

Robust Absolute Rotation Estimation via Low-rank and Sparse Matrix Decomposition

F. Arrigoni, L. Magri
University of Milan
Milano, Italy

federica.arrigoni@studenti.unimi.it
luca.magri@unimi.it

B. Rossi, P. Fragneto
STMicroelectronics
Agrate Brianza, Italy

beatrice.rossi@st.com
pasqualina.fragneto@st.com

A. Fusiello
University of Udine
Udine, Italy

andrea.fusiello@uniud.it

Abstract

This paper proposes a robust method to solve the absolute rotation estimation problem, which arises in global registration of 3D point sets and in structure-from-motion. A novel cost function is formulated which inherently copes with outliers. In particular, the proposed algorithm handles both outlier and missing relative rotations, by casting the problem as a “low-rank & sparse” matrix decomposition. As a side effect, this solution can be seen as a valid and cost-effective detector of inconsistent pairwise rotations. Computational efficiency and numerical accuracy, are demonstrated by simulated and real experiments.

1. Introduction

In this paper we deal with the Absolute Rotation Estimation (ARE) problem, i.e. the problem of recovering the *absolute* attitudes (rotations) – with respect to a global frame of reference – of a set of local reference frames, given their *relative* attitudes. These local frames can be camera reference frames, in which case we are in the context of structure-from-motion, or local coordinates where 3D points are represented, in which case we are dealing with a 3D point set registration problem. This problem is analyzed in depth in [12], under the name “rotation averaging”. Given a redundant number of *relative* rotations R_{ij} between image pairs, the goal is to compute N *absolute* rotations R_1, \dots, R_N (in a given absolute frame) by averaging the R_{ij} in order to satisfy the compatibility constraint $R_{ij} = R_i R_j^T$. This task finds application in the (global) structure-from-motion problem, and in the global registration problem.

Global registration (a.k.a. *N-view point set registration problem*) consists in finding the rigid transformation that brings multiple ($N > 2$) 3-D point sets into alignment. We are interested here only in the rotation component of the

transformation. Global registration can be solved in point (correspondences) space or in frame space. In the former case [20, 15], all the rotations are simultaneously optimized with respect to a cost function that depends on the distance of corresponding points. In the latter case [9, 21], the optimization criterion is related to the internal coherence of the network of rotations (and translations) applied to the local coordinates frame.

The **Structure-from-motion** (SfM) problem (or *block orientation*, in Photogrammetry) consists in recovering both scene structure, i.e. 3D scene points, and camera motion, i.e. absolute positions and attitudes of the cameras. SfM methods can be divided into three categories: *structure-first*, *structure-and-motion*, and *motion-first*. Structure-first approaches (e.g. independent models block adjustment [7]), first build stereo-models and then co-register them, similarly to the 3D registration problem. Structure-*and*-motion techniques (e.g. bundle block adjustment [26], resection-intersection methods [22], hierarchical methods [10]) – which are the most common – solve simultaneously for “structure” and “motion”. Finally, motion-first methods [17, 8, 19, 1, 18] first recover the “motion” and then compute the “structure”. These motion-first methods are *global*, as they take into account simultaneously the entire epipolar graph, whose vertices represent the cameras and edges link images having consistent matching points. Most of them solve the SfM optimization problem in two steps. In the first step, the absolute rotation of each image is computed, and in the second step camera translations are recovered: we are concerned here with the first step.

Several approaches for the ARE problem have been proposed. Sharp [21] distributes the error along all cycles in a cycle basis while [9] casts the problem as the optimization of an objective function where rotations are parameterized as quaternions. Martinec in [17] computes an approximate solution to the ARE problem based on Frobenius minimization, and this approach is extended in [1] using spectral de-

composition or semi-definite programming. A gradient descent method based on matrix completion is presented in [2]. The main drawback of such global techniques is that they suffer the presence of inconsistent/outlier relative rotations, and thus they need a preliminary step to detect and remove such outliers before computing the absolute rotations. A wide overview of methods aimed to the identification of outlier epipolar geometries can be found in [18]. These approaches [8, 18, 19] check for cycle consistency, i.e. the deviation from identity, within the epipolar graph. These strategies are computationally demanding and speed is always traded-off with accuracy (in terms of outliers classification). In particular, RANSAC-based approaches (e.g. [19]) suffer from the limitation of increased computational complexity for large-scale datasets. This is confirmed by results in [19], which show that outlier removal is the most expensive step (after feature extraction and matching) within the entire SfM pipeline.

Recently, a few approaches have been developed to robustly solve the ARE problem without detecting outlier rotations explicitly. Techniques in [11, 6], together with the approach presented in this paper, come under this category. In [11] a cost function based on the ℓ_1 norm is used to average relative rotations, exploiting the fact that the ℓ_1 norm is more robust to outliers than the ℓ_2 norm. More precisely, each absolute rotation is updated in turn by applying the Weiszfeld algorithm to its neighbors. The method proposed in [6] works on the Lie group structure of 3D rotations, and combines an ℓ_1 solution with an iteratively reweighted least squares approach.

A related approach [25] employs rank minimization for simultaneous alignment of range images, without computing matching points.

Our solution to the ARE problem is inspired by recent advances in the research fields of *low-rank & sparse matrix decomposition* and *matrix completion*. The main contribution of this paper is the formulation of a novel cost function that naturally includes the outliers in its definition and a minimization scheme that leverages on the GODEC algorithm [28]. This results in a robust method which copes with outliers and missing data simultaneously. These two entangled problems of matrix completion *and* low-rank recovery in the presence of outliers have been addressed only by a few, recent works [29].

The paper is organized as follows. Section 2 gives an overview of the theoretical background required to define our algorithm, i.e. low-rank & sparse matrix decomposition and matrix completion. Section 3 provides a detailed description of our robust solution to the ARE problem. The method proposed in this section is supported by experimental results on both synthetic and real data, shown in Section 4. The conclusions along with possible further developments are presented in Section 5.

2. Low-rank & sparse matrix decomposition

We will show in the next section that the rotation averaging problem can be reduced to a problem of recovering the entries of a matrix X which is known to admit an approximate decomposition as the sum of a low-rank term L and a sparse term S starting from an incomplete set of measurements of its entries.

The goal of low-rank & sparse matrix decomposition is to find an approximate decomposition of a data matrix X into a low-rank matrix L and a sparse matrix S such that $X = L + S + N$, with N an additive noise. Generally the sparse term S represents gross errors affecting the measurements, while the low-rank part represents some meaningful low-dimensional structure contained into the data. An example of such decomposition techniques is Robust Principal Component Analysis (RPCA) [4] which computes a blind separation of low-rank data and sparse errors by solving the minimization problem

$$\begin{cases} \min_{L,S} \|L\|_* + \lambda \|S\|_1 \\ \text{s.t. } \|X - L - S\|_F \leq \epsilon \end{cases} \quad (1)$$

where $\|\cdot\|_*$ denotes the nuclear norm and $\|\cdot\|_F$ the Frobenius norm. A faster alternative to RPCA is represented by the GODEC algorithm described in [28]. This method requires to know approximately both the rank r of the low-rank term L and the cardinality k of the sparse term S , and solves the following minimization problem

$$\begin{cases} \min_{L,S} \|X - L - S\|_F^2 \\ \text{s.t. } \text{rank}(L) \leq r, \text{ card}(S) \leq k. \end{cases} \quad (2)$$

GODEC alternatively forces L to the rank- r approximation of $X - S$, and forces S to the sparse approximation with cardinality k of $X - L$. The rank- r projection is computed using *Bilateral Random Projections* (BRP) instead of Singular Value Decomposition (SVD) in order to speed up the computation. The updating of S is obtained via entry-wise hard thresholding, keeping the first k largest elements of $X - L$ only. It can be shown that the value of the cost function monotonically decreases and converges to a local minimum, while L and S linearly converge to local optima. The method is described in Algorithm 1. More details can be found in [28].

Low-rank & sparse matrix decomposition methods generally assume that the data matrix X is fully available; however, in practical scenarios, one has to face the problem of missing data. Matrix completion theory [5] is the most natural instrument to manage low-rank matrices containing missing entries. Actually, the goal of matrix completion techniques is to complete a low-rank data matrix X starting from a random subset of its entries $\mathcal{P}_\Omega(X)$ eventually corrupted with a small amount of noise. Here Ω denotes the

Algorithm 1 GODEC

Input: X, r, k, ϵ **Output:** L, S **Initialize:** $L_0 = X, S_0 = 0, t = 0$ **while** $\|X - L_t - S_t\|_F^2 / \|X\|_F^2 > \epsilon$ **do**

1. $t = t + 1$
2. Assign the rank- r approximation of $X - S_{t-1}$ to L_t using BRP
3. Assign the projection onto K of $X - L_t$ to S_t , where K is the nonzero subset of the first k largest entries of $|X - L_t|$

end whileReturn $L = L_t, S = S_t$

sampling index matrix of X , i.e. $\Omega_{ij} = 1$ if X_{ij} is available, $\Omega_{ij} = 0$ otherwise, and $\mathcal{P}_\Omega(\cdot)$ denotes the projection onto Ω . The matrix completion minimization problem can be formulated as

$$\begin{cases} \min_L \|\mathcal{P}_\Omega(X) - \mathcal{P}_\Omega(L)\|_F^2 \\ \text{s.t. rank}(L) \leq r. \end{cases} \quad (3)$$

Conventional solvers for this problem include Grassmann manifold solvers such as OPTSPACE [14], and convex solvers such as Augmented Lagrangian Multiplier [16]. The matrix completion problem can also be solved by slightly modifying the GODEC Algorithm, as explained in [28]. The minimization problem (3) is reformulated by introducing a sparse term S which approximates $-\mathcal{P}_{\Omega^C}(L)$, where Ω^C represents the complementary of Ω

$$\begin{cases} \min_{L,S} \|X - L - S\|_F^2 \\ \text{s.t. rank}(L) \leq r, \text{supp}(S) = \Omega^C. \end{cases} \quad (4)$$

Note that here S does not represent the outliers, but the completion of missing entries. In the GODEC algorithm for matrix completion, the updating of the sparse term is obtained by assigning $\mathcal{P}_{\Omega^C}(X - L) = -\mathcal{P}_{\Omega^C}(L)$ to S . Note that the two versions of the GODEC algorithm, i.e. Algorithm 1 and its modification for matrix completion, are rather orthogonal. On one hand Algorithm 1 handles the presence of outliers but it does not deal with missing data, on the other hand the matrix completion modification can fill missing entries, but it is not robust to outliers.

Here we propose a novel version of the GODEC algorithm (called R-GODEC, where ‘‘R’’ stands for ‘‘robust’’) which manages at the same time both the presence of outliers and missing entries in the data matrix X . To the best of our knowledge, we are among the first to address the problem at the intersection between low-rank & sparse matrix decomposition and matrix completion. A seminal work is

presented in [27], where authors combine a greedy pursuit for the updating of the sparse term, with an SVD-based approximation for the low-rank term. Similarly to GODEC, this method requires to know in advance the cardinality of the sparse term. Instead of following this approach, we leverage on the GODEC algorithm for two main reasons. First of all, using a BRP-based instead of an SVD-based approximation for the low-rank term, the GODEC computational efficiency is inherited; secondly, minimizing a cost function which is not explicitly dependent from the projection onto Ω , it results in a very flexible structure which can be easily modified in order to automatically manage the sparse term cardinality. The next section is devoted to introduce our solution.

3. Our method

The ARE problem consists in recovering the absolute rotations R_1, \dots, R_N in order to satisfy the compatibility constraints $R_{ij} = R_i R_j^T$, where R_{ij} denotes the relative rotation of the pair (j, i) . Suppose that estimates \hat{R}_{ij} of the theoretical relative rotations are available for some index pairs (i, j) in a set $\mathcal{N} \subset \{1, \dots, N\} \times \{1, \dots, N\}$ (hereafter we denote estimates inferred from the input data with the hat accent). These relative rotations will in general not be compatible, thus the goal is to find the absolute rotations such that $\hat{R}_{ij} \approx R_i R_j^T$, resulting in the following minimization problem

$$\min_{R_i \in SO(3)} \sum_{(i,j) \in \mathcal{N}} \left\| \hat{R}_{ij} - R_i R_j^T \right\|_F^2. \quad (5)$$

In order to cast the ARE problem in terms of ‘‘low-rank & sparse’’ matrix decomposition, we observe that problem (5) can be reformulated in a useful equivalent form in the following way. We first consider the case where estimates \hat{R}_{ij} of the theoretical relative rotations R_{ij} are available for all $i, j = 1, \dots, N$ (we will deal with the case of missing relative rotations in Section 3.2). Let R be the $3N \times 3$ block-matrix containing the absolute rotations and let X be the $3N \times 3N$ block-matrix containing the pairwise rotations:

$$R = \begin{bmatrix} R_1 \\ R_2 \\ \dots \\ R_N \end{bmatrix}, \quad X = \begin{pmatrix} I & R_{12} & \dots & R_{1N} \\ R_{21} & I & \dots & R_{2N} \\ \dots & \dots & \dots & \dots \\ R_{N1} & R_{N2} & \dots & I \end{pmatrix}. \quad (6)$$

It is shown in [1] that X admits the decomposition $X = RR^T$ and hence it is symmetric, positive semidefinite and of rank 3. Thus the optimization problem (5) becomes equivalent to minimize the error between the observed \hat{X} and the underlying ground truth X , i.e. (5) coincides with

$$\begin{cases} \min_X \left\| \hat{X} - X \right\|_F^2 \\ \text{s.t. } X = RR^T, R_i \in SO(3). \end{cases} \quad (7)$$

This formulation implicitly assumes that the data matrix \widehat{X} satisfies the properties mentioned above except from an additive noise N , that is

$$\widehat{X} = X + N. \quad (8)$$

In other terms, problem (7) aims at minimizing the noise N . As observed in [12], this is a complex multi-variable non-convex optimization problem, thus a reasonable approach is to relax some constraints over the variable X to make the computation tractable. Two examples of relaxations, i.e. *spectral* and *semidefinite programming* relaxations, are described in [1]. The former forces the entire columns of R to be orthonormal, instead of imposing the orthonormality constraints on each 3×3 block R_i . The latter guarantees that the optimization variable X is symmetric positive semidefinite, and covered by identity blocks along its diagonal. Although these solutions can efficiently average noisy orientations, it is well known that they are highly non-robust and that they can give incorrect results in presence of even a single outlier. Such outliers are very frequent when dealing with real data. In the SfM context, for example, repetitive structures in the images cause mismatches which skew the epipolar geometry. In the global registration of 3D point sets, outliers are caused by faulty pairwise registration, which in turn may be originated by insufficient overlap and/or poor initialization.

Here we introduce a new model for the ARE problem that naturally copes with the presence of outliers. More precisely, we add a new term S in Equation (8), with the property that S is nonzero in correspondence of the inconsistent relative rotations only. Moreover, we consider the *rank* relaxation in which the matrix X is enforced to have rank (at most) 3. This results in the following model

$$\widehat{X} = L + S + N \quad (9)$$

where L is a rank-3 matrix, S is a sparse matrix containing the outliers and N is the noise. We use the notation L instead of X to underline that L will not coincide with X in general, due to the rank relaxation. Equation (9) is the approximated low-rank & sparse matrix decomposition of \widehat{X} . The associated minimization problem is

$$\begin{cases} \min_{L,S} \left\| \widehat{X} - L - S \right\|_F^2 \\ \text{s.t. rank}(L) \leq 3, S \text{ sparse.} \end{cases} \quad (10)$$

It is worth noting that here the outliers are intrinsically included in the cost function. With respect to non robust solutions that rely on a preliminary outlier rejection step, our approach has the great advantage of being intrinsically robust against outliers.

3.1. Decomposition with Soft Thresholding

The approximated low-rank and sparse decomposition of the matrix \widehat{X} can be computed using Algorithm 1. In our case the rank r is known and it is equal to 3, while the value of k is unknown. Rather than estimating the cardinality of the sparse term, which could fall into a thorny outlier rejection problem, we prefer to explore an alternative strategy which does not involve the parameter k . This implies to modify Step 3 in Algorithm 1 by considering the following minimization problem instead of (10)

$$\begin{cases} \min_{L,S} \frac{1}{2} \left\| \widehat{X} - L - S \right\|_F^2 + \lambda \|S\|_1 \\ \text{s.t. rank}(L) \leq 3 \end{cases} \quad (11)$$

where λ is a regularization parameter and $\|S\|_1$ denotes the ℓ_1 -norm of S viewed as a vector. It is well known from sparse representation theory that minimizing the ℓ_1 -norm will in general yield a sparse vector, and hence the solution of the above problem is expected to recover the sparse pattern of the outlier rotations. In this case, the updating of the sparse part is obtained by minimizing the cost function in (11) with respect to S , keeping L constant. Such a problem is known to have an analytical solution called *soft thresholding* or *shrinkage* [3], expressed as

$$S_\lambda(\widehat{X} - L) = \text{sign}(\widehat{X} - L) \cdot \max(0, |\widehat{X} - L| - \lambda) \quad (12)$$

where scalar operations are applied element-wise.

3.2. Dealing with missing data

We now consider the case of missing relative rotations, which frequently occurs in practice, because many local reference frames are not directly related one to each other by a relative rotation. In other words, referring to the SfM terminology, the epipolar graph is not complete.

To handle this situation, the data matrix \widehat{X} is modified by introducing zero blocks in correspondence of the missing pairwise rotations. When noise with a small variance is added to the given data, the matrix \widehat{X} can be completed by exploiting conventional matrix completion methods, as shown in [2], or by modifying the GODEC Algorithm, as explained in Section 2. However, in the presence of outliers, these methods offer no guarantees on a correct recovery of the low-rank matrix.

We propose to fill this gap by extending the model in (9) in order to cope with both outliers and missing data. More in detail, we express the sparse term in (9) as the sum of two terms S_1 and S_2 resulting in the following model

$$\widehat{X} = L + S_1 + S_2 + N. \quad (13)$$

S_1 is a sparse matrix over the sampling set Ω representing the outliers in the measurements. S_2 has support on Ω^C and

it is an approximation of $-\mathcal{P}_{\Omega^C}(L)$, representing the completion of the missing entries. In other words, equation (13) reduces to (9) over the sampling set, since S_2 is zero in Ω ; on the contrary, out of the sampling set, equation (13) turns to $L + S_2 + N = 0$, since both S_1 and \hat{X} are zero in Ω^C , and thus S_2 can be updated according to the modified GODEC algorithm for matrix completion, as explained in Section 2. Starting from (13) we solve the following problem

$$\begin{cases} \min_{L, S_1, S_2} \left\| \hat{X} - L - S_1 - S_2 \right\|_F^2 \\ \text{s.t. rank}(L) \leq 3, \\ \text{supp}(S_1) \subseteq \Omega, S_1 \text{ sparse over } \Omega, \\ \text{supp}(S_2) = \Omega^C \end{cases} \quad (14)$$

by modifying Algorithm 1 according to a block-coordinate minimization scheme. The sparse terms S_1, S_2 are considered separately: S_1 is updated by applying soft thresholding to the matrix $\mathcal{P}_{\Omega}(\hat{X} - L)$, while $-\mathcal{P}_{\Omega^C}(L)$ is assigned to S_2 . This method, called R-GODEC, is summarized in Algorithm 2. As a matter of fact, a formal proof of convergence of the algorithm is out of the scope of this paper. However, the fact that each step does not increase the objective function is a property shared with any block-relaxation technique (modulo the approximation induced by BRP). Once the optimal L is found, we proceed as follows to compute the absolute rotations. Since the solution is defined up to a global rotation, corresponding to a change in the orientation of the world coordinate frame, any block-column of L can be used as an estimate of R . Due to the rank relaxation, each 3×3 block is not guaranteed to belong to $SO(3)$, thus we find the nearest rotation matrix (in the Frobenius norm sense) by using SVD [13].

Algorithm 2 R-GODEC

Input: $\hat{X}, r, \epsilon, \lambda$
Output: L, S_1, S_2
Initialize: $L_0 = \hat{X}, S_1^0 = 0, S_2^0 = 0, t = 0$
while $\left\| \hat{X} - L_t - S_1^t - S_2^t \right\|_F^2 / \left\| \hat{X} \right\|_F^2 > \epsilon$ **do**
 1. $t = t + 1$
 2. Assign the rank- r projection of $\hat{X} - S_1^{t-1} - S_2^{t-1}$ to L_t using BRP
 3. Assign $\mathcal{S}_{\lambda}(\mathcal{P}_{\Omega}(\hat{X} - L_t))$ to S_1^t
 4. Assign $-\mathcal{P}_{\Omega^C}(L_t)$ to S_2^t
end while
Return $L = L_t, S_1 = S_1^t, S_2 = S_2^t$

3.3. Outlier Detection

Although the absolute rotations computed by our algorithm are intrinsically insensitive to outliers, it might be

beneficial for the subsequent steps (e.g., computing translations in SfM) to single out bad relative rotations from the data matrix \hat{X} , which are indicators that the whole rigid transformation is probably faulty. As a matter of fact, outliers can be identified by analyzing the optimal S_1 returned by R-GODEC (Algorithm 2), which, however, operates element-wise on \hat{X} . Thus, we will deem the rotation \hat{R}_{ij} as outlier whenever the number of non-zero entries of the 3×3 block in S_1 associated with \hat{R}_{ij} is greater than a threshold θ , with $\theta \in \{1, 2, \dots, 9\}$.

4. Experiments

To assess our method – R-GODEC, we consider both synthetic and real scenarios. All the experiments are carried out in MATLAB on a dual-core 1.3 GHz PC. We compare R-GODEC with the techniques introduced in [1], i.e. *spectral decomposition* (EIG) and *semidefinite programming* (SDP). To implement such techniques, we use the MATLAB command *eigs* for the former and the SeDuMi toolbox [24] for the latter. We also insert in the comparison the matrix completion algorithm OPTSPACE, whose code is available online [14], and our implementation of the Weiszfeld algorithm [11]. We evaluate the accuracy of rotation recovery by using the *angular distance* between ground truth and estimated absolute rotations. The angular distance between two rotations A and B is defined as $d_{\angle}(A, B) = d_{\angle}(BA^T, I) = 1/\sqrt{2} \left\| \log(BA^T) \right\|_2$. Other distances in $SO(3)$ can be considered with comparable results. As for the validation of outlier detection (which is not strictly part of R-GODEC, though), we consider the customary receiver operating characteristic (ROC) curve, where the parameter is the threshold θ .

4.1. Simulated Data

In this section we show experimental results on synthetic data by analyzing the performances of R-GODEC in the presence of outliers among the relative rotations.

Following the experiments described in [1], we consider $N = 100$ rotation matrices representing the ground truth absolute orientations, and we perturb the relative rotations with Gaussian noise with SNR of 30dB, corresponding to a mean angular error of 2 degrees. These noisy matrices are then projected onto $SO(3)$. We consider a realistic case in which a fraction p of the relative rotations is missing (drawn randomly with the constraint that the resulting epipolar graph remains connected). The rest of the pairwise orientations are either true rotations corrupted by noise or drawn uniformly from $SO(3)$, simulating outliers. We analyze the cases $p = 0$ (no missing pairwise rotations), $p = 0.5$, $p = 0.7$ and $p = 0.9$. In the first case we also insert in the comparison the original GODEC (Algorithm 1), which only works with the full data matrix and assumes that the parameter k is given. The value of λ is set equal to

0.05 in the cases $p = 0$, $p = 0.5$, equal to 0.1 in the case $p = 0.7$, and equal to 0.15 in the case $p = 0.9$. Indeed, if the data matrix is highly incomplete, it is preferable to choose a higher value for λ , in order to give more importance to the outlier term S_1 rather than the completion term S_2 .

Figure 1 shows the results, averaged over 30 trials. R-GODEC performs significantly better than the other analyzed state-of-the-art techniques. In the cases $p = 0, 0.5, 0.7$, when outliers do not exceed inliers (the second last value corresponds to 50% of effective outliers), the error of R-GODEC remains almost constant, showing no sensitivity to outliers. When the percentage of outliers exceeds that of inliers, the error starts to grow, which suggest empirically that R-GODEC might have a 0.5 breakdown point. On the contrary, the Weiszfeld algorithm, which belong to the category of robust methods together with R-GODEC, has a 0.3 breakdown point. Indeed, by using the ℓ^1 -norm in (5) in place of the ℓ^2 -norm, the influence of outliers is reduced but not canceled. Our approach, on the contrary, uses the ℓ^1 -norm as a sparsity promoter. The standard deviation of R-GODEC is also very small, compared with the other methods. In particular, note that in the case $p = 0$ there are no significant differences between R-GODEC and the original GODEC algorithm *with known k*, which demonstrates the effectiveness of the soft thresholding strategy for computing the sparse term. It is worth noting that R-GODEC outperforms all the analyzed techniques also in the challenging situation of a highly incomplete data matrix ($p = 0.9$).

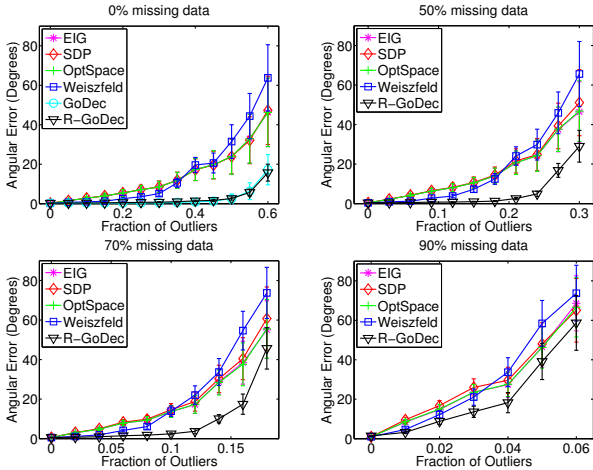


Figure 1. Mean angular errors [deg] and standard deviations of absolute rotations as a function of the fraction of outliers, for different percentages of missing relative rotations. The fraction of outliers in the abscissae is always referred to the total number of rotations, including the missing ones, so as the right extremum corresponds to 60% of effective outliers.

In this experiment we also evaluate the outlier detection performances of our method. With reference to Figure 2, up

to 50% of missing data (first row of the figure), our outlier detector gives a perfect classification, as confirmed by the area under the ROC curve which is equal to 1. This does not hold for the black curves marked with triangles, which refer to the extreme case of 60% of outliers, which is beyond the breakdown point of R-GODEC. However, also in this case a fairly good classification can be appreciated. In the case of a highly incomplete data matrix (second row of Figure 2), although our detector becomes less accurate as the fraction of outliers increases, it maintains quite good performances.

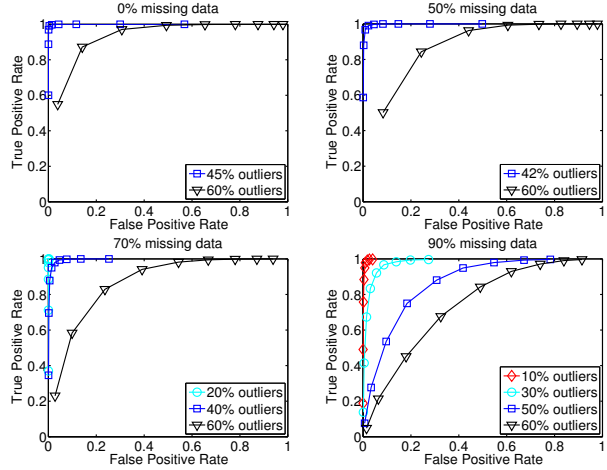


Figure 2. Outlier detection: ROC curves of the classification, for different percentages of missing relative rotations. The fraction of outliers is always referred to the number of available relative rotations.

We conclude this analysis by discussing the performances of R-GODEC in terms of computational time. Figure 3 reports the running time of several algorithms, including R-GODEC, for different values of N . The execution cost includes both operations on the data matrix and the subsequent projection onto $SO(3)$. Our method is comparable to spectral decomposition and OPTSPACE algorithms, which however are not robust, and significantly faster than semidefinite programming and Weiszfeld algorithms. Considering that EIG is the fastest solution to the ARE problem known in the literature, one can see how R-GODEC buys robustness at a very little computational cost. The reason why Weiszfeld curve is not so regular in Figure 3 is as follows. As for the initialization, Weiszfeld algorithm propagates the compatibility constraint along a random spanning tree, starting from the node with the maximum number of incident edges. Thus, the number of iterations required to yield convergence is dependent on the accuracy of such initial guess.

4.2. Real data

In this section we apply R-GODEC to the ARE of real cameras. More precisely, to assess our method, we use the

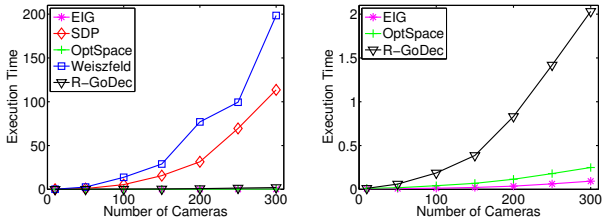


Figure 3. Execution times (in seconds) of ARE as a function of the number of absolute rotations. The parameters defining the stopping criterion (maximum number of iterations and tolerance) are the same for all methods. The right figure is a zoom of the left one.

benchmark in [23] which provides ground-truth rotations. The datasets consist in 8 to 30 images of dimensions 3072×2048 pixels.

To compute the relative rotations we follow a standard SfM pipeline. First, reliable matching points across the input images are computed by using SIFT key-points. Then, each essential matrix is computed through the RANSAC procedure, and it is factorized to obtain a unique \hat{R}_{ij} , which is considered missing if less than m inlier correspondences are found. In our simulations, we consider the cases $m = 100, 250, 500$. As m increases, the percentage of missing data becomes larger and the fraction of outliers decreases. The value of λ is set equal to 0.05.

Results are shown in Table 1. In all the analyzed datasets, R-GODEC outperforms EIG, SDP and OPTSPACE, which are not robust to outliers. Our method gives better results than the Weiszfeld algorithm when contamination of outliers is particularly evident. This is clear in the Castle sequences which contain repetitive structures, such as doors and windows, that generate outliers. When there are no (or few) outliers, R-GODEC and Weiszfeld are comparable, our method is significantly faster, though. In few cases, Weiszfeld yields more accurate results than R-GODEC, probably due to the fact that R-GODEC uses a binary criterion to identify gross errors in the data matrix, causing possible ambiguity between outliers and noise, when outliers are actually absent. Note that the errors in Table 1 are obtained *without* applying bundle adjustment, i.e. the final refinement required in any SfM method. Given the high dimensional optimization involved, successful convergence of bundle adjustment heavily depends on a good initial guess: R-GODEC provides an excellent initial guess to such optimization.

5. Conclusion

In this paper we have presented a robust method to solve the ARE problem. We have formulated a novel cost function for such a problem which naturally includes the outliers in its definition. In particular, we have developed an algorithm which successfully handles both outlier and missing

relative rotations by casting the problem as a “low-rank & sparse” matrix decomposition. R-GODEC, obtained by extending the GODEC algorithm, is efficient and highly accurate, as demonstrated by simulated and real experiments. As a side effect, our solution can be seen as a valid and cost-effective detector of the inconsistent pairwise rotations. As regards possible future work, two directions could be investigated. From one side, one could enforce that the output matrix, beside having a given rank, must also be positive semidefinite, by devising specific techniques for updating the low-rank term. On the other side, one could substitute the ℓ_1 norm in (11) with some mixed norm in order to promote group sparsity and to better highlight the block structure of the data matrix.

References

- [1] M. Arie-Nachimson, S. Z. Kovalsky, I. Kemelmacher-Shlizerman, A. Singer, and R. Basri. Global motion estimation from point matches. *International Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission*, 2012.
- [2] F. Arrigoni, B. Rossi, F. Malapelle, P. Fragneto, and A. Fusiello. Robust global motion estimation with matrix completion. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XL-5:63–70, 2014.
- [3] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1), 2009.
- [4] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM*, 58(3), 2011.
- [5] E. J. Candès and T. Tao. The power of convex relaxation: near-optimal matrix completion. *Transactions on Information Theory*, 56(5), 2010.
- [6] A. Chatterjee and V. M. Govindu. Efficient and robust large-scale rotation averaging. In *International Conference on Computer Vision*, 2013.
- [7] F. Crosilla and A. Beinat. Use of generalised procrustes analysis for the photogrammetric block adjustment by independent models. *ISPRS Journal of Photogrammetry & Remote Sensing*, 56(3), 2002.
- [8] O. Enqvist, F. Kahl, and C. Olsson. Non-sequential structure from motion. In *Eleventh Workshop on Omnidirectional Vision, Camera Networks and Non-classical Camera*, 2011.
- [9] A. Fusiello, U. Castellani, L. Ronchetti, and V. Murino. Model acquisition by registration of multiple acoustic range views. In *European Conference on Computer Vision*. 2002.
- [10] R. Gherardi, M. Farenzena, and A. Fusiello. Improving the efficiency of hierarchical structure-and-motion. In *Conference on Computer Vision and Pattern Recognition*, 2010.
- [11] R. Hartley, K. Aftab, and J. Trunpf. L1 rotation averaging using the Weiszfeld algorithm. *Conference on Computer Vision and Pattern Recognition*, 2011.
- [12] R. I. Hartley, J. Trunpf, Y. Dai, and H. Li. Rotation averaging. *International Journal of Computer Vision*, 2013.

Table 1. Mean and median angular errors [deg] on the absolute rotations, for different values of the threshold m used to define missing pairwise rotations. Top: $m = 100$. Middle: $m = 250$. Bottom: $m = 500$. The percentage of outlier rotations is computed with reference to the complete epipolar graph. In the case $m = 100$ some sequences are missing since the percentage of outliers is greater than 50%.

	miss. %	out. %	EIG		SDP		OPTSPACE		Weiszfeld		R-GoDEC	
			mean	median	mean	median	mean	median	mean	median	mean	median
Entry-P10	27	4	1.65	1.48	1.65	1.49	1.65	1.48	0.7	0.61	1.12	1.08
Fountain-P11	0	15	4.04	1.68	4.36	1.75	4.04	1.69	2.95	0.74	1.71	1.53
Herz-Jesu-P8	0	18	10.92	10.44	11.81	11.32	10.92	10.44	6.99	6.39	3.66	3.85

	miss. %	out. %	EIG		SDP		OPTSPACE		Weiszfeld		R-GoDEC	
			mean	median	mean	median	mean	median	mean	median	mean	median
Castle-P30	43	13	8.77	3.12	8.75	3.21	8.74	3.22	6.45	1.76	2.47	1.14
Castle-P19	40	16	12.44	11.01	13.72	15.15	13.26	13.92	8.04	8.81	3.22	3.92
Entry-P10	27	4	1.65	1.48	1.65	1.49	1.65	1.48	0.7	0.61	1.12	1.08
Fountain-P11	13	4	1.09	0.69	1.09	0.69	1.09	0.69	0.73	0.59	0.72	0.65
Herz-Jesu-P25	36	5	6.67	4.67	6.74	4.76	6.68	4.68	4.1	2.3	2.2	1.9
Herz-Jesu-P8	14	4	5.36	5.7	5.39	5.7	5.36	5.70	2.3	2.5	1.86	1.98

	miss. %	out. %	EIG		SDP		OPTSPACE		Weiszfeld		R-GoDEC	
			mean	median	mean	median	mean	median	mean	median	mean	median
Castle-P30	52	5	7.05	2.37	7.97	2.98	7.46	2.74	4.52	1.35	1.67	0.9
Castle-P19	50	7	6.19	3.77	6.29	6.11	6.02	3.79	3.89	1.73	2.73	1.86
Entry-P10	27	4	1.65	1.48	1.65	1.49	1.65	1.48	0.7	0.61	1.12	1.08
Fountain-P11	16	2	0.97	0.73	0.97	0.72	0.97	0.73	0.74	0.62	0.77	0.63
Herz-Jesu-P25	41	2	2.49	2.37	2.5	2.38	2.5	2.38	1.84	1.99	1.46	1.5
Herz-Jesu-P8	18	0	0.82	0.8	0.82	0.8	0.82	0.8	0.54	0.44	0.74	0.63

- [13] J. Keller. Closest unitary, orthogonal and Hermitian operators to a given operator. *Mathematics Magazine*, 48, 1975.
- [14] R. H. Keshavan, A. Montanari, and S. Oh. OPTSPACE: a matrix completion algorithm. <http://web.engr.illinois.edu/~swoh/software/optspace/code.html>.
- [15] S. Krishnan, P. Y. Lee, J. B. Moore, and S. Venkatasubramanian. Optimisation-on-a-manifold for global registration of multiple 3D point sets. *International Journal of Intelligent Systems Technologies and Applications*, 3(3/4), 2007.
- [16] Z. Lin, M. Chen, and Y. Ma. The augmented lagrange multiplier method for exact recovery of corrupted Low-Rank matrices. eprint arXiv:1009.5055.
- [17] D. Martinec and T. Pajdla. Robust rotation and translation estimation in multiview reconstruction. In *Conference on Computer Vision and Pattern Recognition*, 2007.
- [18] P. Moulon, P. Monasse, and R. Marlet. Global Fusion of Relative Motions for Robust, Accurate and Scalable Structure from Motion. In *International Conference on Computer Vision*, December 2013.
- [19] C. Olsson and O. Enqvist. Stable structure from motion for unordered image collections. In *Proceedings of the 17th Scandinavian conference on Image analysis (SCIA11)*, 2011.
- [20] X. Pennec. Multiple registration and mean rigid shape: Applications to the 3D case. In *16th Leeds Annual Statistical Workshop*, 1996.
- [21] G. C. Sharp, S. W. Lee, and D. K. Wehe. Multiview registration of 3D scenes by minimizing error between coordinate frames. In *European Conference on Computer Vision*, 2002.
- [22] N. Snavely, S. M. Seitz, and R. Szeliski. Photo tourism: Exploring photo collections in 3D. *ACM Transactions on Graphics*, 25(3), 2006.
- [23] C. Strecha, W. von Hansen, L. J. V. Gool, P. Fua, and U. Thoennessen. On benchmarking camera calibration and multi-view stereo for high resolution imagery. In *Conference on Computer Vision and Pattern Recognition*, 2008.
- [24] J. F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optimization Methods and Software*, 1999.
- [25] D. Thomas, Y. Matsushita, and A. Sugimoto. Robust simultaneous 3D registration via rank minimization. *International Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission*, 2012.
- [26] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon. Bundle adjustment - a modern synthesis. In *Proceedings of the International Workshop on Vision Algorithms*, 2000.
- [27] A. E. Waters, A. C. Sankaranarayanan, and R. G. Baraniuk. Sparcs: Recovering low-rank and sparse matrices from compressive measurements. In *Neural Information Processing Systems (NIPS)*, 2011.
- [28] T. Zhou and D. Tao. GoDec: Randomized low-rank & sparse matrix decomposition in noisy case. In *International Conference on Machine Learning*, 2011.
- [29] X. Zhou, C. Yang, H. Zhao, and W. Yu. Low-rank modeling and its applications in image analysis. *CoRR*, 2014.