-1}|$ to allow for a small increase in the support size from $t - 1$ to t . A larger value of k also makes it more likely that elements of the set $(T_t \setminus \hat{T}_{t-1})$ are detected into the support estimate.⁴

The final estimate of the support, \hat{T}_t , is obtained by thresholding on $\hat{S}_{t,add}$ using a threshold ω . If ω is appropriately chosen, this step helps to delete some of the extra elements from \hat{T}_{add} and this ensures that the size of \hat{T}_t does not keep increasing (unless the object's size is increasing). An LS estimate computed on \hat{T}_t gives us the final estimate of S_t , i.e. $\hat{S}_t = \text{LS}(y_t, A, \hat{T}_t)$. We use $\omega = \sqrt{\|M_t\|^2/n}$ except in situations where $\|S_t\| \ll \|L_t\|$ —in this case we use $\omega = 0.25\sqrt{\|M_t\|^2/n}$. An alternate approach is to let ω be proportional to the noise magnitude seen by the ℓ_1 step, i.e. to let $\omega = q\|\hat{\beta}_t\|_\infty$, however this approach required different values of q for different experiments (it is not possible to specify one q that works for all experiments).

The complete algorithm with all the above steps is summarized in Algorithm 1.

⁴Due to the larger weight on the $\|x_{(\hat{T}_{t-1}^c)}\|_1$ term as compared to that on the $\|x_{(\hat{T}_{t-1})}\|_1$ term, the solution of (9) is biased towards zero on \hat{T}_{t-1}^c and thus the solution values along $(T_t \setminus \hat{T}_{t-1})$ are smaller than the true ones.

Algorithm 2: $[\hat{P}, \hat{\Sigma}] = \text{inc-SVD}(\hat{P}, \hat{\Sigma}, D)$

- 1) set $D_{\parallel, \text{proj}} \leftarrow \hat{P}'D$ and $D_{\perp} \leftarrow (I - \hat{P}\hat{P}')D$
- 2) compute QR decomposition of D_{\perp} , i.e. $D_{\perp} \stackrel{QR}{=} JK$ (here J is a *basis matrix* and K is an upper triangular matrix)
- 3) compute the SVD: $\begin{bmatrix} \hat{\Sigma} & D_{\parallel, \text{proj}} \\ 0 & K \end{bmatrix} \stackrel{SVD}{=} \tilde{P}\tilde{\Sigma}\tilde{V}'$
- 4) update $\hat{P} \leftarrow [\hat{P} \ J]\tilde{P}$ and $\hat{\Sigma} \leftarrow \tilde{\Sigma}$

Note: As explained in [13], due to numerical errors, step 4 done too often can eventually result in \hat{P} no longer being a basis matrix. This typically occurs when one tries to use inc-SVD at every time t , i.e. when D is a column vector. This can be addressed using the modified Gram Schmidt re-orthonormalization procedure whenever loss of orthogonality is detected [13].

D. Simplifying Subspace Update: Simple Recursive PCA

Even the practical version of p-PCA needs to set K_{\min} and K_{\max} besides also setting b and α . Thus, we also experiment with using PCA to replace p-PCA (it is difficult to prove a performance guarantee with PCA but that does not necessarily mean that the algorithm itself will not work). The simplest way to do this is to compute the top \hat{r} left singular vectors of $[\hat{L}_1, \hat{L}_2, \dots, \hat{L}_t]$ either at each time t or every α frames. While this is simple, its complexity will keep increasing with time t which is not desirable. Even if we use the last d frames instead of all past frames, d will still need to be large compared to \hat{r} to get an accurate estimate. To address this issue, we can use the recursive PCA (incremental SVD) algorithm given in Algorithm 2. We give the complete algorithm that uses this and a rank \hat{r} truncation step every d frames (motivated by [13]) in Algorithm 3.

Algorithm 3: Practical ReProCS-Recursive-PCA

Input: M_t ; **Output:** $\hat{T}_t, \hat{S}_t, \hat{L}_t$; **Parameters:** q, b, α , set $\alpha = 20$ in all experiments, set q, b as explained in Algorithm 1.

Initialization: $[\hat{P}_0, \Sigma_0] \leftarrow$
approx - basis($[M_1, \dots, M_{t_{\text{train}}}], b\%$),
 $\hat{r} \leftarrow \text{rank}(\hat{P}_0)$, $d \leftarrow 3\hat{r}$, initialize $\hat{P}_{t_{\text{mp}}} \leftarrow \hat{P}_0$, $\hat{\Sigma}_{t_{\text{mp}}} \leftarrow \Sigma_0$,
 $\hat{P}_{(t_{\text{train}})} \leftarrow \hat{P}_0$ and $\hat{T}_t \leftarrow [\cdot]$. **For** $t > t_{\text{train}}$ **do**

- 1) Perpendicular Projection: do as in Algorithm 1.
- 2) Sparse Recovery: do as in Algorithm 1.
- 3) Estimate L_t : do as in Algorithm 1.
- 4) Update \hat{P}_t : recursive PCA
 - a) If $\text{mod}(t - t_{\text{train}}, \alpha) = 0$,
 - i) $[\hat{P}_{t_{\text{mp}}}, \hat{\Sigma}_{t_{\text{mp}}}] \leftarrow$
inc - SVD($\hat{P}_{t_{\text{mp}}}, \hat{\Sigma}_{t_{\text{mp}}}, [\hat{L}_{t-\alpha+1}, \dots, \hat{L}_t]$)
where inc-SVD is given in Algorithm 2.
 - ii) $\hat{P}_t \leftarrow (\hat{P}_{t_{\text{mp}}})_{1:\hat{r}}$
Else $\hat{P}_t \leftarrow \hat{P}_{t-1}$.
 - b) If $\text{mod}(t - t_{\text{train}}, d) = 0$,
 - i) $\hat{P}_{t_{\text{mp}}} \leftarrow (\hat{P}_{t_{\text{mp}}})_{1:\hat{r}}$ and $\hat{\Sigma}_{t_{\text{mp}}} \leftarrow (\hat{\Sigma}_{t_{\text{mp}}})_{1:\hat{r}, 1:\hat{r}}$

IV. COMPRESSIVE MEASUREMENTS: RECOVERING S_t

Consider the problem of recovering S_t from

$$M_t := AS_t + BL_t$$

when A and B are $m \times n$ and $m \times n_2$ matrices, S_t is an n length vector and L_t is an n_2 length vector. In general m can be larger, equal or smaller than n or n_2 . In the compressive measurements' case, $m < n$. To specify the assumptions needed in this case, we need to define the basis matrix for $\text{range}(BP_{(j)})$ and we need to define a generalization of the denseness coefficient.

Definition 4.1: Let $Q_j := \text{basis}(BP_{(j)})$ and let $Q_{j, \text{new}} := \text{basis}((I - Q_{j-1}Q_{j-1}')BP_{(j)})$.

Definition 4.2: For a matrix or a vector M , define

$$\kappa_{s,A}(M) = \kappa_{s,A}(\text{range}(M)) := \max_{|T| \leq s} \|A_T' \text{basis}(M)\|_2 \quad (10)$$

where $\|\cdot\|_2$ is the vector or matrix 2-norm. This quantifies the incoherence between the subspace spanned by any set of s columns of A and the range of M .

We assume the following.

- 1) L_t and S_t satisfy the assumptions of Section II-A, II-C.
- 2) The matrix A satisfies the restricted isometry property [39], i.e. $\delta_s(A) \leq \delta_* \ll 1$.
- 3) The denseness assumption is replaced by: $\kappa_{2s,A}(Q_j) \leq \kappa_*$, $\kappa_{2s,A}(Q_{j, \text{new}}) \leq \kappa_{\text{new}} < \kappa_*$ for a κ_* that is small compared to one. Notice that this depends on the L_t 's and on the matrices A and B .

Assume that we are given an initial training sequence that satisfies $M_t = BL_t$ for $t = 1, 2, \dots, t_{\text{train}}$. The goal is to recover S_t at each time t . It is not possible to recover L_t unless B is time-varying (this case is studied in [44]). In many imaging applications, e.g. undersampled fMRI or single-pixel video imaging, $B = A$ ($B = A$ is a partial Fourier matrix for MRI and is a random Gaussian or Rademacher matrix for single-pixel imaging). On the other hand, if L_t is large but low-dimensional sensor noise, then $B = I$ (identity matrix), while A is the measurement matrix.

Let $\tilde{L}_t := BL_t$. It is easy to see that if L_t lies in a slowly changing low-dimensional subspace, the same is also true for the sequence \tilde{L}_t . Consider the problem of recovering S_t from $M_t := AS_t + \tilde{L}_t$ when an initial training sequence $M_t := \tilde{L}_t$ for $t = 1, 2, \dots, t_{\text{train}}$ is available. Using this sequence, it is possible to estimate its approximate basis \hat{Q}_0 as explained earlier. If we then project M_t into the subspace orthogonal to $\text{range}(\hat{Q}_0)$, the resulting vector $y_t := \Phi_t M_t$ satisfies

$$y_t = (\Phi_t A)S_t + \beta_t$$

where $\beta_t = \Phi_t BL_t$ is small noise for the same reasons explained earlier. Thus, one can still recover S_t from y_t by ℓ_1 or weighted ℓ_1 minimization followed by support recovery and LS. Then, \tilde{L}_t gets recovered as $\hat{\tilde{L}}_t \leftarrow M_t - A\hat{S}_t$ and this is used for updating its subspace estimate. We summarize the resulting algorithm in Algorithm 4. This is being analyzed in ongoing work [45].

The following lemma explains why some of the extra assumptions are needed for this case.

Lemma 4.3: [45] For a *basis matrix*, Q ,

$$\delta_s((I - QQ')A) \leq \kappa_{s,A}(Q)^2 + \delta_s(A)$$

Using the above lemma with $Q \equiv Q_j$, it is clear that incoherence of Q_j w.r.t. any set of $2s$ columns of A along with RIP of A ensures that any s sparse vector x can be recovered from




Fig. 1. (a) Verification of slow subspace change assumption. (b) Verification of denseness assumption. (c) Verification of small support size, small support change.

$y := (I - Q_j Q_j')Ax$ by ℓ_1 minimization. In compressive ReProCS, the measurement matrix uses \hat{Q}_j instead of Q_j and also involves small noise. With more work, these arguments can be extended to this case as well [see [45]].

Algorithm 4: Compressive ReProCS

Use Algorithm 1 or 3 with the following changes.

- Replace Φ_t in step 2 by $\hat{\Phi}_t$.
 - Replace step 3 by $\hat{L}_t \leftarrow M_t - A\hat{S}_t$.
 - Use \hat{L}_t in place of L_t and \hat{Q} in place of \hat{P} everywhere.
-

V. MODEL VERIFICATION

A. Low-Dimensional and Slow Subspace Change Assumption

We used two background image sequence datasets. The first was a video of lake water motion. For this sequence, $n = 6480$ and the number of images were 1500. The second was an indoor video of window curtains moving due to the wind. There was also some lighting variation. The latter part of this sequence also contains a foreground (various persons coming in, writing on the board and leaving). For this sequence, the image size was $n = 5120$ and the number of background-only images were 1755. Both sequences are posted at <http://www.ece.iastate.edu/~hanguo/PracReProCS.html>.

First note that any given background image sequence will never be exactly low-dimensional, but only be approximately so. Secondly, in practical data, the subspace does not just change as simply as in the model of Section II-A. Typically there are some changes to the subspace at every time t . Moreover, with just one training sequence of a given type, it is not possible to estimate the covariance matrix of L_t at each t and thus one cannot detect the subspace change times. The only thing one can do is to assume that there may be a change every τ frames, and that during these τ frames the L_t 's are stationary and ergodic; estimate the covariance matrix of L_t for this period using a time average; compute its eigenvectors corresponding to $b\%$ energy (or equivalently compute the $b\%$ approximate basis of $[L_{t-\tau+1}, \dots, L_t]$) and use these as $P_{(j)}$. These can be used to test our assumptions.

Testing for slow subspace change can be done in various ways. In [27, Fig 6], we do this after low-rankifying the video data first. This helps to very clearly demonstrate slow subspace change, but then it is not checking what we really do on real

video data. In this work, we proceed without low-rankifying the data. We let $t_0 = 0$ and $t_j = t_0 + j\tau$ with $\tau = 725$. Let L_t denote the mean subtracted background image sequence, i.e. $L_t = B_t - \mu$ where $\mu = (1/t_1) \sum_{t=0}^{t_1} B_t$. We computed $P_{(j)}$ as $P_{(j)} = \text{approx-basis}([L_{t_j}, \dots, L_{t_{j+1}-1}], 95\%)$. We observed that $\text{rank}(P_{(j)}) \leq 38$ for curtain sequence, while $\text{rank}(P_{(j)}) \leq 33$ for lake sequence. In other words, 95% of the energy is contained in only 38 or lesser directions in either case, i.e. both sequences are approximately low-dimensional. Notice that the dimension of the matrix $[L_{t_j}, \dots, L_{t_{j+1}-1}]$ is $n \times \tau$ and $\min(n, \tau) = \tau = 725$ is much larger than 38. To test for slow subspace change, in Fig. 1(a), we plot $\|(I - P_{(j-1)}P'_{(j-1)})L_t\|_2 / \|L_t\|_2$ when $t \in [t_j, t_{j+1})$. Notice that, after every change time ($t_j = 725, 1450$), this quantity is initially small for the first 100–150 frames and then increases gradually. It later decreases also but that is allowed (all we need is that it be small initially and increase slowly).

B. Denseness Assumption

Exactly verifying the denseness assumption is impossible since computing $\kappa_s(\cdot)$ has exponential complexity (one needs to check all sets T of size s). Instead, to get some idea if it holds even just for T replaced by T_t , in Fig. 1(b), we plot $\max_i \|J_{T_t}'(P_{(j)})i\|_2$ where T_t is the true or estimated support of S_t at time t . For the lake sequence, T_t is simulated and hence known. For the curtain sequence, we select a part of the sequence in which the person is wearing a black shirt (against a white curtains' background). This part corresponds to $t = 35$ to $t = 80$. For this part, ReProCS returns a very accurate estimate of T_t , and we use this estimated support as a proxy for the true support T_t .

C. Support Size, Support Change and Slow Support Change

For real video sequences, it is not possible to get the true foreground support. Thus we used \hat{T}_t for the part of the curtain sequence described above in Section V-B as a proxy for T_t . We plot the support size normalized by the image size $|T_t|/n$, and we plot the number of additions and removals normalized by the support size, i.e. $|\Delta_t|/|T_t|$ and $|\Delta_{e,t}|/|T_t|$ in Fig. 1(c). Notice from the figure that the support size is at most 10.2% of the image size. Notice also that at least at every 3 frames, there is at least a 1% support change. Thus there is some support change every few frames, thus exposing the part of the background behind the foreground. Finally notice that the maximum number

TABLE I

COMPARISON OF RECONSTRUCTION ERRORS OF DIFFERENT ALGORITHMS FOR SIMULATED DATA. HERE, $|T_t|/n$ IS THE SPARSITY RATIO OF S_t , $\mathbb{E}[\cdot]$ DENOTES THE MONTE CARLO AVERAGE COMPUTED OVER 100 REALIZATIONS AND $\|\cdot\|_F$ IS THE FROBENIUS NORM OF A MATRIX. ALSO, $S = [S_1, S_2, \dots, S_{t_{\max}}]$ AND \hat{S} IS ITS ESTIMATE; $(O_t)_i = (M_t)_i$ IF $i \in T_t$ AND $(O_t)_i = 0$ OTHERWISE AND \hat{O}_t IS DEFINED SIMILARLY WITH THE ESTIMATES. O AND \hat{O} ARE THE CORRESPONDING MATRICES. WE SHOW ERROR FOR O FOR iRSL AND ADAPTED-iSVD SINCE THESE ALGORITHMS CAN ONLY RETURN AN ESTIMATE OF THE OUTLIER SUPPORT T_t ; THEY DO NOT RETURN THE BACKGROUND ESTIMATE

(a) $(S_t)_i = 100$ for $i \in T_t$ and $(S_t)_i = 0$ for $i \in T_t^c$						
	$\mathbb{E}\ S - \hat{S}\ _F^2 / \mathbb{E}\ S\ _F^2$				$\mathbb{E}\ O - \hat{O}\ _F^2 / \mathbb{E}\ O\ _F^2$	
$ T_t /n$	ReProCS-pPCA	RSL	PCP	GRASTA	adapted-iSVD	iRSL
9%	1.52×10^{-4}	0.0580	0.0021	3.75×10^{-4}	0.0283	0.9105
27%	1.90×10^{-4}	0.0198	0.6852	0.1043	0.0637	0.9058

(b) $(S_t)_i = 10$ for $i \in T_t$ and $(S_t)_i = 0$ for $i \in T_t^c$						
	$\mathbb{E}\ S - \hat{S}\ _F^2 / \mathbb{E}\ S\ _F^2$				$\mathbb{E}\ O - \hat{O}\ _F^2 / \mathbb{E}\ O\ _F^2$	
$ T_t /n$	ReProCS-pPCA	RSL	PCP	GRASTA	adapted-iSVD	iRSL
9%	0.0344	8.7247	0.2120	0.1390	0.2346	0.9739
27%	0.0668	3.3166	0.6456	0.1275	0.3509	0.9778

of support changes is only 9.9% of the support size, i.e. slow support change holds for this piece.

VI. SIMULATION AND EXPERIMENTAL RESULTS

In this section, we show comparisons of ReProCS with other batch and recursive algorithms for robust PCA. For implementing the ℓ_1 or weighted ℓ_1 minimizations, we used the YALL1 ℓ_1 minimization toolbox [46], its code is available at <http://yall1.blogs.rice.edu/>.

Code and data for our algorithms and for all experiments given below is available at http://www.ece.ias.tate.edu/~hanguo/ReProCS_demo.rar.

Simulated Data: In this experiment, the measurement at time t , $M_t := L_t + S_t$, is an $n \times 1$ vector with $n = 100$. We generated L_t using the autoregressive model described in [1] with auto-regression parameter 0.1, and the decay parameter $f_d = 0.1$. The covariance of a direction decayed to zero before being removed. There was one change time t_1 . For $t < t_1$, $P_t = P_0$ was a rank $r_0 = 20$ matrix and $Cov(a_t)$ was a diagonal matrix with entries $10^4, 0.7079 \times 10^4, 0.7079^2 \times 10^4, \dots, 14.13$. At $t = t_1$, $c = c_{1,\text{new}} = 2$ new directions, $P_{1,\text{new}}$, got added with $Cov(a_{t,\text{new}})$ being diagonal with entries 60 and 50. Also, the variance along two existing directions started to decay to zero exponentially. We used $t_{\text{train}} = 2000$ and $t_1 = t_{\text{train}} + 5$. The matrix $[P_0 \ P_{1,\text{new}}]$ was generated as the first 22 columns of an $n \times n$ random orthonormal matrix (generated by first generating an $n \times n$ matrix random Gaussian matrix and then orthonormalizing it). For $1 \leq t \leq t_{\text{train}}$, $S_t = 0$ and hence $M_t = L_t$. For $t > t_{\text{train}}$, the support set, T_t , was generated in a correlated fashion: S_t contained one block of size 9 or 27 (small and large support size cases). The block stayed static with probability 0.8 and move up or down by one pixel with probability 0.1 each independently at each time. Thus the support sets were highly correlated. The magnitude of the nonzero elements of S_t is fixed at either 100 (large) or 10 (small).

For the small magnitude S_t case, $\|L_t\|_2$ ranged from 150 to 250 while $\|S_t\|_2$ was equal to 30 and 52, i.e. in this case $\|S_t\|_2 \ll \|L_t\|_2$. For the large magnitude case, $\|S_t\|_2$ was 300 and 520. We implemented ReProCS (Algorithm 1) with $b = 99.99$ since this data is exactly low-rank. We used $q = 0.25$ for the small magnitude S_t case and $q = 1$ for the other case.

We compared with three recursive robust PCA methods—incremental robust subspace learning (iRSL) [15] and adapted (outlier-detection enabled) incremental SVD (adapted-iSVD) [13] and GRASTA [31]—and with two batch methods—Principal Components Pursuit (PCP) [5]⁵ and robust subspace learning (RSL)⁶ [4]. Results are shown in Table I.

From these experiments, we can conclude that ReProCS is able to successfully recover both small magnitude and fairly large support-sized S_t 's; iRSL has very bad performance in both cases; RSL, PCP and GRASTA work to some extent in certain cases, though not as well as ReProCS. ReProCS operates by first approximately nullifying L_t , i.e. computing y_t as in (6), and then recovering S_t by exploiting its sparsity. iRSL and RSL also compute y_t the same way, but they directly use y_t to detect or soft-detect (and down-weight) the support of S_t by thresholding. Recall that y_t can be rewritten as $y_t = S_t + (-\hat{P}_{t-1}\hat{P}'_{t-1}S_t) + \beta_t$. As the support size of S_t increases, the interference due to $(-\hat{P}_{t-1}\hat{P}'_{t-1}S_t)$ becomes larger, resulting in wrong estimates of S_t . For the same reason, direct thresholding is also difficult when some entries of S_t are small while others are not. Adapted-iSVD is our adaptation of iSVD [13] in which we use the approach of iRSL described above to provide the outlier locations to iSVD (iSVD is an algorithm for recursive PCA with missing entries or what can be called recursive low-rank matrix completion). It fills in the corrupted locations of L_t by imposing that L_t lies in $\text{range}(\hat{P}_{t-1})$. We used a threshold of $0.5 \min_{i \in T_t} |(S_t)_i|$ for both iRSL and adapted-iSVD (we also tried various other options for thresholds but with similar results). Since adapted-iSVD and iRSL are recursive methods, a wrong \hat{S}_t , in turn, results in wrong subspace updates, thus also causing β_t to become large and finally causing the error to blow up.

RSL works to some extent for larger support size of S_t 's but fails when the magnitude of the nonzero S_t 's is small. PCP fails since the support sets are generated in a highly correlated fashion and the support sizes are large (resulting in the matrix S_t being quite rank deficient also). GRASTA [31] is a recent recursive method from 2012. It was implemented using code

⁵We use the Accelerated Proximal Gradient algorithm [47] and Inexact ALM algorithm [48] (designed for large scale problems) to solve PCP (1). The code is available at http://perception.csl.uiuc.edu/matrix-rank/sample_code.html.

⁶The code of RSL is available at <http://www.salleurl.edu/~fitorre/papers/rpca/rpca.zip>.

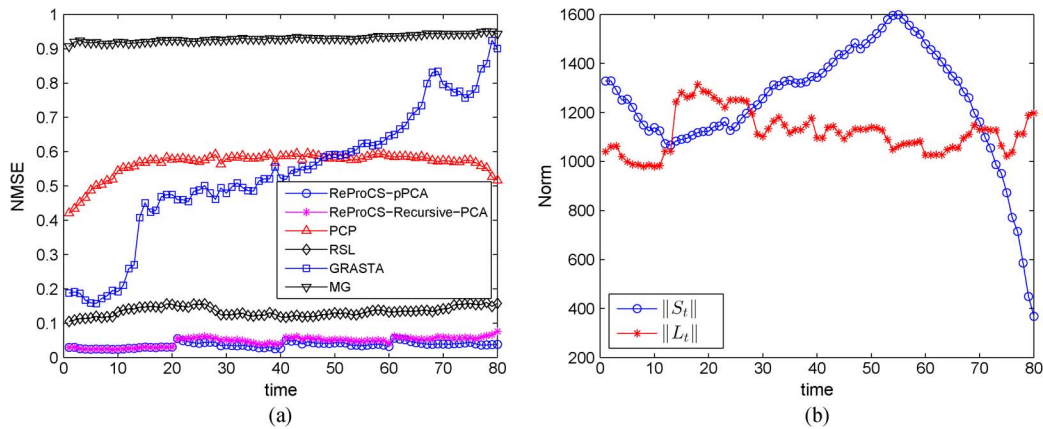


Fig. 2. Experiments on partly simulated video. (a) Normalized mean squared error in recovering S_t for realizations. (b) Comparison of $\|S_t\|_2$ and $\|L_t\|_2$ for one realization. MG refers to the batch algorithm of [34], [35] implemented using code provided by the authors. There was not enough information in the papers or in the code to successfully implement the recursive algorithm.

posted at <https://sites.google.com/site/hejunzz/grasta>. We tried two versions of GRASTA: the demo code as it is and the demo code modified so that it used as much information as ReProCS used (i.e. we used all available frames for training instead of just 100; we used all measurements instead of just 20% randomly selected pixels; and we used \hat{r} returned by ReProCS as the rank input instead of using rank = 5 always). In this paper, we show the latter case. Both experiments are shown on our supplementary material page <http://www.ece.iastate.edu/~hanguo/PracReProCS.html>.

Partly Simulated Video: Lake Video With Simulated Foreground: In the comparisons shown next, we only compare with PCP, RSL and GRASTA. To address a reviewer comment, we also compare with the batch algorithm of [34], [35] (referred to as MG in the figures) implemented using code provided by the authors. There was not enough information in the papers or in the code to successfully implement the recursive algorithm.

We implemented ReProCS (Algorithm 1 and Algorithm 3) with $b = 95$ since the videos are only approximately low-rank and we used $q = 1$ since the magnitude of S_t is not small compared to that of L_t . The performance of both ReProCS-pPCA (Algorithm 1) and ReProCS-recursive-PCA (Algorithm 1) was very similar. Results with using the latter are shown in Fig. 6.

We used the lake sequence described earlier to serve as a real background sequence. Foreground consisting of a rectangular moving object was overlaid on top of it using (3). The use of a real background sequence allows us to evaluate performance for data that only approximately satisfies the low-dimensional and slow subspace change assumptions. The use of the simulated foreground allows us to control its intensity so that the resulting S_t is small or of the same order as L_t (making it a difficult sequence), see Fig. 2(b).

The foreground F_t was generated as follows. For $1 \leq t \leq t_{\text{train}}$, $F_t = 0$. For $t > t_{\text{train}}$, F_t consists of a 45×25 moving block whose centroid moves horizontally according to a constant velocity model with small random acceleration [49, Example V.B.2]. To be precise, let p_t be the horizontal location of the block's centroid at time t , let v_t denote its horizontal velocity. Then $g_t := \begin{bmatrix} p_t \\ v_t \end{bmatrix}$ satisfies $g_t = Gg_{t-1} + \begin{bmatrix} 0 \\ n_t \end{bmatrix}$ where $G := \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$ and n_t is a zero mean truncated Gaussian with

variance Q and with $-2\sqrt{Q} < |n_t| < 2\sqrt{Q}$. The nonzero pixels' intensity is i.i.d. over time and space and distributed as uniform(b_1, b_2), i.e. $(F_t)_i \sim \text{uniform}(b_1, b_2)$ for $i \in T_t$. In our experiments, we generated the data with $t_{\text{train}} = 1420$, $b_1 = 170$, $b_2 = 230$, $p_{t_0+1} = 27$, $v_{t_0+1} = 0.5$ and $Q = 0.02$. With these values of b_1, b_2 , as can be seen from Fig. 2(b), $\|S_t\|_2$ is roughly equal or smaller than $\|L_t\|_2$ making it a difficult sequence. Since it is not much smaller, ReProCS used $q = 1$; since background data is approximately low-rank it used $b = 95$.

We generated 50 realizations of the video sequence using these parameters and compared all the algorithms to estimate S_t, L_t and then the foreground and the background sequences. We show comparisons of the normalized mean squared error (NMSE) in recovering S_t in Fig. 2(a). Visual comparisons of both foreground and background recovery for one realization are shown in Fig. 3. The recovered foreground image is shown as a white-black image showing the foreground support: pixels in the support estimate are white. PCP gives large error for this sequence since the object moves in a highly correlated fashion and occupies a large part of the image. GRASTA also does not work. RSL is able to recover a large part of the object correctly, however it also recovers many more extras than ReProCS. The reason is that the magnitude of the nonzero entries of S_t is quite small (recall that $(S_t)_i = (F_t - B_t)_i$ for $i \in T_t$) and is such that $\|L_t\|_2$ is about as large as $\|S_t\|_2$ or sometimes larger (see Fig. 2(b)).

Real Video Sequences: Next we show comparisons on two real video sequences. These are originally taken from http://perception.i2r.a-star.edu.sg/bk_model/bk_index.html and <http://research.microsoft.com/en-us/um/people/jckrumm/wallflower/testimages.htm>, respectively. The first is the curtain sequence described earlier. For $t > 1755$, in the foreground, a person with a black shirt walks in, writes on the board and then walks out, then a second person with a white shirt does the same and then a third person with a white shirt does the same. This video is challenging because (i) the white shirt color and the curtains' color is quite similar, making the corresponding S_t small in magnitude; and (ii) because the background variations are quite large while the foreground person moves slowly. As can be seen from Fig. 4, ReProCS's performance is significantly better than that of the other algorithms for both




Fig. 3. Original video at $t = t_{\text{train}} + 30, 60, 70$ and its foreground (fg) and background (bg) layer recovery results using ReProCS (ReProCS-pCA) and other algorithms. MG refers to the batch algorithm of [34], [35] implemented using code provided by the authors. There was not enough information in the papers or in the code to successfully implement the recursive algorithm. For fg, we only show the fg support in white for ease of display.




Fig. 4. Original video sequence at $t = t_{\text{train}} + 60, 120, 199, 475, 1148$ and its foreground (fg) and background (bg) layer recovery results using ReProCS (ReProCS-pCA) and other algorithms. For fg, we only show the fg support in white for ease of display.

foreground and background recovery. This is most easily seen from the recovered background images. One or more frames of the background recovered by PCP, RSL and GRASTA contains the person, while none of the ReProCS ones does.

The second sequence consists of a person entering a room containing a computer monitor that contains a white moving region. Background changes due to lighting variations and due to the computer monitor. The person moving in the foreground occupies a very large part of the image, so this is an example of a sequence in which the use of weighted ℓ_1 is essential (the support size is too large for simple ℓ_1 to work). As can be seen from Fig. 5, for most frames, ReProCS is able to recover the person correctly. However, for the last few frames which consist of the person in a white shirt in front of the white part of the screen, the resulting S_t is too small even for ReProCS to correctly recover. The same is true for the other algorithms. Videos of all above experiments and of a few others are posted at <http://www.ece.ias-tate.edu/~hanguo/PracReProCS.html>.

Time Comparisons: The time comparisons are shown in Table II. In terms of speed, GRASTA is the fastest even though its performance is much worse. ReProCS is the second fastest. We expect that ReProCS can be speeded up by using mex files

(C/C++ code) for the subspace update step. PCP and RSL are slower because they jointly process the entire image sequence. Moreover, ReProCS and GRASTA have the advantage of being recursive methods, i.e. the foreground/background recovery is available as soon as a new frame appears while PCP or RSL need to wait for the entire image sequence.

Compressive ReProCS: Comparisons for Simulated Video: We compare compressive ReProCS with SpaRCS [20] which is a batch algorithm for undersampled robust PCA/separation of sparse and low-dimensional parts (its code is downloaded from <http://www.ece.rice.edu/~aew2/sparcs.html>). No code is available for most of the other compressive batch robust PCA algorithms such as [24], [25]. SpaRCS is a greedy approach that combines ideas from CoSaMP [50] for sparse recovery and ADMiRA [51] for matrix completion. The comparison is done for compressive measurements of the lake sequence with foreground simulated as explained earlier. The matrix $B = A$ is $m \times n$ random Gaussian with $m = 0.7n$. Recall that $n = 6480$. The SpaRCS code required the background data rank and foreground sparsity as inputs. For rank, we used \hat{r} returned by ReProCS, for sparsity we used the true size of the simulated foreground. As can be seen from Fig. 7, SpaRCS does not work




Fig. 5. Original video sequence at $t = t_{\text{train}} + 42, 44, 52$ and its foreground (fg) and background (bg) layer recovery results using ReProCS (ReProCS-pCA) and other algorithms. For fg, we only show the fg support in white for ease of display.




Fig. 6. Foreground layer estimated by ReProCS-Recursive-PCA for the lake, curtain and person videos shown in Figs. 3, 4 and 5. As can be seen the recovery performance is very similar to that of ReProCS-pPCA (Algorithm 1). (a) $t = 30, 60, 70$; (b) $t = 60, 120, 475$; (c) $t = 42, 44, 52$.

TABLE II

COMPARISON OF SPEED OF DIFFERENT ALGORITHMS. EXPERIMENTS WERE DONE ON A 64 bit WINDOWS 8 LAPTOP WITH 2.40 GHZ I7 CPU AND 8G RAM. SEQUENCE LENGTH REFERS TO THE LENGTH OF SEQUENCE FOR TRAINING PLUS THE LENGTH OF SEQUENCE FOR SEPARATION. FOR REPROCS AND GRASTA, THE TIME IS SHOWN AS TRAINING TIME+RECOVERY TIME

DataSet	Image Size	Sequence Length	ReProCS-pPCA	ReProCS-Recursive-PCA	PCP	RSL	GRASTA
Lake	72×90	1420 + 80	2.99 + 19.97 sec	2.99 + 19.43 sec	245.03 sec	213.36 sec	39.47 + 0.42 sec
Curtain	64×80	1755 + 1209	4.37 + 159.02 sec	4.37 + 157.21 sec	1079.59 sec	643.98 sec	40.01 + 5.13 sec
Person	120×160	200 + 52	0.46 + 42.43 sec	0.46 + 41.91 sec	27.72 sec	121.31 sec	13.80 + 0.64 sec




Fig. 7. Original video frames at $t = t_{\text{train}} + 30, 60, 70$ and foreground layer recovery by ReProCS and SparCS.

while compressive ReProCS is able to recover S_t fairly accurately, though of course the errors are larger than in the full sampled case. All experiments shown in [20] are for very slow changing backgrounds and for foregrounds with very small support sizes, while neither is true for our data.

VII. CONCLUSIONS AND FUTURE WORK

This work designed and evaluated Prac-ReProCS which is a practically usable modification of its theoretical counterpart that was studied in our earlier work [27]–[29]. We showed that Prac-ReProCS has excellent performance for both simulated data and for a real application (foreground-background separation in videos) and the performance is better than many of the

state-of-the-art algorithms from recent work. Moreover, most of the assumptions used to obtain its guarantees are valid for real videos. Finally we also proposed and evaluated a compressive prac-ReProCS algorithm. In ongoing work, on one end, we are working on performance guarantees for compressive ReProCS [45] and on the other end, we are developing and evaluating a related approach for functional MRI. In fMRI, one is allowed to change the measurement matrix at each time. However if we replace $B = A$ by A_t in Section IV the compressive ReProCS algorithm does not apply because $A_t L_t$ is not low-dimensional [44].

APPENDIX

A. Detailed Discussion of Why ReProCS Works

Define the subspace estimation error as

$$\text{SE}(P, \hat{P}) := \left\| (I - \hat{P}\hat{P}')P \right\|_2$$

where both P and \hat{P} are *basis matrices*. Notice that this quantity is always between zero and one. It is equal to zero when $\text{range}(\hat{P})$ contains $\text{range}(P)$ and it is equal to one if $\hat{P}'P_i = 0$ for at least one column of P .

Recall that for $t \in [t_j, t_{j+1} - 1]$, $P_t = [P_{(j-1)}R_j \setminus P_{(j),\text{old}}, P_{(j),\text{new}}]$. Thus, L_t can be rewritten as

$$L_t = P_{(j-1)}a_{t,*} + P_{(j),\text{new}}a_{t,\text{new}}$$

where $a_{t,\text{new}} := P'_{(j),\text{new}}L_t$ and $a_{t,*} := P'_{(j-1)}L_t$.

Let $c := c_{\max}$ and $r_j := r_0 + jc$. We explain here the key idea of why ReProCS works [27], [29]. Assume the following hold besides the assumptions of Section II.

- 1) Subspace change is detected immediately, i.e. $\hat{t}_j = t_j$ and $c_{j,\text{new}}$ is known.
- 2) Pick a $\zeta \ll 1$. Assume that $\|L_t\|_2 \leq \gamma_*$ for a γ_* that satisfies $\gamma_* \leq 1/\sqrt{r_j\zeta}$. Since ζ is very small, γ_* can be very large.
- 3) Assume that $(t_{j+1} - t_j) \geq K\alpha$ for a K as defined below.
- 4) Assume the following model on the gradual increase of $a_{t,\text{new}}$: for $t \in [t_j + (k-1)\alpha, t_j + k\alpha - 1]$, $\|a_{t,\text{new}}\|_2 \leq v^{k-1}\gamma_{\text{new}}$ for a $1 < v \leq 1.2$ and $\gamma_{\text{new}} \ll \gamma_*$.
- 5) Assume that projection PCA “works” i.e. its estimates satisfy $\text{SE}(P_{(j),\text{new}}, \hat{P}_{(j),\text{new},k-1}) \leq 0.6^{k-1} + 0.4c\zeta$. The proof of this statement is long and complicated and is given in [27], [29].
- 6) Assume that projection PCA is done K times with K chosen so that $0.6^{K-1} + 0.4c\zeta \leq c\zeta$.

Assume that at $t = t_j - 1$, $\text{SE}(P_{(j-1)}, \hat{P}_{(j-1)}) \leq r_{j-1}\zeta \ll 1$. We will argue below that $\text{SE}(P_{(j)}, \hat{P}_{(j)}) \leq r_j\zeta$. Since $r_j \leq r_0 + Jc$ is small, this error is always small and bounded.

First consider a $t \in [t_j, t_j + \alpha)$. At this time, $\hat{P}_t = \hat{P}_{(j-1)}$. Thus,

$$\begin{aligned} \|\beta_t\|_2 &= \left\| \left(I - \hat{P}_{t-1} \hat{P}'_{t-1} \right) L_t \right\|_2 \\ &\leq \text{SE} \left(P_{(j-1)}, \hat{P}_{(j-1)} \right) \|a_{t,*}\|_2 + \|a_{t,\text{new}}\|_2 \\ &\leq (r_{j-1}\zeta)\gamma_* + \gamma_{\text{new}} \\ &\leq \sqrt{\zeta} + \gamma_{\text{new}}. \end{aligned} \quad (11)$$

By construction, $\sqrt{\zeta}$ is very small and hence the second term in the bound is the dominant one. By the slow subspace assumption $\gamma_{\text{new}} \ll \|S_t\|_2$. Recall that β_t is the “noise” seen by the sparse recovery step. The above shows that this noise is small compared to $\|S_t\|_2$. Moreover, using Lemma 2.2 and simple arguments [see [27, Lemma 6.6]], it can be shown that

$$\delta_s(\Phi_t) \leq \kappa_*^2 + r_{j-1}\zeta$$

is small. These two facts along with any RIP-based result for ℓ_1 minimization, e.g. [39], ensure that S_t is recovered accurately in this step. If the smallest nonzero entry of S_t is large enough, it is possible get a support threshold ω that ensures exact support recovery. Then, the LS step gives a very accurate final estimate of S_t and it allows us to get an exact expression for $e_t := S_t - \hat{S}_t$. Since $\hat{L}_t = M_t - \hat{S}_t$, this means that L_t is also recovered accurately and $e_t = \hat{L}_t - L_t$. This is then used to argue that p-PCA at $t = t_j + \alpha - 1$ “works”.

Next consider $t \in [t_j + (k-1)\alpha, t_j + k\alpha - 1]$. At this time, $\hat{P}_t = [\hat{P}_{(j-1)}, \hat{P}_{(j),\text{new},k-1}]$. Then, it is easy to see that

$$\begin{aligned} \|\beta_t\|_2 &= \left\| \left(I - \hat{P}_{t-1} \hat{P}'_{t-1} \right) L_t \right\|_2 \\ &\leq \text{SE} \left(P_{(j-1)}, \hat{P}_{(j-1)} \right) \|a_{t,*}\|_2 \\ &\quad + \text{SE} \left(P_{(j),\text{new}}, \hat{P}_{(j),\text{new},k-1} \right) \|a_{t,\text{new}}\|_2 \\ &\leq (r_{j-1}\zeta)\gamma_* + (0.6^{k-1} + 0.4c\zeta) v^{k-1}\gamma_{\text{new}} \\ &\leq \sqrt{\zeta} + 0.72^{k-1}\gamma_{\text{new}}. \end{aligned} \quad (12)$$

Ignoring the first term, in this interval, $\|\beta_t\|_2 \leq 0.72^{k-1}\gamma_{\text{new}}$, i.e. the noise seen by the sparse recovery step decreases exponentially with every p-PCA step. This, along with a bound on $\delta_s(\Phi_t)$ (this bound needs a more complicated argument than that for $k=1$, see [27, Lemma 6.6]), ensures that the recovery error of S_t , and hence also of $L_t = M_t - S_t$, decreases roughly exponentially with k . This is then used to argue that the p-PCA error also decays roughly exponentially with k .

Finally for $t \in [t_j + K\alpha, t_{j+1} - 1]$, because of the choice of K , we have that $\text{SE}(P_{(j),\text{new}}, \hat{P}_{(j),\text{new},K}) \leq c\zeta$. At this time, we set $\hat{P}_{(j)} = [\hat{P}_{(j-1)}, \hat{P}_{(j),\text{new},K}]$. Thus, $\text{SE}(P_{(j)}, \hat{P}_{(j)}) \leq \text{SE}(P_{(j-1)}, \hat{P}_{(j-1)}) + \text{SE}(P_{(j),\text{new}}, \hat{P}_{(j),\text{new},K}) \leq r_{j-1}\zeta + c\zeta = r_j\zeta$.

REFERENCES

- [1] C. Qiu and N. Vaswani, “Real-time robust principal components’ pursuit,” in *Proc. Allerton Conf. Commun., Contr. Comput.*, 2010, pp. 591–598.
- [2] C. Qiu and N. Vaswani, “Recursive sparse recovery in large but correlated noise,” in *Proc. Allerton Conf. Commun., Control, Comput.*, 2011, pp. 752–759.
- [3] H. Guo, C. Qiu, and N. Vaswani, “Practical reprocs for separating sparse and low-dimensional signal sequences from their sum—Part 1,” presented at the IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP), Florence, Italy, 2014.
- [4] F. De La Torre and M. J. Black, “A framework for robust subspace learning,” *Int. J. Comput. Vision*, vol. 54, pp. 117–142, 2003.
- [5] E. J. Candès, X. Li, Y. Ma, and J. Wright, “Robust principal component analysis?,” *J. ACM*, vol. 58, no. 3, 2011.
- [6] J. Wright and Y. Ma, “Dense error correction via l1-minimization,” *IEEE Trans. Inf. Theory*, vol. 56, no. 7, pp. 3540–3560, 2010.
- [7] I. T. Jolliffe, *Principal Component Analysis*, 2nd ed. New York, NY, USA: Springer, 2002.
- [8] J. Wright and Y. Ma, “Dense error correction via l1-minimization,” *IEEE Trans. Inf. Theory*, vol. 56, no. 7, pp. 3540–3560, Jul. 2010.
- [9] S. Roweis, “EM algorithms for PCA and SPCA,” *Adv. Neural Inf. Process. Syst.*, pp. 626–632, 1998.
- [10] T. Zhang and G. Lerman, “A novel M-estimator for robust PCA,” *J. Mach. Learn. Res.*, 2013 [Online]. Available: <http://arxiv.org/abs/1112.4863>, arXiv:1112.4863v3, to be published.
- [11] H. Xu, C. Caramanis, and S. Sanghavi, “Robust PCA via outlier pursuit,” *IEEE Trans. Inf. Theory*, vol. 58, no. 5, pp. 3047–3064, 2012.
- [12] M. McCoy and J. A. Tropp, “Two proposals for robust PCA using semidefinite programming,” *Electron. J. Statist.*, vol. 5, pp. 1123–1160, 2011.
- [13] M. Brand, “Incremental singular value decomposition of uncertain data with missing values,” in *Proc. Eur. Conf. Comput. Vision*, 2002, pp. 707–720.
- [14] D. Skocaj and A. Leonardis, “Weighted and robust incremental method for subspace learning,” in *Proc. IEEE Int. Conf. Comput. Vision (ICCV)*, Oct. 2003, vol. 2, pp. 1494–1501.
- [15] Y. Li, L. Xu, J. Morphet, and R. Jacobs, “An integrated algorithm of incremental and robust PCA,” in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, 2003, pp. 245–248.
- [16] V. Chandrasekaran, S. Sanghavi, P. A. Parrilo, and A. S. Willsky, “Rank-sparsity incoherence for matrix decomposition,” *SIAM J. Optim.*, vol. 21, pp. 572–596, 2011.
- [17] M. B. McCoy and J. A. Tropp, “Sharp recovery bounds for convex deconvolution, with applications,” *J. Found. Comput. Math.*, 2012 [Online]. Available: <http://arxiv.org/abs/1205.1580>, arXiv:1205.1580, to be published.
- [18] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, “The convex geometry of linear inverse problems,” *Found. Comput. Math.*, no. 6, pp. 805–849, 2012.
- [19] Y. Hu, S. Goud, and M. Jacob, “A fast majorize-minimize algorithm for the recovery of sparse and low-rank matrices,” *IEEE Trans. Image Process.*, vol. 21, no. 2, pp. 742–753, Feb. 2012.
- [20] A. E. Waters, A. C. Sankaranarayanan, and R. G. Baraniuk, “Spars: Recovering low-rank and sparse matrices from compressive measurements,” in *Proc. Neural Inf. Process. Syst. (NIPS)*, 2011, pp. 1089–1097.
- [21] E. Richard, P.-A. Savalle, and N. Vayatis, “Estimation of simultaneously sparse and low rank matrices,” in *Proc. 29th Int. Conf. Mach. Learn. (ICML 2012)* [Online]. Available: <http://arxiv.org/abs/1206.6474>, arXiv:1206.6474.
- [22] D. Hsu, S. M. Kakade, and T. Zhang, “Robust matrix decomposition with sparse corruptions,” *IEEE Trans. Inf. Theory*, vol. 57, no. 11, pp. 7221–7234, 2011.

- [23] M. Mardani, G. Mateos, and G. B. Giannakis, "Recovery of low-rank plus compressed sparse matrices with application to unveiling traffic anomalies," *IEEE Trans. Inf. Theory*, vol. 59, no. 8, pp. 5186–5205, 2013.
- [24] J. Wright, A. Ganesh, K. Min, and Y. Ma, "Compressive principal component pursuit," *Inf. Inference*, vol. 2, no. 1, pp. 32–68, 2013.
- [25] A. Ganesh, K. Min, J. Wright, and Y. Ma, "Principal component pursuit with reduced linear measurements," in *Proc. IEEE Int. Symp. Inf. Theory Process. (ISIT)*, 2012, pp. 1281–1285.
- [26] M. Tao and X. Yuan, "Recovering low-rank and sparse components of matrices from incomplete and noisy observations," *SIAM J. Optim.*, vol. 21, no. 1, pp. 57–81, 2011.
- [27] C. Qiu, N. Vaswani, B. Lois, and L. Hogben, "Recursive robust PCA or recursive sparse recovery in large but structured noise," *IEEE Trans. Inf. Theory*, vol. 60, no. 8, pp. 5007–5039, 2014.
- [28] C. Qiu and N. Vaswani, "Recursive sparse recovery in large but structured noise—Part 2," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2013, pp. 864–868.
- [29] "Blinded title," in *Double Blind Conf.*, 2014, submitted for publication.
- [30] N. Vaswani and W. Lu, "Modified-CS: Modifying compressive sensing for problems with partially known support," *IEEE Trans. Signal Process.*, vol. 59, no. 9, pp. 4595–4607, Sep. 2010.
- [31] J. He, L. Balzano, and A. Sztam, "Incremental gradient on the Grassmannian for online foreground and background separation in subsampled video," in *Proc. IEEE Conf. Comp. Vis. Pattern Recog. (CVPR)*, 2012, pp. 1568–1575.
- [32] J. Feng, H. Xu, and S. Yan, "Online robust PCA via stochastic optimization," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2013, pp. 404–412.
- [33] J. Feng, H. Xu, S. Mannor, and S. Yan, "Online pca for contaminated data," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2013, pp. 764–772.
- [34] G. Mateos and G. Giannakis, "Robust PCA as bilinear decomposition with outlier-sparsity regularization," *IEEE Trans. Signal Process.*, vol. 60, no. 10, pp. 5176–5190, Oct. 2012.
- [35] M. Mardani, G. Mateos, and G. Giannakis, "Dynamic anomalography: Tracking network anomalies via sparsity and low rank," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 1, pp. 50–66, Feb. 2013.
- [36] K. Jia, T.-H. Chan, and Y. Ma, "Robust and practical face recognition via structures sparsity," in *Proc. Eur. Conf. Comp. Vis. (ECCV)*, 2012, pp. 331–344.
- [37] E. Candes and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, Dec. 2005.
- [38] E. J. Candès and B. Recht, "Exact matrix completion via convex optimization," *Commun. ACM*, vol. 55, no. 6, pp. 111–119, 2012.
- [39] E. Candes, "The restricted isometry property and its implications for compressed sensing," *Compte Rendus de l'Acad. des Sci., Paris, Ser. I*, pp. 589–592, 2008.
- [40] S. Chen, D. Donoho, and M. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Scientif. Comput.*, vol. 20, pp. 33–61, 1998.
- [41] E. Candes and T. Tao, "The Dantzig selector: Statistical estimation when p is much larger than n ," *Ann. Statist.*, vol. 35, no. 6, pp. 2313–2351, 2007.
- [42] A. Khajehnejad, W. Xu, A. Avestimehr, and B. Hassibi, "Weighted ℓ_1 minimization for sparse recovery with prior information," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2009, pp. 483–487.
- [43] M. P. Friedlander, H. Mansour, R. Saab, and O. Yilmaz, "Recovering compressively sampled signals using partial support information," *IEEE Trans. Inf. Theory*, vol. 58, no. 2, pp. 1122–1134, 2012.
- [44] J. Zhan and N. Vaswani, "Separating sparse and low-dimensional signal sequence from time-varying undersampled projections of their sums," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, 2013, pp. 5905–5909.
- [45] B. Lois, N. Vaswani, and C. Qiu, "Performance guarantees for under-sampled recursive sparse recovery in large but structured noise," in *Proc. GlobalSIP*, 2013, pp. 1061–1064.
- [46] J. Yang and Y. Zhang, "Alternating direction algorithms for ℓ_1 problems in compressive sensing," Rice Univ., Houston, TX, USA, Tech. Rep., Jun. 2010.
- [47] Z. Lin, A. Ganesh, J. Wright, L. Wu, M. Chen, and Y. Ma, "Fast convex optimization algorithms for exact recovery of a corrupted low-rank matrix," Univ. Illinois at Urbana-Champaign, Tech. Rep., Aug. 2009.
- [48] Z. Lin, M. Chen, and Y. Ma, "Alternating direction algorithms for ℓ_1 problems in compressive sensing," Univ. Illinois at Urbana-Champaign, Tech. Rep., Nov. 2009.
- [49] H. Vincent Poor, *An Introduction to Signal Detection and Estimation*, 2nd ed. New York, NY, USA: Springer, 1994.
- [50] D. Needell and J. A. Tropp, "Cosamp: Iterative signal recovery from incomplete and inaccurate samples," *Appl. Comp. Harmon. Anal.*, vol. 26, no. 3, pp. 301–321, May 2009.
- [51] K. Lee and Y. Bresler, "ADMIRA: Atomic decomposition for minimum rank approximation," *IEEE Trans. Inf. Theory*, vol. 56, no. 9, pp. 4402–4416, Sep. 2010.



Han Guo received the B.S. degree from Northwestern Polytechnical University in China in 2012.

He is currently a Ph.D. student in Electrical and Computer Engineering at Iowa State University, Ames. His research interests are in sparse recovery, robust PCA, and video analysis.



Chenlu Qiu received a B.S. from Southeast University in China in 2006 in Information Engineering and a Ph.D. from Iowa State University in 2013 in Electrical Engineering. She is currently with the Traffic Management Research Institute of the Ministry of Public Security, China. Her research interests include robust PCA and video analysis.



Namrata Vaswani received the B.Tech. degree from the Indian Institute of Technology (IIT), Delhi, in 1999 and the Ph.D. degree from the University of Maryland, College Park, in 2004, both in electrical engineering.

During 2004–2005, she was a Research Scientist at Georgia Tech. Since fall 2005, she has been with the Iowa State University (ISU), Ames, where she is currently an Associate Professor of Electrical and Computer Engineering. She held the Harpole-Pentair Assistant Professorship at ISU during 2008–2009 and received the Early Career Engineering Faculty Research Award in 2014. Her research interests are in signal and information processing for problems motivated by big-data and bio-imaging applications. Her current work focuses on recursive sparse recovery, robust PCA, matrix completion and applications in video and medical imaging.

Dr. Vaswani served an Associate Editor for IEEE TRANSACTIONS ON SIGNAL PROCESSING from 2009 to 2013.