Iterative Grassmannian Optimization for Robust Image Alignment

Jun He[a,*], Dejiao Zhang[a], Laura Balzano[b], Tao Tao[a]

[a]*School of Electronic and Information Engineering, Nanjing University of Information Science and Technology, Nanjing, 210044, China*
[b]*Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, USA*

Abstract

Robust high-dimensional data processing has witnessed an exciting development in recent years. Theoretical results have shown that it is possible using convex programming to optimize data fit to a low-rank component plus a sparse outlier component. This problem is also known as Robust PCA, and it has found application in many areas of computer vision. In image and video processing and face recognition, the opportunity to process massive image databases is emerging as people upload photo and video data online in unprecedented volumes. However, data quality and consistency is not controlled in any way, and the massiveness of the data poses a serious computational challenge. In this paper we present t-GRASTA, or "Transformed GRASTA (Grassmannian Robust Adaptive Subspace Tracking Algorithm)". t-GRASTA iteratively performs incremental gradient descent constrained to the Grassmann manifold of subspaces in order to simultaneously estimate three components of a decomposition of a collection of images: a low-rank subspace, a sparse part of occlusions and foreground objects, and a transformation such as rotation or translation of the image. We show that t-GRASTA is 4× faster than state-of-the-art algorithms, has half the memory requirement, and can achieve alignment for face images as well as jittered camera surveillance images.

*Corresponding author. Tel. +86 13913873052
  Email addresses: jhe@nuist.edu.cn (Jun He), dejiaozhang@gmail.com (Dejiao Zhang), girasole@umich.edu (Laura Balzano)
  URL: http://sites.google.com/site/hejunzz/ (Jun He)

---

1. INTRODUCTION

With the explosion of image and video capture, both for surveillance and personal enjoyment, and the ease of putting these data online, we are seeing photo databases grow at unprecedented rates. On record we know that in July 2010, Facebook had 100 million photo uploads per day [1] and Instagram had a database of 400 million photos as of the end of 2011, with 60 uploads per second [2]; since then both of these databases have certainly grown immensely. In 2010, there were an estimated minimum 10,000 surveillance cameras in the city of Chicago and in 2002 an estimated 500,000 in London [3, 4].

These enormous collections pose both an opportunity and a challenge for image processing and face recognition: The opportunity is that with so much data, it should be possible to assist users in tagging photos, searching the image database, and detecting unusual activity or anomalies. The challenge is that the data are not controlled in any way so as to ensure data quality and consistency across photos, and the massiveness of the data poses a serious computational challenge.

In video surveillance, many recently proposed algorithms model the foreground and background separation problem as one of "Robust PCA"– decomposing the scene as the sum of a low-rank matrix of background, which represents the global appearance and illumination of the scene, and a sparse matrix of moving foreground objects [5, 6, 7, 8, 9]. These popular algorithms and models work very well for a stationary camera. However, in the case of camera jitter, the background is no longer low-rank, and this is problematic for Robust PCA methods [10, 11, 12]. Robustly and efficiently detecting moving objects from an unstable camera is a challenging problem, since we need to accurately estimate both the background and the transformation of each frame. Fig. 1 shows that for a video sequence generated by a simulated unstable camera, GRASTA [13, 6] (Grassmannian Robust Adaptive Subspace Tracking Algorithm) fails to do the separation, but the approach we propose here, t-GRASTA, can successfully separate the background and moving objects despite camera jitter.

Further recent work has extended the Robust PCA model to that of the "Transformed Low-Rank + Sparse" model for face images with occlu-

Figure 1: Video background and foreground separation by t-GRASTA despite camera jitter. 1^{st} row: misaligned video frames by simulating camera jitters; 2^{nd} row: images aligned by t-GRASTA; 3^{rd} row: background recovered by t-GRASTA; 4^{th} row: foreground separated by t-GRASTA; 5^{th} row: background recovered by GRASTA; 6^{th} row: foreground separated by GRASTA.

sions that have come under transformations such as translations and rotations [14, 15, 16, 17]. Without the transformations, this can be posed as a convex optimization problem and therefore convex programming methods can be used to tackle such a problem. In RASL [15] (Robust Alignment by Sparse and Low-Rank decomposition), the authors posed the problem with transformations as well, and though it is no longer convex it can be linearized in each iteration and proven to reach a local minimum.

Though the convex programming methods used in [15] are polynomial in the size of the problem, that complexity can still be too demanding for very large databases of images. We propose Transformed GRASTA, or t-GRASTA for short, to tackle this optimization with an incremental or online optimization technique. The benefit of this approach is three-fold: First, it will improve speeds of image alignment both in batch mode or in online mode, as we show in Section 3. Second, the memory requirement is small, which makes alignment for very large databases realistic, since t-GRASTA only needs to maintain low-rank subspaces throughout the alignment process. Finally, the proposed online version of t-GRASTA allows for alignment and occlusion removal on images as they are uploaded to the database, which is

3

especially useful in video processing scenarios.

1.1. Robust Image Alignment

The problem of robust image alignment arises regularly in real data, as large illumination variations and gross pixel corruptions or partial occlusions often occur, such as sunglasses or a scarf for a human subject. The classic batch image alignment approaches, such as congealing [18, 19] or least squares congealing algorithms [20, 21] cannot simultaneously handle such severe conditions, causing the alignment task to fail.

With the breakthrough of convex relaxation theory applied to decomposing matrices into a sum of low-rank and sparse matrices [22, 5], the recently proposed algorithm "Robust Alignment by Sparse and Low-rank decomposition," or RASL [15], poses the robust image alignment problem as a transformed version of Robust PCA. The transformed batch of images can be decomposed as the sum of a low-rank matrix of recovered aligned images and a sparse matrix of errors. RASL seeks the optimal domain transformations while trying to minimize the rank of the matrix of the vectorized and stacked aligned images and while keeping the gross errors sparse. While the rank minimization and ℓ^0 minimization can be relaxed to their convex surrogates– minimize the corresponding nuclear norm $\|\|\|_*$ and ℓ^1 norm $\|\|\|_1$– the relaxed problem (1) is still highly non-linear due to the complicated domain transformation.

$$\min_{A,E,\tau} \|A\|_* + \lambda\|E\|_1 \quad s.t. \ D \circ \tau = A + E \qquad (1)$$

Here, $D \in \mathbb{R}^{n \times N}$ represents the data (n pixels per each of N images), $A \in \mathbb{R}^{n \times N}$ is the low-rank component, $E \in \mathbb{R}^{n \times N}$ is the sparse additive component, and τ are the transformations. RASL proposes to tackle this difficult optimization problem by iteratively locally linearizing the non-linear image transformation $D \circ (\tau + \triangle\tau) \approx D \circ \tau + \sum_{i=1}^{n} J_i \triangle\tau_i \epsilon_i^T$, where J_i is the Jacobian of image i with respect to transformation i; then in each iteration the linearized problem is convex. The authors have shown that RASL works perfectly well for batch aligning the linearly correlated images despite large illumination variations and occlusions.

In order to improve the scalability of robust image alignment for massive image datasets, [23] proposes an efficient ALM-based (Augmented Lagrange Multiplier-based) iterative convex optimization algorithm ORIA (Online Robust Image Alignment) for online alignment of the input images. Though the

4

proposed approach can scale to large image datasets, it requires the subspace of the aligned images as a prior, and for this it uses RASL to train the initial aligned subspace. Once the input images cannot be well aligned by the current subspace, the authors use an heuristic method to update the basis. In contrast, with t-GRASTA we include the subspace in the cost function, and update the subspace using a gradient geodesic step on the Grassmannian, as in [6, 24]. We discuss this in more detail in the next section.

1.2. Online Robust Subspace Learning

Subspace learning has been an area important to signal processing for a few decades. There are many applications in which one must track signal and noise subspaces, from computer vision to communications and radar, and a survey of the related work can be found in [25, 26].

The GROUSE algorithm, or "Grassmannian Rank-One Update Subspace Estimation," is an online subspace estimation algorithm that can track changing subspaces in the presence of Gaussian noise and missing entries [24]. GROUSE was developed as an online variant of low-rank matrix completion algorithms. It uses incremental gradient methods that have been receiving extensive attention in the optimization community [27]. However, GROUSE is not robust to gross outliers, and the follow-up algorithm GRASTA [13, 6], can estimate a changing low-rank subspace as well as identify and subtract outliers. Still problematic is that, as we showed in Fig. 1, even GRASTA cannot handle camera jitter. Our algorithm includes the estimation of transformations in order to align frames first before separating foreground and background.

2. ROBUST IMAGE ALIGNMENT VIA ITERATIVE ONLINE SUBSPACE LEARNING

2.1. Model

2.1.1. Batch mode

In order to robustly align the set of linearly correlated images despite sparse outliers, we consider the following matrix factorization model (2) where the low-rank matrix U has orthonormal columns that span the low-dimensional subspace of the well-aligned images.

$$\min_{U,W,E,\tau} \quad \|E\|_1 \tag{2}$$
$$s.t. \quad D \circ \tau = UW + E$$
$$U \in \mathcal{G}(d,n)$$

We have replaced the variable A with the product of two smaller matrices UW, and the orthonormal columns of $U \in \mathbb{R}^{n \times d}$ span the low-rank subspace of the images. The set of all subspaces of $\mathbb{R}^n$ of fixed dimension d is called the Grassmannian, which is a compact Riemannian manifold and is denoted by $\mathcal{G}(d,n)$. In this optimization model, U is constrained to the Grassmannian $\mathcal{G}(d,n)$. Though problem (2) can not be directly solved [15] due to the nonlinearity of image transformation, if the misalignments are not too large, by locally linearly approximating the image transformation $D \circ (\tau + \triangle \tau) \approx D \circ \tau + \sum_{i=1}^{N} J_i \triangle \tau_i \epsilon_i^T$, the iterative model (3) can work well as a practical approach.

$$\min_{U^k,W,E,\triangle\tau} \quad \|E\|_1 \tag{3}$$
$$s.t. \quad D \circ \tau^k + \sum_{i=1}^{N} J_i^k \triangle \tau_i \epsilon_i^T = U^k W + E$$
$$U^k \in \mathcal{G}(d^k,n)$$

At algorithm iteration k, $\tau^k \doteq [\tau_1^k|,\ldots,|\tau_N^k]$ are the current estimated transformations at iteration k, J_i^k is the Jacobian of the i-th image with respect to the transformation τ_i^k, and $\{\epsilon_i\}$ denotes the standard basis for $\mathbb{R}^n$. Note, at different iterations the subspace may have different dimensions, i.e. U^k is constrained on different Grassmannian $\mathcal{G}(d^k,n)$.

At each iteration of the iterative model (3), we consider this optimization problem as the subspace learning problem. That is, our goal is to robustly estimate the low-dimensional subspace U^k which best represents the locally transformed images $D \circ \tau^k + \sum_{i=1}^{N} J_i^k \triangle \tau_i$ despite sparse outliers E. In order to solve this subspace learning problem both efficiently with regards to both computation and memory, we propose to learn U^k at each iteration k in model (3) via the online robust subspace learning approach [6].

2.1.2. Online mode

In order to perform online video processing tasks, for example video stabilization, it is desirable to design an efficient approach that can handle image misalignment frame by frame. As in the previous discussion regarding batch mode processing, for each video frame I, we may model the ℓ^1 minimization problem as follows:

$$\min_{U,w,e,\tau} \quad \|e\|_1 \tag{4}$$
$$s.t. \quad I \circ \tau = Uw + e$$
$$U \in \mathcal{G}(d,n)$$

Note that with the constraint $I \circ \tau = Uw + e$ in the above minimization problem, we suppose for each frame the transformed image is well aligned to the low-rank subspace U. However, due to the nonlinear geometric transform $I \circ \tau$, directly exploiting online subspace learning techniques [24, 6] is not possible.

Here we approach this as a manifold learning problem, supposing that the low-dimensional image subspace under nonlinear transformations forms a nonlinear manifold. We propose to learn the manifold approximately using a union of subspaces model U^ℓ, $\ell = 1, \ldots, L$. The basic idea is illustrated in Fig. 2, and the locally linearized model for the nonlinear problem (4) is as follows:

$$\min_{w,e,\triangle\tau} \quad \|e\|_1 \tag{5}$$
$$s.t. \quad I \circ \tau^\ell + J^\ell \triangle \tau = U^\ell w + e \ .$$
$$U^\ell \in \mathcal{G}(d^\ell,n)$$

Intuitively, from Fig. 2, it is reasonable to think that the initial misaligned image sequence should be high rank; then after iteratively approximating the nonlinear transform with a locally linear approximation, the rank of the new subspaces U^ℓ, $\ell = 1, \ldots, L$, should be decreasing as the images become more and more aligned. Then for each misaligned image I and the unknown transformation τ, we iteratively update the union of subspaces U^ℓ, $\ell = 1, \ldots, L$, and estimate the transformation τ. Details of the online mode of t-GRASTA will be discussed in Section 2.4.2

Figure 2: The illustration of iteratively approximating the nonlinear image manifold using a union of subspaces.

The use of a union of subspaces U^ℓ, $\ell = 1, \ldots, L$, to approximate the nonlinear manifold is a crucial innovation for this fully online model. Though we use the symbols U^k and U^ℓ in both the batch mode and the online mode, they have two different interpretations. For batch mode, U^k is the iteratively learned aligned subspace in each iteration; while for online mode, U^ℓ, $\ell = 1, \ldots, L$, is a collection of subspaces which are used for approximating the nonlinear transform, and they are updated iteratively for each video frame.

2.2. ADMM Solver for the Locally Linearized Problem

Whether operating in batch mode or online mode, the key problem is how to quantify the subspace error robustly for the locally linearized problem. Considering batch mode, at iteration k, given the i-th image I_i, its estimate of transformation τ_i^k, the Jacobian J_i^k, and the current estimate of U_t^k, we use the ℓ^1 norm as follows:

$$F(S; t, k) = \min_{w, \triangle\tau} \|U_t^k w - (I_i \circ \tau_i^k + J_i^k \triangle\tau)\|_1 \tag{6}$$

With U_t^k known (or estimated, but fixed), this ℓ^1 minimization problem is a variation of the least absolute deviations problem, which can be solved efficiently by ADMM (Alternating Direction Method of Multipliers) [28]. We rewrite the right hand of (6) as the equivalent constrained problem by introducing a sparse outlier vector e:

$$\min_{w, e, \triangle\tau} \quad \|e\|_1 \tag{7}$$
$$s.t. \quad I_i \circ \tau_i^k + J_i^k \triangle\tau = U_t^k w + e .$$

The augmented Lagrangian of problem (7) is

$$\mathcal{L}(U_t^k, w, e, \triangle\tau, \lambda) = \|e\|_1 \quad + \quad \lambda^T h(w, e, \triangle\tau)$$
$$+ \quad \frac{\mu}{2}\|h(w, e, \triangle\tau)\|_2^2 \tag{8}$$

where $h(w, e, \triangle\tau) = U_t^k w + e - I_i \circ \tau_i^k - J_i^k \triangle\tau$, and $\lambda \in \mathbb{R}^n$ is the Lagrange multiplier or dual vector.

Given the current estimated subspace U_t^k, transformation parameter τ_i^k, and the Jacobian matrix J_i^k with respect to the i-th image I_i, the optimal $(w^*, e^*, \triangle\tau^*, \lambda^*)$ can be computed by the ADMM approach as follows:

$$\begin{cases} \triangle\tau^{p+1} = (J_i^k J_i^{kT})^{-1} J_i^{kT}(U_t^k w^p + e^p - I_i \circ \tau_i^k + \frac{1}{\mu}\lambda^p) \\ w^{p+1} = (U_t^k U_t^{kT})^{-1} U_t^{kT}(I_i \circ \tau_i^k + J_i^k \triangle\tau^{p+1} - e^p - \frac{1}{\mu^p}\lambda^p) \\ e^{p+1} = \mathsf{S}_{\frac{1}{\mu}}(I_i \circ \tau_i^k + J_i^k \triangle\tau^{p+1} - U_t^k w^{p+1} - \frac{1}{\mu^p}\lambda^p) \\ \lambda^{p+1} = \lambda^p + \mu^p h(w^{p+1}, e^{p+1}, \triangle\tau^{p+1}) \\ \mu^{p+1} = \rho\mu^p \end{cases} \tag{9}$$

where $\mathsf{S}_{\frac{1}{\mu}}$ is the elementwise soft thresholding operator [29], and $\rho > 1$ is the ADMM penalty constant enforcing $\{\mu^p\}$ to be a monotonically increasing positive sequence. The iteration (9) indeed converges to the optimal solution of the problem (7) [30]. We summarize this ADMM solver as Algorithm 2 in Section 2.4.

2.3. Subspace Update

Whether identifying the best U^{k*} in the batch mode (3) or estimating the union of subspaces U^{ℓ}, $\ell = 1, \ldots, L$, in the online mode (5), optimizing the orthonormal matrix U along the geodesic of Grassmannian is our key technique. For clarity of exposition in this section, we remove the superscript k or ℓ from U, as the core gradient step along the geodesic of the Grassmannian for both batch mode and online mode is the same. We seek a sequence $\{U_t\} \in \mathcal{G}(d, n)$ such that $U_t \longrightarrow U^*$ (as $t \to \infty$). We now face the choice of an effective subspace loss function. Regarding U as the variable, the loss function (6) is not differentiable everywhere. Therefore, we choose to instead use the augmented Lagrangian (8) as the subspace loss function once we have estimated $(w^*, e^*, \triangle\tau^*, \lambda^*)$ by ADMM (9) from the previous U_t [13, 6].

In order to take a gradient step along the geodesic of the Grassmannian, according to [26], we first need to derive the gradient formula of the real-valued loss function (8) $\mathcal{L} : \mathcal{G}(d, n) \to \mathbb{R}$. The gradient $\nabla\mathcal{L}$ can be determined from the derivative of $\mathcal{L}$ with respect to the components of U:

$$\frac{d\mathcal{L}}{dU} = (\lambda^* + \mu h(w^*, e^*, \triangle\tau^*))\, w^{*T} \qquad (10)$$

Then the gradient is $\nabla\mathcal{L} = (I - UU^T)\frac{d\mathcal{L}}{dU}$ [26]. From Step 6 of Algorithm 1, we have that $\nabla\mathcal{L} = \Gamma w^{*T}$ (see the definition of Γ in Alg. 1). It is easy to verify that $\nabla\mathcal{L}$ is rank one since Γ is a $n \times 1$ vector and w^* is a $d \times 1$ weight vector. The following derivation of geodesic gradient step is similar to GROUSE [24] and GRASTA [13, 6]. We rewrite the important steps of the derivation here for completeness.

The sole non-zero singular value is $\sigma = \|\Gamma\|\|w^*\|$, and the corresponding left and right singular vectors are $\frac{\Gamma}{\|\Gamma\|}$ and $\frac{w^*}{\|w^*\|}$ respectively. Then we can write the SVD of the gradient explicitly by adding the orthonormal set $x_2, \ldots, x_d$ orthogonal to Γ as left singular vectors and the orthonormal set $y_2, \ldots, y_d$ orthogonal to w^* as right singular vectors as follows:

$$\nabla\mathcal{L} = \quad \left[\frac{\Gamma}{\|\Gamma\|} \quad x_2 \quad \ldots \quad x_d \right] \times \mathrm{diag}(\sigma, 0, \ldots, 0)$$
$$\times \left[\frac{w^*}{\|w^*\|} \quad y_2 \quad \ldots \quad y_d \right]^T .$$

Finally, following Equation (2.65) in [26], a geodesic gradient step of length η in the direction $-\nabla\mathcal{L}$ is given by

$$U(\eta) = U \;\; + \;\; (\cos(\eta\sigma) - 1)\frac{Uw_t^*}{\|w_t^*\|}\frac{w_t^{*T}}{\|w_t^*\|}$$

$$- \;\; \sin(\eta\sigma)\frac{\Gamma}{\|\Gamma\|}\frac{w_t^{*T}}{\|w_t^*\|} \; . \tag{11}$$

2.4. Algorithms

2.4.1. Batch Mode

From the discussion of of Sections 2.2 and 2.3, given the batch of unaligned images D, their estimate of transformation τ^k and their Jacobian J^k at iteration k, we can robustly identify the subspace U^{k*} by incrementally updating U_t^k along the geodesic of Grassmannian $\mathcal{G}(d^k, n)$ (11). When $U_t^k \longrightarrow U^{k*}$ (as $t \to \infty$), the estimate of $\triangle\tau_i$ for each initially aligned image $I_i \circ \tau_i^k$ also approaches its optimal value $\triangle\tau_i^*$. Once the subspace U^k is accurately learned, we will update the estimate of the transformation for each image using $\tau_i^{k+1} = \tau_i^k + \triangle\tau_i^*$. Then in the next iteration, the new subspace U^{k+1} can also be learned from $D \circ \tau^{k+1}$, and the algorithm iterates until we reach the stopping criterion, e.g. if $\frac{\|\triangle\tau\|_2}{\|\tau^k\|_2} < \epsilon$ or we reach the maximum iteration K.

We summarize our algorithms as follows. Algorithm 1 is the batch image alignment approach via iterative online robust subspace learning. For Step 7, there are many ways to pick the step-size. For some examples, you may consider the diminishing and constant step-sizes adopted in GROUSE [24], or the multi-level adaptive step-size used for fast convergence in GRASTA [13].

Algorithm 2 is the ADMM solver for the locally linearized problem (7). From our extensive experiments, if we set the ADMM penalty parameter $\rho = 2$ and the tolerance $\epsilon^{tol} = 10^{-7}$, Algorithm 2 has always converged in fewer than 20 iterations.

2.4.2. Online Mode

In Section 2.1.2, we propose to tackle the difficult nonlinear online subspace learning problem by iteratively learning online a union of subspaces U^ℓ, $\ell = 1, \ldots, L$. For a sequence of video frames $I_i, i = 1, \ldots, N$, the union of subspaces U^ℓ are updated iteratively as illustrated in Fig. 3.

Specifically, at i-th frame I_i, for the locally approximated subspace U_i^1 at the first iteration, given the initial roughly estimated transformation τ_i^0, the

11

Algorithm 1 Transformed GRASTA - batch mode

Require: An initial $n \times d^0$ orthogonal matrices U^0. A sequence of unaligned images I_i and the corresponding initial transformation parameters τ_i^0, $i = 1, \ldots, N$. The maximum iteration K.

Return: The estimated well-aligned subspace U^{k*} for the well-aligned images. The transformation parameters τ_i^k for each well-aligned image.

1: **while** not converged **and** $k < K$ **do**
2:     Update the Jacobian matrix of each image :

$$J_i^k = \frac{\partial(I_i \circ \zeta)}{\partial \zeta}\big|_{\zeta = \tau_i^k} \quad (i = 1 \ldots N)$$

3:     Update the wrapped and normalized images:

$$I_i \circ \tau_i^k = \frac{vec(I_i \circ \tau_i^k)}{\|vec(I_i \circ \tau_i^k)\|_2}$$

4:     **for** $j = 1 \rightarrow N, \ldots,$ *until converged* **do**
5:         Estimate the weight vector w_j^k, the sparse outliers e_j^k, the locally linearized transformation parameters $\triangle \tau_j^k$, and the dual vector λ_j^k via the ADMM algorithm 2 from $I_i \circ \tau_i^k$, J_i^k, and the current estimated subspace U_t^k

$$(w_j^k, e_j^k, \triangle \tau_j^k, \lambda_j^k) = \arg \min_{w,e,\triangle\tau,\lambda} \mathcal{L}(U_t^k, w, e, \lambda)$$

6:         Compute the gradient $\nabla\mathcal{L}$ as follows:
          $\Gamma_1 = \lambda_j^k + \mu h(w_j^k, e_j^k, \triangle \tau_j^k),$
          $\Gamma = (I - U_t^k U_t^{k^T})\Gamma_1, \qquad \nabla\mathcal{L} = \Gamma w_j^{k^T}$
7:         Compute step-size η_t.
8:         Update subspace:
        $U_{t+1}^k = U_t^k + \Big((\cos(\eta_t\sigma) - 1)U_t \frac{w_j^k}{\|w_j^k\|}$
        $- \sin(\eta_t\sigma)\frac{\Gamma}{\|\Gamma\|} \Big) \frac{w_j^{k^T}}{\|w_j^k\|}, \quad$ where $\sigma = \|\Gamma\|\|w_j^k\|$.
9:     **end for**
10:    Update the transformation parameters:

$$\tau_i^{k+1} = \tau_i^k + \triangle \tau_i^k, \quad (i = 1 \ldots N)$$

11: **end while**

Algorithm 2 ADMM Solver for the Locally Linearized Problem (7)
___
Require: An $n \times d$ orthogonal matrix U, a wrapped and normalized image $I \circ \tau \in \mathbb{R}^n$, the corresponding Jacobian matrix J, and a structure OPTS which holds four parameters for ADMM: ADMM penalty constant ρ, the tolerance ϵ^{tol}, and ADMM maximum iteration K.

Return: weight vector $w^* \in \mathbb{R}^d$; sparse outliers $e^* \in \mathbb{R}^n$; locally linearized transformation parameters $\triangle\tau^*$; and dual vector $\lambda^* \in \mathbb{R}^n$.

1: Initialize $w, e, \triangle\tau, \lambda, and\ \mu$: $e^1 = 0, w^1 = 0, \triangle\tau^1 = 0,\ \lambda^1 = 0,\ \mu = 1$
2: Cache $P = (U^T U)^{-1} U^T$ and $F = (J^T J)^{-1} J^T$
3: **for** $k = 1 \rightarrow K$ **do**
4:    Update $\triangle\tau$: $\triangle\tau^{k+1} = F(Uw^k + e^k - I \circ \tau + \frac{1}{\mu}\lambda^k)$
5:    Update weights: $w^{k+1} = P(I \circ \tau + J\triangle\tau^{k+1} - e^k - \frac{1}{\mu}\lambda^k)$
6:    Update sparse outliers:
     $e^{k+1} = \mathsf{S}_{\frac{1}{\mu}}(I \circ \tau + J\triangle\tau^{k+1} - Uw^{k+1} - \frac{1}{\mu}\lambda^k)$
7:    Update dual: $\lambda^{k+1} = \lambda^k + \mu h(w^{k+1}, e^{k+1}, \triangle\tau^{k+1})$
8:    Update μ: $\mu = \rho\mu$
9:    **if** $\|h(w^{k+1}, e^{k+1}, \triangle\tau^{k+1})\|_2 \leq \epsilon^{tol}$ **then**
10:       Converge and break the loop.
11:    **end if**
12: **end for**
13: $w^* = w^{k+1}$, $e^* = e^{k+1}$, $\triangle\tau^* = \triangle\tau^{k+1}$, $\lambda^* = y^{k+1}$
___

ADMM solver Algorithm 2 gives us the locally estimated $\triangle\tau_i^1$, and the updated subspace U_{i+1}^1 is obtained by taking a gradient step along the geodesic of the Grassmannian $\mathcal{G}(d^1, n)$ as discussed in Section 2.3. The transformation τ_i^1 of the next iteration is updated by $\tau_i^1 = \tau_i^0 + \triangle\tau_i^1$. Then for the next locally approximated subspace U_i^2, we also estimate $\triangle\tau_i^2$ and update the subspace along the geodesic of the Grassmannian $\mathcal{G}(d^2, n)$ to U_{i+1}^2. Repeatedly, we will update U_i^ℓ in the same way to get U_{i+1}^ℓ and the new transformation $\tau_i^\ell = \tau_i^{\ell-1} + \triangle\tau_i^\ell$. After completing the update for all L subspaces, the union of subspaces $U_{i+1}^\ell (\ell = 1, \ldots, L)$ will be used for approximating the nonlinear transform of the next video frame I_{i+1}.

We summarize the above statements as Algorithm 3, and we call this approach the *fully online mode* of t-GRASTA.
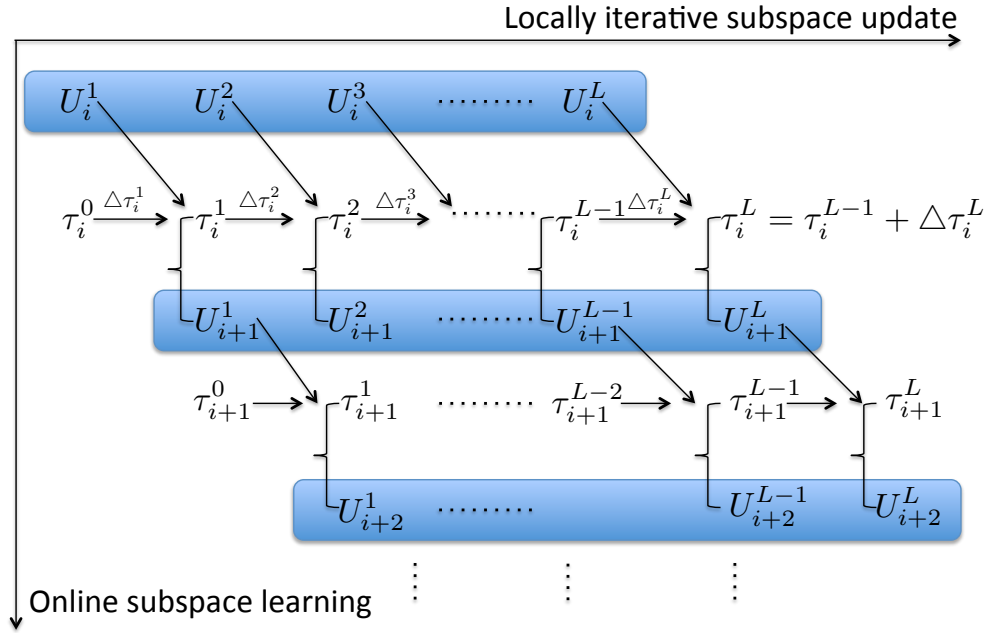

Figure 3: The diagram of the *fully online mode* of t-GRASTA.

2.4.3. Discussion of Online Image Alignment

If the subspace U^k of the well-aligned images is known as a prior, for example if U^k is trained by Algorithm 1 from a "well selected" dataset of one

Algorithm 3 Transformed GRASTA - *Fully Online Mode*

Require: The initial L $n \times d^\ell$ orthonormal matrices U^ℓ spanning the corresponding subspace $\mathcal{S}^\ell$, $\ell = 1, ..., L$. A sequence of unaligned images I_i and the corresponding initial transformation parameters τ_i^0, $i = 1, \ldots, N$.

Return: The estimated iteratively approximated subspaces U_i^ℓ, $\ell = 1, \ldots, L$, after processing image I_i. The transformation parameters τ_i^L for each well-aligned image.

1: **for** unaligned image $I_i, i = 1, \ldots, N$ **do**
2:     **for** the iterative approximated subspace $U^\ell, \ell = 1, \ldots, L$ **do**
3:         Update the Jacobian matrix of image I_i:

$$J_i^\ell = \frac{\partial(I_i \circ \zeta)}{\partial \zeta}|_{\zeta = \tau_i^\ell}$$

4:         Update the wrapped and normalized images:

$$I_i \circ \tau_i^\ell = \frac{vec(I_i \circ \tau_i^\ell)}{\|vec(I_i \circ \tau_i^\ell)\|_2}$$

5:         Estimate the weight vector w_i^ℓ, the sparse outliers e_i^ℓ, the locally linearized transformation parameters $\triangle \tau_i^\ell$, and the dual vector λ_i^ℓ via the ADMM algorithm 2 from $I_i \circ \tau_i^\ell$, J_i^ℓ, and the current estimated subspace U_i^ℓ

$$(w_i^\ell, e_i^\ell, \triangle \tau_i^\ell, \lambda_i^\ell) = \arg \min_{w,e,\triangle\tau,\lambda} \mathcal{L}(U_i^\ell, w, e, \lambda)$$

6:         Compute the gradient $\nabla \mathcal{L}$ as follows:
          $\Gamma_1 = \lambda_i^\ell + \mu h(w_i^\ell, e_i^\ell, \triangle\tau_i^\ell),$
          $\Gamma = (I - U_t^\ell U_t^{\ell T})\Gamma_1,$ $\nabla \mathcal{L} = \Gamma w_i^{\ell T}$
7:         Compute step-size η_i^ℓ.
8:         Update subspace:
          $U_{i+1}^\ell = U_i^\ell + \left((\cos(\eta_i^\ell \sigma) - 1)U_t \frac{w_i^\ell}{\|w_i^\ell\|} - \sin(\eta_i^\ell \sigma) \frac{\Gamma}{\|\Gamma\|}\right) \frac{w_i^{\ell T}}{\|w_i^\ell\|},$
          where $\sigma = \|\Gamma\|\|w_i^\ell\|$.
9:         Update the transformation parameters:

$$\tau_i^{\ell+1} = \tau_i^\ell + \triangle \tau_i^\ell$$

10:     **end for**
11: **end for**

category, we can simply use U^k to align the rest of the unaligned images of the same category. Here "well selected" means the training dataset should cover enough of the global appearance of the object, such as different illuminations, which can be represented by the low-dimensional subspace structure. By category, we mean a particular object of interest or a particular background scene in the video surveillance data.

For massive image processing tasks, it is easy to collect such good training datasets by simply randomly sampling a small fraction of the whole image set. Once U^k is learned from the training set, we can use a variation of Algorithm 1 to align each unaligned image I without updating the subspace, since we have the assumption that the remaining images also lie in the trained subspace. We call Algorithm 4 the *trained online mode*.

However, we note that for a very large streaming dataset such as is typical in real-time video processing, the *trained online mode* may be less well-defined, as the subspace of the streaming video data may change over time. For this scenario, our *fully online mode* for t-GRASTA could gradually adapt to the changing subspace and then accurately estimate the transformation τ.

2.5. Discussion of Memory Usage

We compare the memory usage of our *fully online mode* of t-GRASTA to that of RASL. RASL requires storage of A, E, a Lagrange multiplier matrix Y, the data D, and $D \circ \tau$, each of which require storage of the size nN. To compare fairly to t-GRASTA, which assumes a d-dimensional model, we suppose RASL uses a thin singular value decomposition of size d, which requires $nd + Nd + d^2$ memory elements. Finally for the Jacobian per image, RASL needs nNp, and for τ RASL needs Np, but we will assume p is a small constant independent of dimension and ignore it. Therefore RASL's total memory usage is $6nN + nd + Nd + d^2 + N$.

t-GRASTA must also store the Jacobian, τ, and the data as well as the data with transformation, using memory size $3nN + N$. Otherwise, t-GRASTA needs to store the union of subspaces U^ℓ, $\ell = 1, \ldots, L$ matrices of size $Lnd (L \ll N)$, and the vectors e, λ, Γ, and w for $3n + d$ memory elements. Thus t-GRASTA's memory total is $3nN + Lnd + 3n + d + N$.

For a problem size of 100 images, each with 100×100 pixels, and assuming $d = 10$, $L = 10$, t-GRASTA uses 66.1% of the memory of RASL. For 10000 mega-pixel images, t-GRASTA uses 50.1% of the memory of RASL. The scaling remains about half throughout mid-range to large problem sizes.

Algorithm 4 *Trained Online Mode* of Image Alignment

Require: A well-trained $n \times d$ orthogonal matrix U. An unaligned image I and the corresponding initial transformation parameters τ^0. The maximum iteration K.

Return: The transformation parameters τ^k for the well-aligned image.

1: **while** not converged **and** $k < K$ **do**

2:   Update the Jacobian matrix :

$$J^k = \frac{\partial(I \circ \zeta)}{\partial \zeta}|_{\zeta = \tau^k}$$

3:   Update the wrapped and normalized image:

$$I \circ \tau^k = \frac{vec(I \circ \tau^k)}{\|vec(I \circ \tau^k)\|_2}$$

4:   Estimate the weight vector w^k, the sparse outliers e^k, the locally linearized transformation parameters $\triangle \tau^k$, and the dual vector λ^k via the ADMM algorithm 2 from $I \circ \tau^k$, J^k, and the well-trained subspace U

$$(w^k, e^k, \triangle \tau^k, \lambda^k) = \arg \min_{w,e,\triangle\tau,\lambda} \mathcal{L}(U, w, e, \lambda)$$

5:   Update the transformation parameters:

$$\tau^{k+1} = \tau^k + \triangle \tau^k$$

6: **end while**

---

3. PERFORMANCE EVALUATION

In this section, we conduct comprehensive experiments on a variety of alignment tasks to verify the efficiency and superiority of our algorithm. We first demonstrate the ability of the proposed approach to cope with occlusion and illumination variation during the alignment process. After that, we further demonstrate the robustness and generality of our approach by testing it on handwritten digits and face images taken from the Labeled Faces in the Wild database [31]. Finally, we apply our approach to dealing with video jitters and solving the interesting background foreground separation problem.

3.1. Occlusion and illumination variation

We first test our approach on the dataset 'dummy' described in [15]. Here, we want to verify the ability of our approach to effectively align the images despite occlusion and illumination variation. The dataset contains 100 images of a dummy head taken under varying illumination and with artificially generated occlusions created by adding a square patch at a random location of the image. Fig. 4 shows 10 misaligned images of the dummy. We align these images by Algorithm 1 (the batch mode of t-GRASTA). The canonical frame is chosen to be 49×49 pixels and the subspace dimension is set to 5. Here and in the rest of our experiments, for simplicity we set d^k of Algorithm 1 to a fixed d in every iteration. The last three rows of Fig. 4 show the results of alignment, from which we can see that our approach is successful at aligning the misaligned images while removing the occlusion at the same time.

3.2. Robustness

In order to further demonstrate the robustness of our approach, we apply it on more realistic images taken from the Labeled Faces in the Wild database [31]. The LFW contains more severely misaligned images, for it also includes remarkable variations in pose and expression aside from illumination and occlusion, which can be seen in Fig. 5(c). We chose 16 subjects from LFW, each of them with 35 images. Each image is aligned to an 80×60 canonical frame using τ which are from the group of affine transformations $\mathbb{G} = Aff(2)$, as in [15]; these are translations, rotations, and scale transformations. For each subject, we set the subspace dimension $= 15$ and use Algorithm 1 to align each image. In this example, we demonstrate the robustness of our approach by comparing the average face of each subject before

Figure 4: The first row shows the original misaligned images with occlusions and illumination variation; the second row shows the images aligned by t-GRASTA; the third row shows the recovered aligned images without occlusion; and the bottom row is the occlusion removed by our approach.

and after alignment, which are shown in Fig. 5(a)-(b). We can see that the average faces after alignment are much clearer than those before alignment. Fig. 5(c)-(d) provides more detailed information, showing the unaligned and aligned images of John Ashcroft (marked by red boxes in Fig. 5(a)-(b)).

3.3. Generality

The previous experiments have demonstrated the effectiveness and robustness of t-GRASTA. Here we wish to show the generality of t-GRASTA by applying it to aligning a different type of images – handwritten digits taken from MINST database. For this experiment, we again use Algorithm 1 to align 100 images of a handwritten "3" to a 29×29 canonical frame size. We use Euclidean transformation $\mathbb{G} = E(2)$ and set the dimension of the subspace to be 5.

Fig. 6 shows that t-GRASTA can successfully align the misaligned digits and learn the low dimensional subspace, even though the original digits have significant variation. We can see that the outliers separated by t-GRASTA are generated by variations in the digits that are not consistent with the global appearance. The outliers (d) would be even more sparse if the subspace representation in (c) were to capture more of this variation; If desired, we could achieve this tradeoff by increasing the dimension of the subspace.

Figure 5: (a) Average of 16 misaligned subjects randomly selected from LFW database; (b) average of each subject aligned by t-GRASTA; (c) initial images of John Ashcroft (marked by red boxs in (a) and (b)); (d) images aligned by t-GRASTA.

3.4. Video Jitter

In this section, we apply t-GRASTA to separation problems made difficult by video jitter. Here we apply both the *fully online mode* Algorithm 3 and the *trained online mode* Algorithm 4 to different datasets. We show the superiority of t-GRASTA regarding both the speed and memory requirement of the algorithms.

Figure 6: (a) 100 misaligned digits; (b) digits aligned by t-GRASTA; (c) subspace representation of corresponding digits; (d) outliers.

3.4.1. Hall

Here we apply t-GRASTA to the task of separating moving objects from static background in the video footage recorded by an unstable camera. We note that in [6], the authors simulate a virtual panning camera to show

that GRASTA can quickly track sudden changes in the background subspace caused by a moving camera. Their low-rank subspace tracking model is well-defined, as the camera after panning is still stationary, and thus the recorded video frames are accurately pixelwise aligned. However, for an unstable camera, the recorded frames are no longer aligned; the background cannot be well represented by a low-rank subspace unless the jittered frames are first aligned. In order to show that t-GRASTA can tackle this separation task, we consider a highly jittered video sequence generated by a simulated unstable camera. To simulate the unstable camera, we randomly translate the original well-aligned video frames in x- / y- axis and rotate them in the plane.

In this experiment, we compare t-GRASTA with RASL and GRASTA. We use the first 200 frames of the "Hall" dataset[1], each 144×176 pixels. We first perturb each frame artificially to simulate camera jitter. The rotation of each frame is random, uniformly distributed within the range of $[-\theta_0/2, \theta_0/2]$, and the ranges of x- and y-translations are limited to $[-x_0/2, x_0/2]$ and $[-y_0/2, y_0/2]$. In this example, we set the perturbation size parameters $[x_0, y_0, \theta_0]$ with the values of $[\ 20, 20, 10°]$.

For comparing with RASL, unlike [23], we just let RASL run its original batch model without forcing it into an online algorithm framework. The task we give to RASL and t-GRASTA is to align each frame to a 62×75 canonical frame, again using $\mathbb{G} = Aff(2)$. The dimension of the subspace in t-GRASTA is set to be 10. We first randomly select 30 frames of the total 200 frames to train the subspace by Algorithm 1 and then align the rest using the *trained online mode*. The visual comparison between RASL and t-GRASTA are shown in Fig. 7. Table 1 illustrates the numerical comparison of RASL and t-GRASTA, for which we ran each algorithm 10 times to get the statistics. From Table 1 and Fig. 7 we can see that the two algorithms achieve a very similiar effect, but t-GRASTA runs much faster than RASL: On a PC with Intel P9300 2.27GHz CPU and 2 GB of RAM, the average time for aligning a newly arrived frame is 1.1 second, while RASL needs more than 800 seconds to align the total batch of images, or 4 seconds per frame. Moreover, our approach is also superior to RASL regarding memory efficiency. These superiorities become more dramatic as one increases the

---

[1] Find these along with the videos at `http://perception.i2r.a-star.edu.sg/bk_model/bk_index.html`.

size of the image database.

Table 1: Statistics of errors in two pixels P_1 and P_2, selected from the original video frames and traced through the jitter simulation process to the RASL and t-GRASTA output frames. Max error and mean error are calculated as the distances from the estimated P_1 and P_2 to their statistical center $E(P_1)$ and $E(P_2)$. Std are calculated as the standard deviation of four coordinate value (X_1, Y_1) for P_1 and (X_2, Y_2) for P_2 across all frames.

	Max error	Mean error	X1 std	Y1 std	X2 std	Y2 std
Initial misalignment	11.24	5.07	3.35	3.01	3.34	4.17
RASL	**2.96**	1.73	0.56	**0.71**	0.90	1.54
t-GRASTA	6.62	**0.84**	**0.48**	1.11	**0.57**	**0.74**

(a)   (b)   (c)   (d)   (e)

Figure 7: Comparison between t-GRASTA and RASL. (a) Average of initial misaligned images; (b) average of images aligned by t-GRASTA; (c)average of background recovered by t-GRASTA; (d) average of images aligned by RASL; (e) average of background recovered by RASL.

In order to compare with GRASTA, we use 200 perturbed images to recover the background and separate the moving objects for both algorithms; Fig. 8 illustrates the comparison. For both GRASTA and t-GRASTA, we set the subspace rank $= 10$ and randomly selected 30 images to train the subspace first. For t-GRASTA, we use the affine transformation $\mathbb{G} = Aff(2)$. From Fig. 8, we can see that our approach successfully separates the foreground and the background and simultaneously align the perturbed images. But GRASTA fails to learn a proper subspace, thus, the separation of background and foreground is poor. Although GRASTA has been demonstrated to successfully track a dynamic subspace, e.g. the panning camera, the dynamics of an unstable camera are too fast and unpredictable for the GRASTA subspace tracking model to succeed in this context without pre-alignment of the video frames.

23

Figure 8: Video background and foreground separation with jittered video. 1^{st} row: 8 misaligned video frames randomly selected from artificially perturbed images; 2^{nd} row: images aligned by t-GRASTA; 3^{rd} row: background recovered by t-GRASTA; 4^{th} row: foreground separated by t-GRASTA; 5^{th} row: background recovered by GRASTA; 6^{th} row: foreground separated by GRASTA.

3.4.2. Gore

In this example, we show the capability of t-GRASTA for video stabilization applied to the dataset "Gore" described in [15]. In [15], the original face images are obtained by a face detector, and the jitters are caused by the inherent imprecision of the detector. In contrast, for t-GRASTA, we simply crop the face from each image by a constant rectangle with size 68×44, which has the same position parameters for all frames. So in our case, the jitters are caused by the differences between the motion and pose variation of the target and the stabilization of the constant rectangle.

Figure 9: The first row shows the original misaligned images; the second row shows the images aligned by t-GRASTA; the third row shows the recovered aligned images without outliers; and the bottom row shows the outliers removed by our approach.

For this experiment, the dimension of the subspace is set to be 10, and we again choose the affine transformation $\mathbb{G} = Aff(2)$. We first use the Algorithm 1 to train an initial subspace by 20 images randomly selected from the whole set of 140 images. We then use the *fully online mode* to align the rest of the images. Fig. 9 show the results. t-GRASTA did well for this dataset with better speed than RASL: On a PC with Intel P9300 2.27GHz CPU and 2 GB of RAM, t-GRASTA aligned these images at 5 frames per second. This is 5 times faster than RASL and 3 times faster than ORIA as described in [23].

Although t-GRASTA was not designed as a face detector, the experimental results suggest that t-GRASTA can be transformed into a face detector, or more generally target tracker, if the variation of pose of the target is limited in a certain range (usually 45°). In this case, we can further improve the efficiency of t-GRASTA by choosing a tight frame for the canonical image.

3.4.3. Sidewalk

In the last experiment, we use misaligned frames caused by real camera jitter to test t-GRASTA. Here we align all 1200 frames of "Sidewalk" dataset[2] to 50×78 canonical frames, again using $\mathbb{G} = Aff(2)$ and subspace dimension 5. We also use the first 20 frames to train the initial subspace using the batch mode Algorithm 1, and then use the *fully online mode* to align the rest of the frames. Here we can see that aligning the total 1200 frames is a heavy task for RASL – for our PC with Intel P9300 2.27GHz CPU and 2 GB of RAM, it was necessary to divide the dataset into four parts each containing 300 frames. We then let RASL separately run on each sub-dataset. The total time needed by RASL was around 1000 seconds for 1.2 frames per second, while t-GRASTA achieved more than 4 frames per second without partitioning the data.

Compared to the *trained online mode*, the *fully online mode* can track changes of the subspace over time. This is an important asset of the *fully online mode*, especially when it comes to large streaming datasets containing considerable variations. We see that we usually need no more than 20 frames for *fully online mode* to adapt to the changes of the subspace, such as illumination changes or dynamic background caused by the motion of the subspace. Moreover, if the changes are slow, i.e the natural illumination changes from daylight or the camera moving slowly, then t-GRASTA needs no extra frames to track such changes; it incorporates such information with each iteration during the slowly changing process.

4. CONCLUSIONS AND FUTURE WORK

4.1. Conclusions

In this paper we have presented an iterative Grassmannian optimization approach to simultaneously identify an optimal set of image domain transformations for image alignment and the low-rank subspace matching the aligned images. These are such that the vector of each transformed image can be decomposed as the sum of a low-rank part of the recovered aligned image and a sparse part of errors. This approach can be regarded as an extension of GRASTA and RASL: We extend GRASTA to transformations, and extend

---

[2]Find it along with other datasets containing misaligned frames caused by real video jitters at `http://wordpress-jodoin.dmi.usherb.ca/dataset`.

26

Figure 10: Video background and foreground separation with jittered video. 1^{st} row: 8 original misaligned video frames caused by video jitter; 2^{nd} row: images aligned by t-GRASTA; 3^{rd} row: background recovered by t-GRASTA; 4^{th} row: foreground separated by t-GRASTA.

RASL to the incremental gradient optimization framework. Our approach is faster than RASL and more robust to alignment than GRASTA. We can effectively and computationally efficiently learn the low-rank subspace from misaligned images, which is very practical for computer vision applications.

4.2. Future Work

Though this work presents an approach for robust image alignment more computationally efficient than state-of-the-art, a foremost remaining problem is how to scale the proposed approach to a very large streaming dataset such as is typical in real-time video processing. The fully online t-GRASTA algorithm presented here is a first step towards a truly large-scale real-time algorithm, but several practical implementation questions remain, including online parameter selection and error performance cross-validation. Another question of interest is regarding the estimation of d^k for the subspace update. Though we fix the rank d in this paper, estimating d^k and switching between Grassmannians is a very interesting future direction.

While preparing the conference version of this work [32], we noticed an interesting alignment approach proposed in [33]. Though the two approaches of ours and [33] are both obtained via optimization over a manifold, they

perform alignment for very different scenarios. For example, the approach in [33] focuses on semantically meaningful videos or signals, and then it can successfully align the videos of the same object from different views; t-GRASTA manipulates the set of misaligned images or the video of unstable camera to robustly identify the low-rank subspace, and then it can align these images according to the subspace. An intriguing future direction would be to merge these two approaches.

A final direction of future work is toward applications which require more aggressive background tracking than is possible by a GRASTA-type algorithm. For example, if a camera is following an object around different parts of a single scene, even though the background may be quickly varying from frame to frame, the camera will get multiple shots of different pieces of the background. Therefore, it may be possible to still build a model for the entire background scene using low-dimensional modeling. Incorporating camera movement parameters and a dynamical model into GRASTA would be a natural way to solve this problem, merging classical adaptive filtering algorithms with modern manifold optimization.

5. ACKNOWLEDGEMENTS

References

[1] S. Odio, Making facebook photos better, accessed July 2010 at `https://www.facebook.com/blog/blog.php?post=403838582130`.

[2] Instagram, Year in review: 2011 in numbers, accessed January 2012 at `http://blog.instagram.com/post/15086846976/year-in-review-2011-in-numbers`.

[3] D. Babwin, Cameras make chicago most closely watched U.S. city (April 6 2010).

[4] M. McCahill, C. Norris, Cctv in london, Working Paper 6, Centre for Criminology and Criminal Justice, University of Hull, United Kingdom (June 2002).

[5] E. J. Candès, X. Li, Y. Ma, J. Wright, Robust principal component analysis?, J. ACM 58 (3) (2011) 11:1–11:37. `doi:10.1145/1970392.1970395`.

[6] J. He, L. Balzano, A. Szlam, Incremental gradient on the grassmannian for online foreground and background separation in subsampled video, in: Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on, 2012, pp. 1568–1575. `doi:10.1109/CVPR.2012.6247848`.

[7] G. Mateos, G. Giannakis, Sparsity control for robust principal component analysis, in: Signals, Systems and Computers (ASILOMAR), 2010 Conference Record of the Forty Fourth Asilomar Conference on, 2010, pp. 1925–1929. `doi:10.1109/ACSSC.2010.5757875`.

[8] R. Sivalingam, A. D'Souza, M. Bazakos, R. Miezianko, V. Morellas, N. Papanikolopoulos, Dictionary learning for robust background modeling, in: Robotics and Automation (ICRA), 2011 IEEE International Conference on, 2011, pp. 4234–4239. `doi:10.1109/ICRA.2011.5979981`.

[9] F. De la Torre, M. Black, A framework for robust subspace learning, International Journal of Computer Vision 54 (1-3) (2003) 117–142. `doi:10.1023/A:1023709501986`.

[10] P. Jodoin, J. Konrad, V. Saligrama, V. Veilleux-Gaboury, Motion detection with an unstable camera, in: Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on, 2008, pp. 229–232. `doi:10.1109/ICIP.2008.4711733`.

[11] G. Puglisi, S. Battiato, A robust image alignment algorithm for video stabilization purposes, Circuits and Systems for Video Technology, IEEE Transactions on 21 (10) (2011) 1390–1400. `doi:10.1109/TCSVT.2011.2162689`.

[12] K. Simonson, T. Ma, Robust real-time change detection in high jitter, Sandia Report SAND2009-5546 (2009) 1–41.

[13] J. He, L. Balzano, J. Lui, Online robust subspace tracking from partial information, Arxiv preprint arXiv:1109.3827.

[14] Y. Peng, A. Ganesh, J. Wright, W. Xu, Y. Ma, Rasl: Robust alignment by sparse and low-rank decomposition for linearly correlated images, in: Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on, 2010, pp. 763–770. `doi:10.1109/CVPR.2010.5540138`.

[15] Y. Peng, A. Ganesh, J. Wright, W. Xu, Y. Ma, Rasl: Robust alignment by sparse and low-rank decomposition for linearly correlated images, Pattern Analysis and Machine Intelligence, IEEE Transactions on 34 (11) (2012) 2233–2246. `doi:10.1109/TPAMI.2011.282`.

[16] Z. Zhang, A. Ganesh, X. Liang, Y. Ma, Tilt: Transform invariant low-rank textures, International Journal of Computer Vision 99 (1) (2012) 1–24. `doi:10.1007/s11263-012-0515-x`.

[17] A. Wagner, J. Wright, A. Ganesh, Z. Zhou, H. Mobahi, Y. Ma, Toward a practical face recognition system: Robust alignment and illumination by sparse representation, Pattern Analysis and Machine Intelligence, IEEE Transactions on 34 (2) (2012) 372–386. `doi:10.1109/TPAMI.2011.112`.

[18] G. Huang, V. Jain, E. Learned-Miller, Unsupervised joint alignment of complex images, in: Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on, 2007, pp. 1–8. `doi:10.1109/ICCV.2007.4408858`.

[19] E. Learned-Miller, Data driven image models through continuous joint alignment, Pattern Analysis and Machine Intelligence, IEEE Transactions on 28 (2) (2006) 236–250. `doi:10.1109/TPAMI.2006.34`.

[20] M. Cox, S. Sridharan, S. Lucey, J. Cohn, Least squares congealing for unsupervised alignment of images, in: Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on, 2008, pp. 1–8. `doi:10.1109/CVPR.2008.4587573`.

[21] M. Cox, S. Sridharan, S. Lucey, J. Cohn, Least-squares congealing for large numbers of images, in: Computer Vision, 2009 IEEE 12th International Conference on, 2009, pp. 1949–1956. `doi:10.1109/ICCV.2009.5459430`.

[22] V. Chandrasekaran, S. Sanghavi, P. Parrilo, A. Willsky, Rank-sparsity incoherence for matrix decomposition, SIAM Journal on Optimization 21 (2) (2011) 572–596. `doi:10.1137/090761793`.

[23] Y. Wu, B. Shen, H. Ling, Online robust image alignment via iterative convex optimization, in: Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on, 2012, pp. 1808 –1814. `doi: 10.1109/CVPR.2012.6247878`.

[24] L. Balzano, R. Nowak, B. Recht, Online identification and tracking of subspaces from highly incomplete information, in: Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on, 2010, pp. 704–711. `doi:10.1109/ALLERTON.2010.5706976`.

[25] P. Comon, G. Golub, Tracking a few extreme singular values and vectors in signal processing, Proceedings of the IEEE 78 (8) (1990) 1327–1343. `doi:10.1109/5.58320`.

[26] A. Edelman, T. Arias, S. Smith, The geometry of algorithms with orthogonality constraints, SIAM Journal on Matrix Analysis and Applications 20 (2) (1998) 303–353. `doi:10.1137/S0895479895290954`.

[27] D. P. Bertsekas, Incremental gradient, subgradient, and proximal methods for convex optimization: A survey, Tech. Rep. LIDS-P-2848, MIT Lab for Information and Decision Systems (August 2010).

[28] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, Distributed optimization and statistical learning via the alternating direction method of multipliers, Found. Trends Mach. Learn. 3 (1) (2011) 1–122. `doi: 10.1561/2200000016`.

[29] S. Boyd, L. Vandenberghe, Convex optimization, Cambridge university press, 2004.

[30] D. P. Bertsekas, Nonlinear Programming, Athena Science, 2004.

[31] G. B. Huang, M. Mattar, T. Berg, E. Learned-Miller, Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments, in: Proceedings of the European Conference on Computer Vision, Workshop on Faces in 'Real-Life' Images: Detection, Alignment, and Recognition, 2008.

[32] J. He, D. J. Zhang, L. Balzano, T. Tao, Iterative online subspace learning for robust image alignment, in: Face and Gesture Recognition (FG), 2013 IEEE 10th Conference on, 2013.

[33] R. Li, R. Chellappa, Spatiotemporal alignment of visual signals on a special manifold, Pattern Analysis and Machine Intelligence, IEEE Transactions on 35 (3) (2013) 697–715. `doi:10.1109/TPAMI.2012.144`.