

Online Algorithms for Factorization-Based Structure from Motion

Ryan Kennedy, Laura Balzano, Stephen J. Wright, and Camillo J. Taylor

Abstract

We present a family of online algorithms for real-time factorization-based structure from motion, leveraging a relationship between incremental singular value decomposition and recently proposed methods for online matrix completion. Our methods are orders of magnitude faster than previous state of the art, can handle missing data and a variable number of feature points, and are robust to noise and sparse outliers. We demonstrate our methods on both real and synthetic sequences and show that they perform well in both online and batch settings. We also provide an implementation which is able to produce 3D models in real time using a laptop with a webcam.

Index Terms

structure from motion, matrix completion, incremental singular value decomposition



1 INTRODUCTION

THE problem of structure from motion — recovering the 3D structure of an object and locations of the camera from a monocular video stream — has been studied extensively in computer vision. For the rigid case, many algorithms are based on the seminal work of Tomasi and Kanade [1], in which it was shown that a noise-free measurement matrix of point tracks has rank at most 3 for an affine camera when the data are centered at the origin. The 3D locations of all tracked points and camera positions can be easily obtained from a factorization of this matrix. Due to occlusion, however, the matrix is typically missing many entries, so standard matrix factorization techniques cannot be applied.

Recent work in *low-rank matrix completion* has explored conditions under which the missing entries in a low-rank matrix can be determined, even when the matrix is corrupted with noise or sparse outliers [2], [3], [4], [5]. Most algorithms for matrix completion are static in nature, working with a “batch” of matrix columns. In this paper, by contrast, we focus on *online* structure from motion, in which the 3D model of point locations must be updated in real time as the video is taken. The algorithm must be efficient enough to run in real time, yet still must deal effectively with missing data and noisy observations. The online problem has received little attention in comparison to batch algorithms, although online algorithms have been developed for matrix completion [6] that have shown promise in other real-time computer vision applications [7]. In this paper, we extend these online algorithms to the problem of rigid structure from motion. Our algorithms naturally address difficulties that have been observed in this field, specifically: (1) our method is inherently online, (2) we naturally deal with data that are offset from the origin, (3) we directly deal with missing data, (4) we are able to handle a dynamically changing number of features, (5) our algorithms can be made robust to outliers, and (6) our method is extremely fast.

In addition to testing with online data, we compare our methods to batch algorithms, showing that they are competitive with – and often orders of magnitude faster than – state-of-the-art batch algorithms. To demonstrate the utility of our approach, we describe a laptop implementation of our algorithm that is able to create 3D models of objects in real time using video from an attached webcam.

2 RELATED WORK

2.1 Structure From Motion

Much research on the rigid structure-from-motion problem is based on Tomasi and Kanade’s factorization algorithm [1] for orthographic cameras, and subsequent work [8] that extended its applicability to other camera models [9], [10], [11], [12], [13], [14]. Comparatively little work has been done on *online* structure from motion,

• Here I believe we say who we are, who is the corresponding author, etc.

apart from algorithms that employ batch methods or local bundle adjustment in an online framework [15], [16].

In [13], Mortia and Kanade proposed a sequential version of the factorization algorithm but cannot deal with missing data, nor can they handle outliers or a dynamically changing set of features. The approach of McLauchlan and Murray [17] can handle missing data but uses simplifying heuristics to achieve a low computational complexity. The related algorithm of Trajković and Hedley [18] dispenses with the heuristics, but they focus on tracking moving objects within a scene. More recently, Cabral et al. [19] proposed a method that performs matrix completion iteratively, in an online manner. However, it is difficult to add new features dynamically, and their approach can be numerically unstable. The algorithm most similar to our own is the incremental SVD (ISVD) approach of Bunch and Nielsen [20], which has been previously used for sequential structure from motion by Brand [21]. In Section 4.2 we show how ISVD is related to our algorithms.

2.2 Matrix Completion and Subspace Tracking

Low-rank matrix completion is the problem of recovering a low-rank matrix from an incomplete sample of the entries. It was shown in [3], [22] that under assumptions on the number of observed entries and on incoherence of the singular vectors of this matrix with respect to the canonical coordinate axes, the nuclear norm minimization convex optimization problem solves the NP-hard rank minimization problem exactly. Since this breakthrough, a flurry of research activity has centered around developing faster algorithms to solve this convex optimization problem, both exact and approximate; see [23], [24] for two examples. The online algorithm Grouse [6] (Grassmannian Rank-One Update Subspace Estimation) outperforms all non-parallel algorithms in computational efficiency, often by an order of magnitude, while remaining competitive in terms of estimation error.

The Grouse algorithm was also developed for low-dimensional subspace tracking, which has been an active area of research. Comprehensive reference lists for complete-data adaptive methods for tracking subspaces and extreme singular values and vectors of covariance matrices can be found in [25], [26]; the authors discuss methods from the matrix computation literature as well as gradient-based methods from the signal processing literature.

A useful extension of these algorithms is toward “Robust PCA,” in which we seek to recover a low-rank matrix in the presence of outliers. This work has had application in computer vision, decomposing a scene from several video frames 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 [4], [27]. The algorithm Grasta [7], [28] (Grassmannian Robust Adaptive Subspace Tracking Algorithm) is a robust extension of Grouse that performs online estimation of a changing low-rank subspace while subtracting outliers.

3 PROBLEM DESCRIPTION

Given a set of points tracked through a video, the *measurement matrix* W is defined as

$$W = \begin{bmatrix} x_{1,1} & \cdots & x_{1,m} \\ \vdots & \ddots & \vdots \\ x_{n,1} & \cdots & x_{n,m} \end{bmatrix}, \quad (1)$$

where $x_{i,j} \in \mathbb{R}^{1 \times 2}$ is the projection of point i in frame j , giving W a dimension of $n \times 2m$, where n is the number of points and m is the number of frames. If we assume that the data are centered at the origin, then in the absence of noise this matrix has rank at most 3 and can be factored as

$$W = \tilde{S}\tilde{M}^T, \quad (2)$$

where \tilde{S} is the $n \times 3$ *structure* matrix containing 3D point locations, while \tilde{M} is the $2m \times 3$ *motion* matrix, which is the set of affine camera matrices corresponding to each frame [1]. Data that are offset from the origin have an additional translation vector τ ; we can write

$$W = [\tilde{S} \quad \mathbf{1}] [\tilde{M} \quad \tau]^T = SM^T, \quad (3)$$

where W now has rank at most 4 and the constant ones vector is necessarily part of its column space. We use this fact to naturally deal with offset data in our algorithms (Section 5.1).

In the presence of noise, one can use the singular value decomposition (SVD) [1] in order to find S and M such that SM^T most closely approximates W in either the Frobenius or operator norm. In this paper we address the problem of online structure from motion in which we have an estimated factorization at time t — $\hat{W}_t = \hat{S}_t \hat{M}_t^T$ — and wish to update our estimate at time $t + 1$ as we track points into the next video frame.

The new information takes the form of two additional columns for W_t , leading to successive updates of the form $W_{t+1} = [W_t \ v_t]$. The algorithm for updating the approximate factorization must be efficient enough to be used in real time, able to handle missing data as points become occluded, and able to incorporate new points.

4 MATRIX COMPLETION ALGORITHMS FOR SFM

4.1 The Grouse Algorithm [6]

Here we review the Grouse algorithm in the context of structure from motion in order to set the stage for our proposed algorithms, which are given in Section 4.2.

As in Section 3, we use the subscript t to denote values at time t . Let n_t be the number of tracks started up to time t and $U_t \in \mathbb{R}^{n_t \times 4}$ be an orthogonal matrix whose columns span the rank-4 column space of the measurement matrix W_t , and let v_t be a new column such that $W_{t+1} = [W_t \ v_t]$. (We consider the two new columns provided by each video frame one at a time.) The indices of v_t that are observed are given by $\Omega_t \subseteq \{1, \dots, n_t\}$ such that v_{Ω_t} and U_{Ω_t} are row submatrices of v_t and U_t corresponding to Ω_t .

The Grouse algorithm [6] measures the error between the current subspace-spanning matrix U_t and the new vector v_t using the ℓ_2 distance

$$\mathcal{E}(U_{\Omega_t}, v_{\Omega_t}) = \|U_{\Omega_t} w_t - v_{\Omega_t}\|_2^2, \quad (4)$$

where¹

$$w_t = U_{\Omega_t}^+ v_{\Omega_t} \quad (5)$$

is the set of weights that project v_{Ω_t} orthogonally onto the subspace given by U_{Ω_t} . We would like to replace U_t by an updated orthogonal matrix of the same dimensions, that reduces the error \mathcal{E} in accordance with the new observation v_t . Denoting by r_t the residual vector for the latest observation, we define the subvector corresponding to the indices in Ω_t by

$$r_{\Omega_t} = v_{\Omega_t} - U_{\Omega_t} w_t, \quad (6)$$

and set the components of r_t whose indices are not in Ω_t to zero. We can express the sensitivity of the error \mathcal{E} to U_t as follows:

$$\frac{\partial \mathcal{E}}{\partial U_t} = -2r_t w_t^T, \quad (7)$$

Grouse essentially performs projected gradient descent on the Grassmann manifold, taking a step in the negative gradient direction while staying on the manifold to maintain orthonormality. For details, see [6].

Because it is known *a priori* that $\text{rank}(W) \leq 4$, Grouse can be applied directly to the factorization problem, but there are several issues that prevent Grouse from being easily used for online structure from motion. First, Grouse only maintains an estimate of the column space of W_t , so the corresponding matrix $R_t \in \mathbb{R}^{2m_t \times 4}$ for which $\hat{W}_t = U_t R_t^T$ must be computed whenever a final reconstruction is needed. This can be a problem for online applications since it requires keeping all data until the algorithm is complete. Additionally, the choice of a critical parameter in Grouse — the step size for gradient descent — can affect the rate of convergence strongly. In the following section, we take advantage of a recent result on the relationship between Grouse and incremental SVD to resolve these difficulties.

4.2 The Incremental SVD (ISVD) Formulation

The incremental SVD algorithm [20] is a simple method for computing the SVD of a collection of data by updating an initial decomposition one column at a time. Given a matrix W_t of rank k whose SVD is $W_t = U_t \Sigma_t V_t^T$, we wish to compute the SVD of a new matrix with a single column added: $W_{t+1} = [W_t \ v_t]$. Defining $w_t = U_t^T v_t$ and $r_t = v_t - U_t w_t$, we have

$$W_{t+1} = \begin{bmatrix} U_t & \frac{r_t}{\|r_t\|} \end{bmatrix} \begin{bmatrix} \Sigma_t & w_t \\ 0 & \|r_t\| \end{bmatrix} \begin{bmatrix} V_t^T & 0 \\ 0 & 1 \end{bmatrix}, \quad (8)$$

where one can verify that the left and right matrices are still orthogonal. We compute an SVD of the center matrix:

$$\begin{bmatrix} \Sigma_t & w_t \\ 0 & \|r_t\| \end{bmatrix} = \tilde{U} \tilde{\Sigma} \tilde{V}^T, \quad (9)$$

1. Our notation $A^+ b$ denotes the least-residual-norm solution to the least-squares problem $\min_x \|Ax - b\|_2^2$, obtainable from a singular value decomposition of A or a factorization of $A^T A$.

where \tilde{U} , $\tilde{\Sigma}$, and \tilde{V} are of size $(k+1) \times (k+1)$. If only the first k singular vectors are needed, then as a heuristic the smallest singular value and its associated singular vectors can be dropped. Let \hat{U} , $\hat{\Sigma}$, and \hat{V} be the resulting $(k+1) \times k$, $k \times k$ and $(k+1) \times k$ matrices. The updated rank- k SVD estimate is given by $W_{t+1} = U_{t+1}\Sigma_{t+1}V_{t+1}^T$, where

$$U_{t+1} = \begin{bmatrix} U_t & \frac{r_t}{\|r_t\|} \end{bmatrix} \hat{U} \quad ; \quad \Sigma_{t+1} = \hat{\Sigma} \quad ; \quad V_{t+1} = \begin{bmatrix} V_t & 0 \\ 0 & 1 \end{bmatrix} \hat{V}. \quad (10)$$

In the case of missing data when only entries $\Omega_t \subseteq \{1, \dots, n_t\}$ are observed, we can define the weights and residual vector in these update formulae as in Equations (5) and (6), respectively.

We now examine the relationship of this algorithm to Grouse, in the context of SFM. Let $\hat{W}_t = U_t R_t^T$ be an estimated rank-4 factorization of W_t such that U_t has orthonormal columns. Given a new column v_t with observed entries Ω_t , if w_t and r_t are the least-squares weight and residual vector, respectively, defined with respect to the set of observed indices Ω_t as in Equations (5) and (6), then we can write

$$\begin{bmatrix} U_t R_t^T & \tilde{v}_t \end{bmatrix} = \begin{bmatrix} U_t & \frac{r_t}{\|r_t\|} \end{bmatrix} \begin{bmatrix} I & w_t \\ 0 & \|r_t\| \end{bmatrix} \begin{bmatrix} R_t & 0 \\ 0 & 1 \end{bmatrix}^T, \quad (11)$$

where $\tilde{v}_t \in \mathbb{R}^{n_t}$ and the subvector of \tilde{v}_t corresponding to Ω_t is set to v_{Ω_t} , while the remaining entries in v_t are imputed as the inner product of w_t and the rows $i \notin \Omega_t$ of U_t . Stated again in mathematics, letting $[\tilde{v}_t]_i$ refer to the i^{th} component of the vector \tilde{v}_t , we have:

$$[\tilde{v}_t]_i := \begin{cases} [v_t]_i & i \in \Omega_t \\ [U_t w_t]_i & i \in \Omega_t^C \end{cases}.$$

Define the SVD of the center matrix to be

$$\begin{bmatrix} I & w_t \\ 0 & \|r_t\| \end{bmatrix} = \tilde{U} \tilde{\Sigma} \tilde{V}^T. \quad (12)$$

Let \hat{U} be the 5×4 matrix obtained by dropping the last column of \tilde{U} , corresponding to the smallest singular value of $\tilde{\Sigma}$. Similarly, let \hat{V} be 5×4 matrix resulting from dropping the last column of \tilde{V} , and let $\hat{\Sigma}$ be the 4×4 diagonal matrix obtained by dropping the last column and row from $\tilde{\Sigma}$.

In [29], it was shown that updating U_t to

$$U_{t+1} = \begin{bmatrix} U_t & \frac{r_t}{\|r_t\|} \end{bmatrix} \hat{U} \quad (13)$$

is equivalent to Grouse for a particular data-dependent choice of step size. Additionally, combining Equations (11) and (12), R_t is updated as

$$R_{t+1} = \begin{bmatrix} R_t & 0 \\ 0 & 1 \end{bmatrix} \hat{V} \hat{\Sigma}. \quad (14)$$

The result is a new rank-4 factorization $\hat{W}_{t+1} = U_{t+1} R_{t+1}^T$. This algorithm is again described as Algorithm 1 below.

The advantages of this method are two-fold. First, by updating both U and R simultaneously, there is no need to calculate R when a reconstruction is needed. Instead, we keep a running estimate of both U and R ; estimates of the motion and structure matrices can be easily found at any point in time. We thus have a truly online SFM approach because we no longer need to store the entire observation matrix W in order to solve for R . Keeping a subset of the data is still useful so that old data can be revisited, but this can be limited to a fixed amount if memory becomes an issue. Second, this formulation uses an implicit step size; we no longer are required to specify a step size as in the original Grouse formulation, although the residual vector can still be scaled to affect the step size if needed. In Section 6.2, we show that the parameter-free algorithm outperforms other online SFM algorithms on real data. In Section 6.1, we show that the scaled version markedly outperforms all other batch algorithms for SFM on both real and synthetic data in terms of both error and speed.

4.3 Robust SFM

The Grasta algorithm [7], [28] is an extension of Grouse which is robust to outliers. Instead of minimizing the ℓ_2 cost function given in Equation (4), Grasta uses the robust ℓ_1 cost function $\mathcal{E}_{grasta}(U_{\Omega_t}, v_{\Omega_t}) = \|U_{\Omega_t} w_t - v_{\Omega_t}\|_1$. Grasta estimates the weights w_t of this ℓ_1 projection as well as the sparse additive component using ADMM [30] and then updates U_t using Grassmannian geodesics, replacing r_t with a variable Γ_t which is a function of the sparse component and the ℓ_1 weights [7]. Thus, we can use these w_t and Γ_t in place of w_t and r_t in the ISVD formulation of Grouse (Section 4.2), resulting in an ISVD formulation of Grasta.

4.4 The Resulting Family of Algorithms

The algorithms that we have presented so far make up a family of algorithms that can be applied to SFM in various circumstances. Both Grouse and Grasta are studied here, as well as the ISVD versions which carry forward an estimate for the singular values Σ .

To obtain a final variant within this family, we scale the residual norm in (12) by some value α_t before taking its SVD:

$$\begin{bmatrix} I & w_t \\ 0 & \alpha_t \|r_t\| \end{bmatrix} = \hat{U} \hat{\Sigma} \hat{V}^T. \quad (15)$$

This scaling provides us with more control over the contribution of the residual vector to the new subspace estimate. In our experiments, we took α_t to be decreasing with t . This scaling is useful in the batch setting, but with online data we found that no scaling performs best.

We present in Algorithm 1 one iteration of the algorithm from which our family of algorithms is derived. In Table 1 the reader will find a description of all proposed algorithms relative to this meta-algorithm, along with their names as used in the experimental results.

Algorithm 1 One Iteration of the Meta-Algorithm

At time t , given U_t , an $n_t \times k$ (the rank $k = 4$ for this paper) orthonormal matrix, R_t , an $m_t \times k$ matrix, and scaling parameter α_t ;

Take the new column v_{Ω_t} with only points indexed by Ω_t observed;

Define

$$w_t := \arg \min_w \|U_{\Omega_t} w - v_{\Omega_t}\|; \quad (16)$$

Define

$$[\tilde{v}_t]_i := \begin{cases} [v_t]_i & i \in \Omega_t \\ [U_t w_t]_i & i \in \Omega_t^C \end{cases} \quad ; \quad p_t := U_t w_t \quad ; \quad r_t := \tilde{v}_t - p_t; \quad (17)$$

Noting that

$$\begin{bmatrix} U_t R_t^T & \tilde{v}_t \end{bmatrix} = \begin{bmatrix} U_t & \frac{r_t}{\|r_t\|} \end{bmatrix} \begin{bmatrix} I & w_t \\ 0 & \|r_t\| \end{bmatrix} \begin{bmatrix} R_t & 0 \\ 0 & 1 \end{bmatrix}^T$$

we compute the SVD of the update matrix with singular values in decreasing order:

$$\begin{bmatrix} I & w_t \\ 0 & \alpha_t \|r_t\| \end{bmatrix} = \tilde{U}_t \tilde{\Sigma}_t \tilde{V}_t^T, \quad (18)$$

Define \hat{U}_t to be the $(k+1) \times k$ matrix obtained by removing the last column from \tilde{U}_t . Similarly, remove the last column of \tilde{V}_t to obtain \hat{V}_t and remove the last column and row from $\tilde{\Sigma}_t$ to obtain the $k \times k$ matrix $\hat{\Sigma}_t$.

Set

$$U_{t+1} := \begin{bmatrix} U_t & \frac{r_t}{\|r_t\|} \end{bmatrix} \hat{U}_t. \quad (19)$$

Set

$$R_{t+1} = \begin{bmatrix} R_t & 0 \\ 0 & 1 \end{bmatrix} \hat{V}_t \hat{\Sigma}_t. \quad (20)$$

The result is a new rank- k factorization $\hat{W}_{t+1} = U_{t+1} R_{t+1}^T$.

5 IMPLEMENTATION

5.1 The All-1s Vector

A few more issues need to be addressed for implementation of Grouse for SFM. The first issue is to exploit the fact that the constant vector of ones, which we will denote as $\mathbb{1}$, should always be in the column space of W_t and thus should be in the span of any estimate U_t of the column space of W_t , since the points may

A Family of Matrix Completion SFM Algorithms		
	Grouse (ℓ_2 cost)	Grasta (ℓ_1 cost)
Track singular values Σ	Eqns (9),(10), "isvd"	(9),(10), using w_t from ADMM, and replacing r_t with Γ_t [7], "isvd-grasta"
Use $\Sigma = I$ and setting $\alpha_t = 1$	Eqns (18),(19),(20), "grouse"	(18),(19), (20) using w_t from ADMM, and replacing r_t with Γ_t [7], "grasta"
Use $\Sigma = I$ and scale residual by $\alpha_t \propto \frac{1}{t}$	Eqns (18),(19),(20) "grouse (scaled)"	(18),(19),(20) using w_t from ADMM, and replacing r_t with Γ_t [7], "grasta (scaled)"

TABLE 1: Update steps for our family of matrix completion algorithms for structure from motion. The names in quotes correspond to names used in our experiments.

be offset from the origin. Without loss of generality, suppose that U_t has the ones-vector as its last column (appropriately scaled), so that

$$U_t = [\bar{U}_t \quad \mathbb{1}/\sqrt{n_t}], \quad R_t = [\bar{R}_t \quad \tau_t \sqrt{n_t}]. \quad (21)$$

Here τ_t is the corresponding translation vector, which is the (scaled) last column of R_t . Similarly, let $w_t = [\bar{w}_t^T \quad \gamma_t]^T$. The derivative of the error \mathcal{E} in Equation (4) with respect to just the first three columns is

$$\frac{\partial \mathcal{E}}{\partial \bar{U}_t} = -2r_t \bar{w}_t^T. \quad (22)$$

By not considering the derivative of \mathcal{E} with respect to the ones vector it will remain in the span of U , since the Grouse update will be applied only to \tilde{U} . Our Grouse update for structure from motion is obtained by first setting

$$\bar{U}_{t+1} = \begin{bmatrix} \bar{U}_t & \frac{r_t}{\|r_t\|} \end{bmatrix} \tilde{U}, \quad (23)$$

$$\bar{R}_{t+1} = \begin{bmatrix} \bar{R}_t & 0 \\ 0 & 1 \end{bmatrix} \tilde{V} \tilde{\Sigma}, \quad (24)$$

$$\tau_{t+1} = \begin{bmatrix} \tau_t \\ \gamma_t/\sqrt{n_t} \end{bmatrix}, \quad (25)$$

(where \tilde{U} , $\tilde{\Sigma}$, and \tilde{V} are defined as in Equation (12) using \bar{w}_t), and then dropping the last column of \bar{U}_{t+1} and \bar{R}_{t+1} . Because the residual vector r_t is still based on the full matrix U_t , including the ones-vector, r_t will necessarily be orthogonal to the ones-vector. Therefore, since

$$U_{t+1} = [\bar{U}_{t+1} \quad \mathbb{1}/\sqrt{n_t}],$$

U_{t+1} will retain the ones-vector in its span and will still have orthonormal columns.

5.2 Adding New Points

In online structure from motion, we are initially unaware of the total number of points to be tracked and need to account for newly added points as the video progresses. If $U_t \in \mathbb{R}^{n_t \times 4}$ is the current subspace estimate, then new points will manifest themselves as additional rows of U_t . However, when updating U_t , we have to make sure that the columns of U_t remain orthonormal and the last column continues to be the vector of ones. We thus perform the following update when each new point is added, where we increment $n_{t+1} = n_t + 1$ so that the columns of U_t remain orthonormal:

$$U_t \leftarrow \begin{bmatrix} \bar{U}_t & \mathbb{1}/\sqrt{n_t+1} \\ 0 & 1/\sqrt{n_t+1} \end{bmatrix}. \quad (26)$$

5.3 Updating Past Points

The Grouse update for structure from motion is fast enough that many updates can be done for each new video frame. It is therefore advantageous to be able to revisit old frames and reduce the error more than would be possible using a single pass over the frames. Using the original Grouse formulation described in Section 4.1, we simply run additional Grouse updates using past columns of W . This does not work in the new

online formulation, where we also keep track of the matrix R_t , since running another Grouse update will add a new row to R_t . Instead, we simply drop the associated row of R_t before the update and replace it with the resulting new row. Because we do not impose any orthogonality restrictions on the matrix R_t , no correction is needed. In our experiments, however, we also compare to ISVD, which requires that the the right-side matrix be orthogonal. In this case, we first “downdate” our SVD using the algorithm given by Brand [31] before performing an update.

6 EXPERIMENTS

We used three different datasets for comparison: the Dinosaur sequence from [32], the Teddy Bear sequence from [14], and a sequence of a Giraffe that we created ourselves. The Dinosaur sequence consists of 4983 point tracks over 36 frames. We use the full data matrix and do not remove low-quality tracks. The Teddy Bear sequence has 806 tracks over 200 frames and the Giraffe has 2634 tracks over 343 frames. The three sequences have measurement matrices that are missing 91%, 89%, and 93% of their entries, respectively. Our reconstruction results for these datasets are illustrated in Figure 1. Since no ground truth is available, we report 2D RMSE error. We found this error measure to be a good indicator of the quality of the reconstruction, and results were similar to a synthetic cylinder dataset for which 3D RMSE errors could be calculated. All 3D reconstructions were created by imposing the standard scaled-orthographic metric constraints.

6.1 Batch SFM

We first consider the batch or offline setting where the entire measurement matrix W is input. In this section we focus on the scaled version of our algorithms, because these allow us to diminish the stepsize with continuing passes over the data, guaranteeing more stable convergence. We compare the scaled versions of our algorithms to `pf` [33], `dn` [10], `ga` [12], `csf` [11] and `balm` [34]. For `balm`, we used the scaled-orthographic projector. Each algorithm was run until convergence. All of the methods we compared to are in some way iterative and require initialization, so in order to fairly compare across different methods, we initialized each algorithm to the same random subspace and averaged over 5 random initializations.

Results are shown in Figure 2, where we plot the root-mean-squared error of each algorithm with respect to the measurement matrix over time. Notice that our algorithm `grouse` (scaled) converges significantly faster than most other batch algorithms, with `pf` being the only competitor in terms of convergence rate. However, we found that `pf` was sensitive to initial conditions and did not always converge to the same accuracy as `grouse`, especially for the Dinosaur sequence (Figure 2a). `grouse` also outperforms the more robust `grasta`. For real datasets, then, rather than using a robust algorithm such as `grasta`, it may be preferable to manually remove any extreme outliers and use the faster `grouse` algorithm, especially if the data is of relatively high quality to begin with.

6.1.1 3D Errors

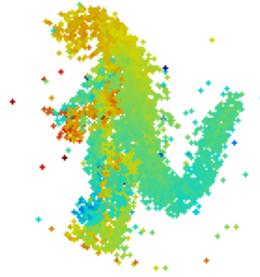
We ran all batch algorithms for 60 seconds on the Synthetic Cylinder. The resulting reconstructions for `grouse` and `damped-newton` (which was the best algorithm from the literature after 60 seconds) are shown in Figure 4. Surprisingly, no other algorithm was able to fully reconstruct the cylinder, which has relatively little noise, even after a full minute of computation. In fact, not all algorithms converged within the full minute, though they may reach lower errors if given more time or a better initial point. In particular, the `damped-newton` method shown in Figure 4b converged after 3.5 minutes to the correct cylinder. We contend, however, that the slowness of these algorithms would make it difficult to use them for a real-time application. In contrast, our method converges rapidly to a good solution from a completely random initialization.

6.1.2 Large Dataset Experiment

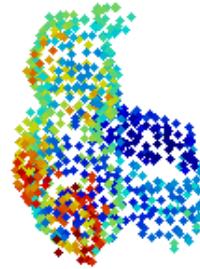
To test our algorithm on a large dataset and directly compare to [34], we used the Venice dataset from Agarwal et. al [35], as modified in [34]. The 3D model was projected onto 100 random orthographic cameras, producing a measurement matrix of size 939551×200 from which 90% of the entries were randomly removed. We measure the error as $\|S - S_{gt}\|_F / \|S_{gt}\|_F$, where S is the resulting 3D reconstruction and S_{gt} is the ground-truth 3D model. The algorithm of [34] reported a reconstruction error of 6.67%, which `grouse` achieved within 50 seconds. `grouse` had reduced the error to 2.48% after 2 minutes, and to 1.37% after 10 minutes.



(a) Dinosaur



(b) Teddy Bear



(c) Giraffe



Fig. 1: Dinosaur, Teddy Bear and Giraffe datasets used in our experiments. We show one video frame and the reconstruction from our online algorithm as viewed from the estimated camera model of each frame. Color indicates depth.

6.2 Online SFM

For online structure-from-motion, we evaluate our algorithm `grouse` along with its robust counterpart `grasta`, and also evaluate the variants of both algorithms in which residual vectors for past frames are scaled (`grouse (scaled)` and `grasta (scaled)`). We also compare to incremental SVD (`isvd`) and a robust version of this approach, which we denote `isvd-grasta`. Our `isvd` algorithm is the same as in [21].

In [6], it was shown that for a static subspace, the step size for `grouse` should diminish over time. With our formulation, we still have control of the step size by scaling the residual; here, at iteration t we scale by the function C/t for a fixed C . We find that this is important for convergence when `grouse` is run in batch mode (Section 6.1). For tests of the online methods, we show results both with and without scaling the residual for past frames, and do not scale the residual for new frames.

Our algorithms were implemented in Matlab and were initialized by finding all points that were fully tracked for the first 5 frames and running SVD on this subset of the measurement matrix.

Results for online SFM are shown in Figure 5. Each algorithm was run for enough iterations so that it maintained a fixed frame rate, which we varied from 1-100 fps. The plots show the root-mean-squared pixel

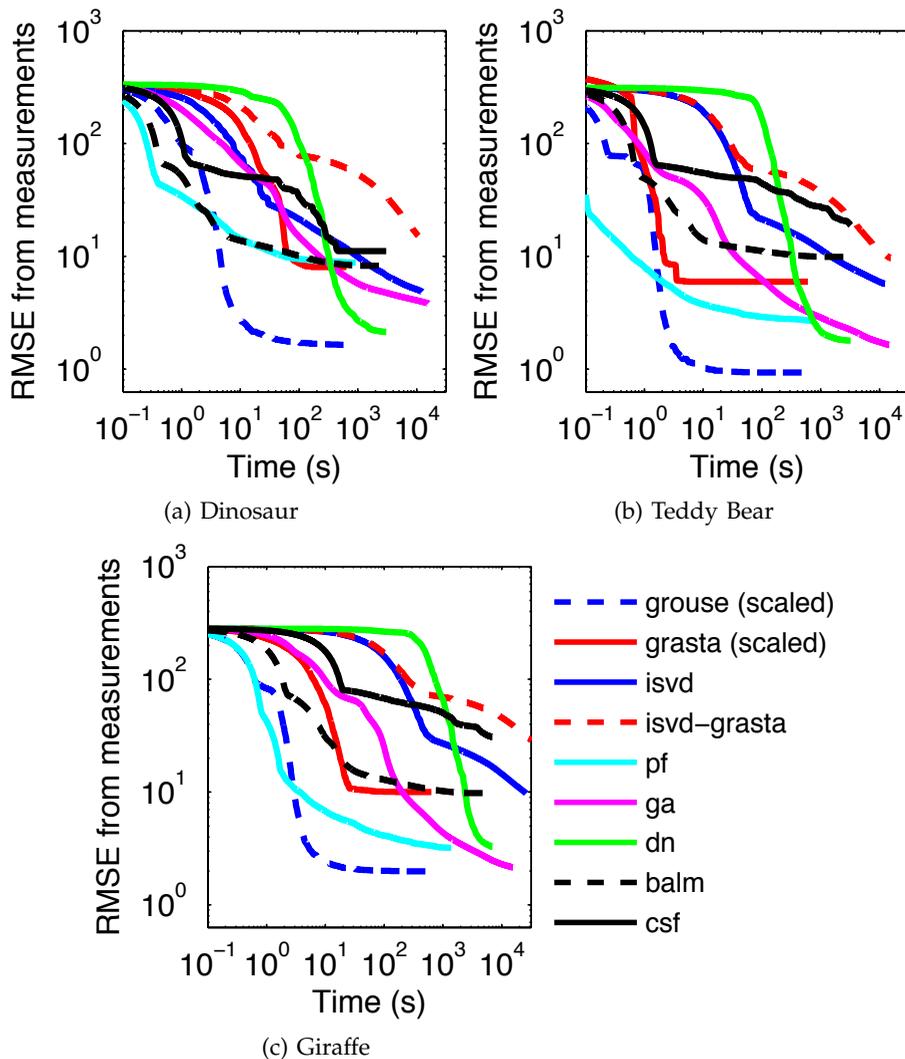


Fig. 2: Comparison of batch SFM algorithms. We show how the root-mean-squared pixel error (RMSE) varies over time for each algorithm, averaged over 5 random initializations. Note that these results are shown on a log-log scale. Power-factorization (pf) has lower error in the first second of execution, but *grouse (scaled)* converges in 5-10 seconds to lower error than the other algorithms reach after 2+ hours.

error (RMSE) from the measurements after all frames have been processed for a given frame rate. For all datasets, *grouse* performed better than either its robust *grasta* counterpart or the *isvd* algorithms. The other methods do eventually catch up to *grouse* as the frame rate is slowed, but *grouse* is by far the best-performing algorithm for frame rates that would correspond to videos streams at 15-30 fps.

6.2.1 Real-time implementation

To demonstrate the utility of our method, we implemented *grouse* in C++ using OpenCV. Our implementation uses two threads running on separate cores: one thread reads frames from an attached webcam and tracks point using the KLT tracker in OpenCV, while the other thread continually runs *grouse* and incorporates new data as it becomes available. We used a MacBook Pro with a 2.66 GHz Intel Core 2 Duo processor and 8 GB of memory and a Logitech C270 webcam.

Our implementation was run at 15 fps and we used it to capture and reconstruct a 3D model of a toy giraffe in real time (Figure 1c). The bottleneck in this process is tracking points between frames, and during this time *grouse* completed an average of 205 iterations per frame where each iteration of *grouse* was run on a randomly-selected past column of the measurement matrix. We envision this algorithm being useful in applications where it is desirable to build a 3D model with low-cost hardware, such as at-home 3D printing; using a simple webcam and a standard computer, we can obtain a 3D model of an object in real-time.

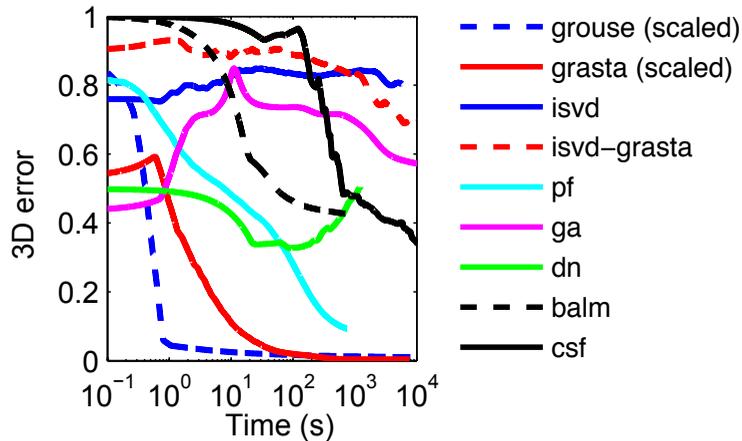


Fig. 3: Comparison of batch SFM algorithms on a synthetic cylinder. The matrix is of size 748×1000 with 66.5% missing data. Let S and S_{gt} be the reconstructed structure and the ground-truth structure respectively, both $n \times 3$ matrices of 3D point locations. S is aligned with the ground truth S_{gt} using Procrustes, and then the error is calculated as $\|S - S_{gt}\|_F / \|S_{gt}\|_F$. Results are averages over 5 random initializations. Only `grouse (scaled)` and `grasta (scaled)` are consistently able to reconstruct the cylinder correctly.

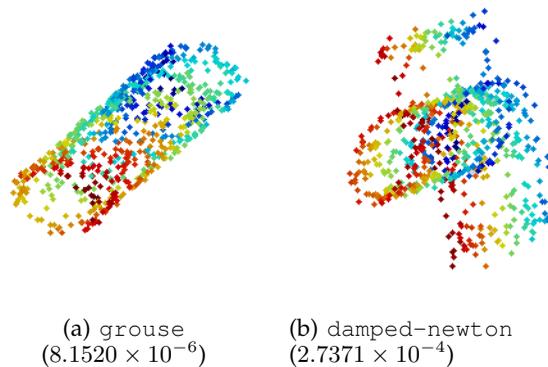


Fig. 4: Reconstructions using `grouse` and the best algorithm from the literature (`damped-newton`) for batch SFM on the Synthetic Cylinder. Each algorithm was given one minute to run. While the `grouse` solution is very close to the ground truth, the `damped-newton` method does not fully wrap into a cylinder. The sum-of-squares error after Procrustes alignment to the ground-truth cylinder is shown in parentheses. Color indicates depth.

6.3 Effect of Noise

Using synthetic data allows us to systematically investigate the effect of noise on our method. We do so by generating a noise-free cylinder dataset, and then adding variable amounts of noise to it. The synthetically-generated cylinder has a radius of 10 pixels and a height of 50 pixels, centered at (50, 50). Two hundred random points were tracked over 100 frames while the cylinder underwent one full rotation. The measurement matrix for this cylinder has 66% of its entries missing.

The results in Figure 6a show how our algorithm performs when run in batch when Gaussian noise with different standard deviations is added. Each algorithm was run for 60 seconds. We plot the resulting mean-squared pixel error (MSE) from the ground-truth, noise-free, measurement matrix. Similarly, we added sparse noise uniformly distributed in the interval $[-50, 50]$ to a variable proportion of observations, and show the results in Figure 6b. When only Gaussian noise is present, we find that the non-robust `grouse` gives the best results, although both `grouse` and `grasta` perform well. With sparse noise, `grasta` is preferable.

Because `grouse` is much faster than `grasta`, it is better to use `grouse` if the point tracks are of high quality. Even when bad tracks are present, due to the speed advantage it may be preferable to improve the

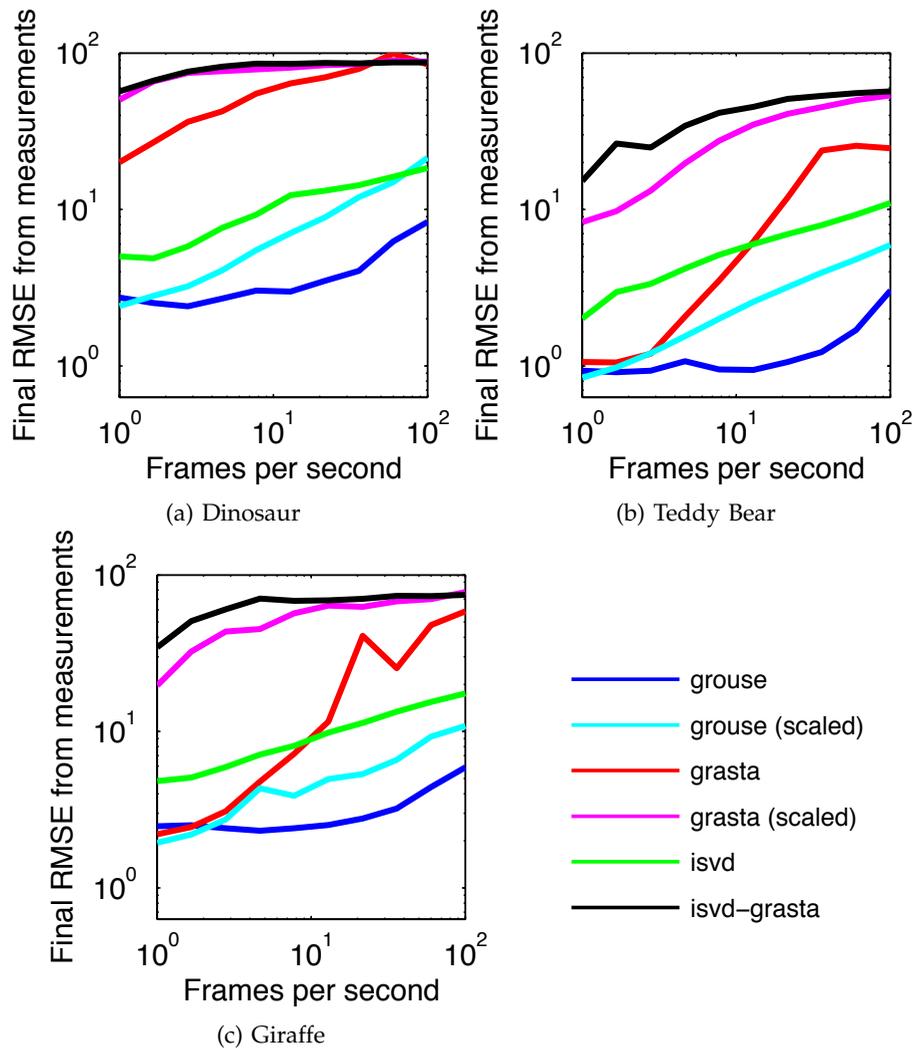


Fig. 5: Comparison of online SFM algorithms. Each algorithm was run with enough iterations to ensure it ran at a fixed frame rate. We plot the final root-mean-squared pixel error (RMSE) between the estimated and actual measurement matrices as the number of frames per second is varied. `grouse` outperforms all other versions of our online SFM algorithm.

tracking performance or throw out bad tracks rather than use `grasta`.

7 DISCUSSION

We have presented a family of online algorithms for factorization-based structure from motion, opening the door to real-time recovery of 3D structure from a single camera video stream on low-power computers such as cell phones. Our algorithms have state-of-the-art error performance in orders of magnitude less computation time for the structure from motion problem. There are many remaining questions for exploration, but one of immediate importance for the online matrix completion approach to SFM is a better understanding of the algorithm’s convergence properties and, in particular, their dependence on the sampled observations. Existing analyses assume that the observed subvector is chosen randomly and independently at each iteration. In SFM, however, there is a great deal of structure to the visible features over time and the observations are not independent at each iteration.

REFERENCES

[1] C. Tomasi and T. Kanade, “Shape and motion from image streams under orthography: a factorization method,” *International Journal of Computer Vision*, vol. 9, no. 2, pp. 137–154, 1992. 1, 2

[2] É. Candès and T. Tao, “The power of convex relaxation: Near-optimal matrix completion,” *IEEE Transactions on Information Theory*, vol. 56, no. 5, pp. 2053–2080, May 2010. 1

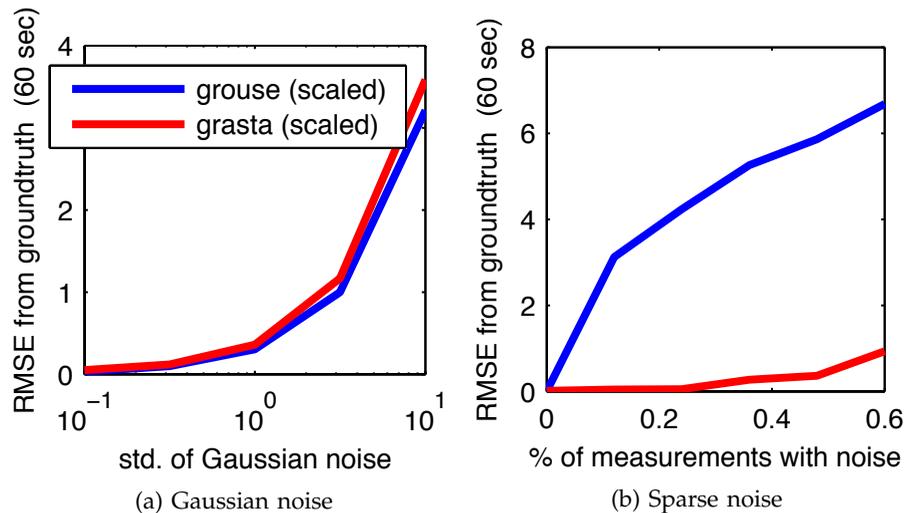


Fig. 6: Effect of noise on our algorithms. In (a), we generate a synthetic cylinder and add a variable level of Gaussian noise, while in (b) we add uniform random scaled noise in the interval $[-50, 50]$ to a variable fraction of the measurements. For Gaussian noise, grouse (scaled) slightly outperforms the more robust grasta (scaled), while grasta (scaled) has a clear advantage when there is sparse noise.

- [3] B. Recht, "A simpler approach to matrix completion," *Journal of Machine Learning Research*, vol. 12, pp. 3413–3430, 2011. 1, 2
- [4] E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" *Journal of the ACM*, vol. 58, no. 1, pp. 1–37, 2009. 1, 2
- [5] E. Candès and Y. Plan, "Matrix completion with noise," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 925–936, 2010. 1
- [6] L. Balzano, R. Nowak, and B. Recht, "Online identification and tracking of subspaces from highly incomplete information," in *Communication, Control, and Computing (Allerton)*. IEEE, 2010, pp. 704–711. 1, 2, 3, 8
- [7] J. He, L. Balzano, and A. Szelam, "Incremental gradient on the grassmannian for online foreground and background separation in subsampled video," in *Computer Vision and Pattern Recognition*, June 2012. 1, 2, 4, 6
- [8] C. Poelman and T. Kanade, "A paraperspective factorization method for shape and motion recovery," *Pattern Analysis and Machine Intelligence*, vol. 19, no. 3, pp. 206–218, 1997. 1
- [9] H. Aanaes, R. Fisker, K. Astrom, and J. Carstensen, "Robust factorization," *Pattern Analysis and Machine Intelligence*, vol. 24, no. 9, pp. 1215–1225, 2002. 2
- [10] A. Buchanan and A. Fitzgibbon, "Damped newton algorithms for matrix factorization with missing data," in *Computer Vision and Pattern Recognition (CVPR)*, vol. 2. IEEE, 2005, pp. 316–322. 2, 7
- [11] P. Gotardo and A. Martinez, "Computing smooth time trajectories for camera and deformable shape in structure from motion with occlusion," *Pattern Analysis and Machine Intelligence*, vol. 33, no. 10, pp. 2051–2065, 2011. 2, 7
- [12] R. Guerreiro and P. Aguiar, "3d structure from video streams with partially overlapping images," in *International Conference on Image Processing*, vol. 3. IEEE, 2002, pp. 897–900. 2, 7
- [13] T. Morita and T. Kanade, "A sequential factorization method for recovering shape and motion from image streams," *Pattern Analysis and Machine Intelligence*, vol. 19, no. 8, pp. 858–867, 1997. 2
- [14] J. Tardif, A. Bartoli, M. Trudeau, N. Guilbert, and S. Roy, "Algorithms for batch matrix factorization with application to structure-from-motion," in *Computer Vision and Pattern Recognition*. IEEE, 2007, pp. 1–8. 2, 7
- [15] G. Klein and D. Murray, "Parallel tracking and mapping for small ar workspaces," in *International Symposium on Mixed and Augmented Reality*. IEEE, 2007, pp. 225–234. 2
- [16] E. Mouragnon, M. Lhuillier, M. Dhome, F. Dekeyser, and P. Sayd, "Generic and real-time structure from motion using local bundle adjustment," *Image and Vision Computing*, vol. 27, no. 8, pp. 1178–1193, 2009. 2
- [17] P. McLauchlan and D. Murray, "A unifying framework for structure and motion recovery from image sequences," in *International Conference on Computer Vision*. IEEE, 1995, pp. 314–320. 2
- [18] M. Trajković and M. Hedley, "A practical algorithm for structure and motion recovery from long sequence of images," in *Image Analysis and Processing*. Springer, 1997, pp. 470–477. 2
- [19] R. Cabral, J. Costeira, F. De la Torre, and A. Bernardino, "Fast incremental method for matrix completion: an application to trajectory correction," in *International Conference on Image Processing*. IEEE, 2011, pp. 1417–1420. 2
- [20] J. R. Bunch and C. P. Nielsen, "Updating the singular value decomposition," *Numerische Mathematik*, vol. 31, pp. 111–129, 1978, 10.1007/BF01397471. [Online]. Available: <http://dx.doi.org/10.1007/BF01397471> 2, 3
- [21] M. Brand, "Incremental singular value decomposition of uncertain data with missing values," *European Conference on Computer Vision*, pp. 707–720, 2002. 2, 8
- [22] E. Candès and B. Recht, "Exact matrix completion via convex optimization," *Foundations of Computational Mathematics*, vol. 9, no. 6, pp. 717–772, December 2009. 2
- [23] R. Keshavan, A. Montanari, and S. Oh, "Matrix completion from noisy entries," *Journal of Machine Learning Research*, vol. 11, pp. 2057–2078, July 2010. 2
- [24] K.-C. Toh and S. Yun, "An accelerated proximal gradient algorithm for nuclear norm regularized least squares problems," *Pacific Journal of Optimization*, vol. 6, pp. 615–640, 2010. 2
- [25] L. Balzano, "Handling missing data in high-dimensional subspace modeling," Ph.D. dissertation, University of Wisconsin, Madison, 2012. 2
- [26] A. Edelman, T. A. Arias, and S. T. Smith, "The geometry of algorithms with orthogonality constraints," *SIAM Journal on Matrix Analysis and Applications*, vol. 20, no. 2, pp. 303–353, 1998. 2

- [27] G. Mateos and G. B. Giannakis, "Sparsity control for robust principal component analysis," in *Asilomar Conference on Signals, Systems, and Computers*, 2010. [2](#)
- [28] J. He, L. Balzano, and J. Lui, "Online robust subspace tracking from partial information," *Arxiv preprint arXiv:1109.3827*, 2011. [2](#), [4](#)
- [29] L. Balzano and S. J. Wright, "On grouse and incremental svd," *To appear in the Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, 2013. [4](#)
- [30] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–123, 2011. [4](#)
- [31] M. Brand, "Fast low-rank modifications of the thin singular value decomposition," *Linear algebra and its applications*, vol. 415, no. 1, pp. 20–30, 2006. [7](#)
- [32] A. Fitzgibbon, G. Cross, and A. Zisserman, "Automatic 3d model construction for turn-table sequences," *3D Structure from Multiple Images of Large-Scale Environments*, pp. 155–170, 1998. [7](#)
- [33] R. Hartley and F. Schaffalitzky, "Powerfactorization: 3d reconstruction with missing or uncertain data," in *Australia-Japan Advanced Workshop on Computer Vision*, vol. 74, 2003, pp. 76–85. [7](#)
- [34] A. Del Bue, J. Xavier, L. Agapito, and M. Paladini, "Bilinear modeling via augmented lagrange multipliers (balm)," *Pattern Analysis and Machine Intelligence*, vol. 34, no. 8, pp. 1496–1508, 2012. [7](#)
- [35] S. Agarwal, N. Snavely, S. M. Seitz, and R. Szeliski, "Bundle adjustment in the large," in *European Conference on Computer Vision*. Springer, 2010, pp. 29–42. [7](#)