Global Registration of 3D Point Sets
via LRS decomposition

Federica Arrigoni[1], Beatrice Rossi[2], and Andrea Fusiello[1]

[1] DPIA - Università di Udine   Via Delle Scienze, 208 - Udine, Italy
arrigoni.federica@spes.uniud.it, andrea.fusiello@uniud.it
[2] AST Lab - STMicroelectronics   Via Olivetti, 2 - Agrate Brianza, Italy
beatrice.rossi@st.com

Abstract. This paper casts the global registration of multiple 3D point-sets into a low-rank and sparse decomposition problem. This neat mathematical formulation caters for missing data, outliers and noise, and it benefits from a wealth of available decomposition algorithms that can be plugged-in. Experimental results show that this approach compares favourably to the state of the art in terms of precision and speed, and it outperforms all the analysed techniques as for robustness to outliers.

Keywords: point-set registration, motion synchronization, matrix completion, low-rank and sparse matrix decomposition

1   Introduction

The goal of multiple point-set registration is to find the rigid transformations that bring multiple ($n \geq 2$) 3D point sets into alignment, where each rigid transformation is represented by an element of the Special Euclidean Group SE(3), namely the semi-direct product of the Special Orthogonal Group SO(3) with $\mathbb{R}^3$. This is a fundamental problem in the reconstruction of 3D models of objects, covering a wide range of applications, including (but not limited to) cultural heritage, engineering modelling and virtual reality.

If $n = 2$ then we are dealing with a *pairwise* (two point-sets) registration problem. The gold standard in this context is the Iterative Closest Point (ICP) Algorithm [1, 2], which computes correspondences between the point sets given an estimate for the rigid transformation, then updates the transformation based on the current correspondences, and iterates through these steps until convergence – to a local minimum – is reached. See [3] for an overview of several variants of the ICP Algorithm.

If $n > 2$ then we are dealing with a *multiple* point-set registration problem, which is more complex than the $n = 2$ case due to the high amount of parameters that have to be estimated. Among the initial attempts to address this problem are the *sequential* techniques introduced in [2, 4], that repeatedly register a new point set into a growing model, until all the sets are considered. This approach however returns suboptimal solutions since it does not take into account all the available constraints, *e.g.* the constraint between the last and first point set is

not used if the sets are obtained using a turntable. *Global* methods, on the other hand, consider simultaneously all the points sets. They are able to exploit the redundancy in the constraints between pairs of point sets, to compensate and distribute the error, thereby preventing drift in the solution. Global registration can be solved in *point space* or in *frame space*. In the former case, all the rigid transformations are computed by optimizing a cost function that depends on the distance between corresponding points. In the latter case, the optimization criterion is related to the internal consistency of the network of rigid transformations applied to the local coordinate frames. This instance is also known as *motion synchronization* [5] or *motion averaging* [6].

Early point-space solutions include the methods presented in [7–10]. More recently, [11] solves the problem on the manifold of rotations with Gauss-Newton iterations and then computes translations through least-squares. Such approach is improved in [12] by reducing computational time and it is embedded in a Bayesian framework in [13], in order to take into account reliability of correspondences. A similar formulation is adopted in [14] where the authors cast the registration problem to a semidefinite program, proving conditions for exact and stable recovery of rigid transformations. In [15] a robust solution is derived by minimizing a cost function based on the ℓ_1-norm. A generalization of the ICP Algorithm to $n > 2$ is described in [16], which builds on the Levenberg-Marquardt ICP formulation of [17], while in [18] multiple point-set registration is solved by combining ICP and Generalized Procrustes Analysis [19]. A related approach employs rank minimization for global registration of multiple depth images [20].

Frame-space methods originate from the pioneering works of [21], that distributes the error along all cycles in a cycle basis, and [22], that casts the problem as the optimization of an objective function in SE(3) where rotations are parameterized as unit quaternions. The authors of [23] represent rigid transformations as dual quaternions, and propose a graph diffusion algorithm where each transformation is updated in turn through linear or geodesic averaging. In [6, 24] the Lie-group structure of SE(3) is exploited and an iterative scheme is proposed in which at each step the rigid transformations are updated by averaging two-view transformations in the tangent space. In [25, 5] motion synchronization is formulated as a null-space problem.

At the border between frame-space and point-space methods is the formulation in [26], where 3D points are used to compute a second-order approximation of the cost function, but they are not involved in subsequent computations.

Among the aforementioned methods, [11–14, 21] first recover the rotation component of the rigid transformations and then compute translations, while [15, 16, 18, 20, 22, 23, 6, 25, 26] and our approach compute rotations and translations simultaneously, as elements of SE(3).

In this paper we concentrate on frame-space methods, for they are faster and less memory-demanding than point-space ones. It goes without saying that any optimal formulation *must* include points in the cost function, in analogy to bundle adjustment in the context of structure from motion. Nevertheless, frame-space approaches yield a fairly accurate registration.

Many frame-space methods are deceived by gross errors (outliers), caused by failure of ICP to estimate the correct transformation between two point-sets. To overcome this drawback, we propose a global frame-space approach to multiple point-set registration which is *robust* to outliers.

We show that the registration problem can be cast to a *low-rank and sparse* (LRS) matrix decomposition. Our proposal is a general framework, not a specific method, since – in principle – any LRS decomposition algorithm can be plugged-in. However, in order to make our approach concrete, we analyse three LRS algorithms, namely R-GoDec [27], Grasta [28] and L1-Alm [29], showing that they can be profitably applied to address the registration problem. Experimental results show that our approach compares favourably to the state of the art as for precision and speed, and it outperforms all the analysed techniques in terms of robustness to outliers. Failure cases appear when the percentage of missing data is extremely high.

This paper builds on [27] and extends it in two respects. From the theoretical point of view, our approach generalizes to SE(3) the original formulation for SO(3). From the experimental point of view, we tackle the problem of multiple point-set registration, while [27] concentrates on structure from motion.

The next section provides the background on LRS decomposition, Sec. 3 defines the multiple point-set registration problem, whereas Sec. 4 describes how such a problem can be translated into a LRS decomposition of an incomplete matrix, corrupted by noise and outliers. Experiments on synthetic and real dataset are reported in Sec. 5, and conclusions are drawn in Sec. 6.

2   Low-rank and sparse decomposition

Low-rank and sparse (LRS) matrix decompositions have become interesting to researchers due to their profitable application in different areas, such as image analysis, pattern recognition, and graph clustering. A survey of such decompositions and a wide overview of available algorithms can be found in [30]. LRS decompositions work by imposing constraints on the rank and sparsity of the addends and have the following general form

$$\mathcal{F}(\widehat{X}) = \mathcal{F}(L) + S + N \tag{1}$$

where $\widehat{X}$ is a known data matrix, L is an unknown low-rank matrix representing some meaningful structure contained into the data, S is an unknown sparse matrix representing outliers, N accounts for a diffuse noise and $\mathcal{F}$ is a linear operator. The goal is to recover L (and possibly S) under conditions on S, N and $\mathcal{F}$ to be further specified.

Examples of LRS decompositions are Robust Principal Component Analysis (RPCA) [31] and Matrix Completion (MC) [32, 33]. RPCA looks for the lowest-rank matrix L and the sparsest matrix S such that a given data matrix $\widehat{X}$ can be decomposed as $\widehat{X} = L + S + N$. Note that this is an instance of Equation (1) with $\mathcal{F}$ being the identity operator. MC is concerned with the problem of

recovering missing entries of an *incomplete* low-rank data matrix $\widehat{X}$. Entries of $\widehat{X}$ are specified on a subset of indices Ω (a.k.a. *sampling set*), namely $\Omega_{ij} = 1$ if $\widehat{X}_{ij}$ is specified and $\Omega_{ij} = 0$ otherwise. MC can be cast as an instance of Equation (1) if $S = 0$ and $\mathcal{F}$ is the projection onto the space of matrices that vanish out of Ω, namely $\mathcal{F}(X) = \mathcal{P}_\Omega(X) = X \circ \Omega$, where $\circ$ is the Hadamard (or entry-wise) product.

We briefly recall some of the available approaches to compute LRS decompositions in the general form (1) with $S, N \neq 0$, L of fixed rank r, and $\mathcal{F} = \mathcal{P}_\Omega$.

The R-GoDec algorithm [27] is a modified version of GoDec [34], which was originally conceived to solve RPCA or MC problems separately. R-GoDec expresses the sparse term S as the sum of two terms S_1 and S_2 having complementary supports: S_1 has support on Ω and represents outliers, while S_2 has support on $\mho$ (the complementary of Ω) and it is an approximation of $-\mathcal{P}_\mho(L)$, *i.e.* it represents completion of missing entries. The associated problem is

$$\min_{L,S_1,S_2} \frac{1}{2} \left\| \mathcal{P}_\Omega(\widehat{X}) - L - S_1 - S_2 \right\|_F^2 + \lambda \left\| S_1 \right\|_1 \tag{2}$$

$$\text{s.t. } \text{rank}(L) \leq r, \text{supp}(S_1) \subseteq \Omega, \text{supp}(S_2) = \mho$$

which is solved using a block-coordinate minimization scheme that alternates the update of L (with S_1, S_2 fixed), and S_1, S_2 (with L fixed). L is computed as the the rank-r projection of $\mathcal{P}_\Omega(\widehat{X}) - S_1 - S_2$ using Bilateral Random Projections [34], S_1 is computed via soft-thresholding [35] of $\mathcal{P}_\Omega(\widehat{X} - L)$, and S_2 is updated as $-\mathcal{P}_\mho(L)$.

Another option is to express the LRS decomposition problem in terms of *subspace identification*, as done in [28, 29]. The goal of subspace identification methods is to identify the column space of the unknown low-rank term L. Indeed, since any matrix L of fixed rank r admits a factorization of the form $L = UY^T$ where U and Y have r columns, the problem of recovering L can be translated into recovering U and Y. Specifically, the GRASTA algorithm [28] solves

$$\min_{S,U,Y} \left\| S \right\|_1 \tag{3}$$

$$\text{s.t. } \mathcal{P}_\Omega(\widehat{X}) = \mathcal{P}_\Omega(UY^T) + S$$

with U belonging to the Grasmannian manifold, i.e. the set of all r-dimensional subspaces of a Euclidean space. GRASTA considers one data vector $\hat{x}$ at a time, representing one column of $\widehat{X}$, and alternates between estimating U and a triple of vectors (s, y, w) which represent respectively the sparse corruptions in $\hat{x}$, the weights for the fit to the subspace U, and the dual vector. For computing U, GRASTA uses gradient descent on the Grasmannian with (s, y, w) fixed, while for computing (s, y, w), it uses the Alternating Direction Method of Multipliers (ADMM) [36]. The L1-ALM algorithm presented in [29] exploits a similar approach and solves instead

$$\min_{U,Y} \left\| \mathcal{P}_\Omega(\widehat{X} - UY^T) \right\|_1 + \lambda \left\| Y^T \right\|_* \tag{4}$$

$$\text{s.t. } U^T U = I_r$$

where $\|\cdot\|_*$ denotes the trace-norm of a matrix, *i.e.* the sum of its singular values. The trace-norm regularization term is introduced to improve convergence. The optimization problem is solved via the augmented Lagrange multiplier (ALM) method [37]. At each iteration, the augmented Lagrange function with orthogonal U is minimized via the Gauss-Seidel iteration, then the Lagrange multiplier and the dual parameter are updated.

3   Problem Definition

Let $P = \{\mathbf{p}_k\}_{k=1}^m$ be a set of 3D points representing a given object expressed in an absolute (world) coordinate system. Let $\{P^i\}_{i=1}^n$ denote multiple views of the object taken from different positions and viewing directions, where each 3D point set $P^i = \{\mathbf{p}_k^i\}_{k \in V_i}$ refers to a subset $V_i \subseteq \{1, \ldots, m\}$ of the original m points. Let $M_i \in SE(3)$ denote the 3D displacement between the local reference frame of view i and the world coordinate system

$$M_i = \begin{pmatrix} R_i & \mathbf{t}_i \\ \mathbf{0}^\mathsf{T} & 1 \end{pmatrix} \in SE(3) \tag{5}$$

where $R_i \in SO(3)$ represents the rotation component of the transformation, and $\mathbf{t}_i \in \mathbb{R}^3$ represents the translation component. In this paper $M_i \in SE(3)$ is referred to as the *absolute motion* of view i. Using this notation, the (homogeneous) coordinates of the k-th point can be expressed in the reference frame of view i as $\mathbf{p}_k^i = M_i \mathbf{p}_k$ and the relation between the coordinates of $\mathbf{p}_k$ in references i and j is given by

$$\mathbf{p}_k^i = M_i M_j^{-1} \mathbf{p}_k^j \tag{6}$$

assuming that $k \in V_i \cap V_j$, where the index set $V_i \cap V_j$ defines corresponding points between P_i and P_j.

The goal of multiple point-set registration is to estimate the absolute transformations $M_i \in SE(3)$ starting from the knowledge of the point sets $\{P^i\}_{i=1}^n$. Since P can be recovered from $\mathbf{p}_k^i = M_i \mathbf{p}_k$ by applying the inverse of absolute motions to each point, the absolute motions can be viewed as the transformations that bring multiple point sets into alignment. The index sets $\{V_i\}_{i=1}^n$ are in general unknown, and therefore they have to be computed beforehand or during the registration process.

The registration problem can be profitably formulated in frame space without involving 3D points [21–25]. Let $M_{ij} \in SE(3)$ denote the rigid transformation between the reference frame of view i and that of view j, which is referred to as the *relative motion* of the pair (i, j). It follows from Equation (6) that the following condition holds

$$M_{ij} = M_i M_j^{-1} \tag{7}$$

which means that the registration problem can be reduced to finding the absolute motions $M_i \in SE(3)$ given measurements of their ratios. Such a problem is known in the literature as *motion synchronization* [5] or *motion averaging* [6].

Let $\mathcal{E} \subseteq \{1, \ldots, n\} \times \{1, \ldots, n\}$ denote the set of available pairs, which can be viewed as the set of edges of a finite simple graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $\mathcal{V} = \{1, \ldots, n\}$, referred to as the *measurement graph*. Obviously, it is possible to recover the absolute motions – up to a global transformation – only if such a graph is connected. If $\mathcal{G}$ is a tree then there is no counteraction of the errors in the solution. However, as soon as redundant measures are considered (*i.e.* the graph has at least one cycle), they are exploited by the synchronization process to globally compensate the errors.

4   Proposed Approach

Let X denote the $4n \times 4n$ block-matrix containing the relative motions, let M denote the $4n \times 4$ block-matrix containing the absolute motions, and let $M^{-\flat}$ denote the $4 \times 4n$ block-matrix containing the inverse of absolute motions, namely

$$
X = \begin{pmatrix} I_4 & M_{12} & \ldots & M_{1n} \\ M_{21} & I_4 & \ldots & M_{2n} \\ & \ldots & & \ldots \\ M_{n1} & M_{n2} & \ldots & I_4 \end{pmatrix}, \quad M = \begin{bmatrix} M_1 \\ M_2 \\ \ldots \\ M_n \end{bmatrix}, \quad M^{-\flat} = \begin{bmatrix} M_1^{-1} & M_2^{-1} & \ldots & M_n^{-1} \end{bmatrix}
\tag{8}
$$

where I_4 denotes the 4×4 identity matrix. Using this notation, the compatibility constraint (7) can be expressed in a compact form as $X = MM^{-\flat}$, which implies that $\mathrm{rank}(X) = 4$.

Let $\widehat{M}_{ij} \in SE(3)$ denote an estimate of the true relative motion $M_{ij} \in SE(3)$. In the presence of noise the measured relative motions will not satisfy Equation (7), thus the goal is to average them so as to maximally satisfy $\widehat{M}_{ij} \approx M_i M_j^{-1}$. A possible approach consists in formulating the following optimization problem

$$
\min_{M_i \in SE(3)} \sum_{(i,j) \in \mathcal{E}} \left\| \widehat{M}_{ij} - M_i M_j^{-1} \right\|_F^2
\tag{9}
$$

where the Frobenius norm $|| \cdot ||_F$ defines a left-invariant metric on $SE(3)$.

Let $\widehat{X}$ denote the $4n \times 4n$ block-matrix containing the measured relative motions $\widehat{M}_{ij} \in SE(3)$, and let A be the adjacency matrix of the measurement graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, *i.e.* $A_{ij} = 1$ if $(i, j) \in \mathcal{E}$ and $A_{ij} = 0$ otherwise. In practical scenarios the measurement graph is not complete, since a point set do not overlap with all the others, thus $\widehat{X}$ has missing entries, which are represented as zero blocks. In other words, the available relative information is represented by $\mathcal{P}_\Omega(\widehat{X})$, where the sampling set has a 4×4 block-structure: $\Omega = A \otimes \mathbb{1}_{4 \times 4}$ ($\mathbb{1}_{4 \times 4}$ is a 4×4 matrix filled by ones, and $\otimes$ denotes the Kronecker product). Using this notation, the minimization problem (9) can be expressed as

$$
\min_X \left\| \mathcal{P}_\Omega(\widehat{X} - X) \right\|_F^2
\tag{10}
$$
$$
\text{s.t. } X = MM^{-\flat}, \; M \in SE(3)^n.
$$

This formulation can successfully average noisy relative motions, but it is not resistant to outliers. For this reason we consider the following problem

$$\min_{X,S} \ \left\| \mathcal{P}_\Omega(\widehat{X} - X) - S \right\|_F^2 \tag{11}$$
$$\text{s.t. } X = MM^{-\flat}, \ M \in SE(3)^n, \ S \text{ is sparse in } \Omega$$

where the additional variable S represents outliers, which are sparse over the measurement graph (by assumption).

If the *rank relaxation* is adopted, *i.e.* all the constraints except of the rank property are ignored, then the following relaxed optimization problem is obtained

$$\min_{L,S} \ \left\| \mathcal{P}_\Omega(\widehat{X} - L) - S \right\|_F^2 \tag{12}$$
$$\text{s.t. } \text{rank}(L) \leq 4, \ S \text{ is sparse in } \Omega$$

where L denotes a low-rank matrix which approximates the theoretical X defined in (8). Note that if the optimal solution to (12) satisfies $L = MM^{-\flat}$ with $M \in SE(3)^n$, then it is the global minimizer of (11). Otherwise, the optimal L only provides an estimate for the theoretical X, since it solves a relaxed version of Problem (11). In particular, the 4×4 blocks of L are not constrained to be Euclidean motions, thus they need to be projected onto $SE(3)$. Specifically, every fourth row is set equal to $[0\ 0\ 0\ 1]$ and 3×3 rotation blocks are projected onto $SO(3)$ through Singular Value Decomposition. Any block-column of the resulting matrix can be taken as an estimate of M, as we already know that the solution is up to a global rigid transformation.

Note that Problem (12) is indeed a LRS decomposition problem, since it is associated to the general formulation (1) with $\mathcal{F} = \mathcal{P}_\Omega$, namely $\mathcal{P}_\Omega(\widehat{X}) = \mathcal{P}_\Omega(L) + S + N$, thus the absolute motions can be recovered by means of any algorithm that computes such decomposition.

5   Experiments

We evaluated our approach on both simulated and real datasets analysing resilience to noise, robustness to outliers, sensitivity to missing data and computational cost. We considered three LRS decomposition algorithms, namely R-GoDec [27], Grasta [28] and L1-Alm [29], showing that they can be successfully applied to perform multiple point-set registration in a robust and efficient manner. We compared such algorithms to other registration techniques which work in frame space (*i.e.* not requiring point correspondences), namely the methods developed by Sharp *et al.* [21], Torsello *et al.* (Diffusion) [23], Bernard *et al.* (Null-space) [25] and Govindu [6]. The codes of Grasta, L1-Alm and Diffusion are available online, the one by Govindu has been provided by the author, while in the other cases we used our implementation. All the simulations were performed in Matlab on a dual-core computer with 1.3 GHz processor.

5.1   Simulated Data

We evaluated the aforementioned methods assuming that pairwise motions have been computed beforehand, thus 3D points are not involved in these simulations. Specifically, we generated a redundant set of relative motions (possibly corrupted by noise and/or outliers) – simulating the output of a pairwise registration algorithm such as ICP – which were given as input to all the techniques, and we evaluated the final estimates of absolute motions.

We considered n absolute motions in which rotations were sampled from random Euler angles and translation coordinates follow a standard Gaussian distribution. The measurement graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a random graph drawn from the Erdős-Rényi model with parameters (n, p), *i.e.* given a vertex set $\mathcal{V} = \{1, 2, \ldots, n\}$ each edge (i, j) is in the set $\mathcal{E}$ with probability $p \in [0, 1]$, independently of all other edges. Thus $(1 - p)$ controls the degree of sparsity of the graph ($p = 1$ corresponds to the complete graph). The inlier pairwise motions were corrupted by a multiplicative noise $\widehat{M}_{ij} = M_{ij} E_{ij}$, with $E_{ij} \in SE(3)$ representing a small perturbation of the identity matrix. The rotation component of E_{ij} has axis uniformly distributed over the unit sphere and angle following a Gaussian distribution with zero mean and standard deviation $\sigma_R \in [1°, 10°]$, and the translation components were sampled from a Gaussian distribution with zero mean and standard deviation $\sigma_T \in [0.01, 0.1]$. All the results were averaged over 50 trials. In order to compare estimated and theoretical absolute motions we computed the transformation that aligns them by applying single averaging [38] for the rotation term and least-squares for the translation term. We used the angular distance and Euclidean norm to measure the accuracy of estimated rotations and translations respectively, where the angular (or geodesic) distance between two rotations $R, S \in SO(3)$ is the angle (in the angle-axis space) of the rotation SR^{T} which lies in the range $[0, 180°]$.

It is hard to evaluate the performances of a registration method as a whole, since several parameters are involved, thus in the following simulations we let one parameter vary at a time and keep the others fixed.

Noise. In this experiment we evaluate the effect of noise on relative motions in the absence of outliers. We considered $n = 100$ absolute motions and $p = 0.3$, which corresponds to about 70% of missing pairs. Results with higher values of p correspond to better conditioned problems, with the same qualitative behaviour as $p = 0.3$, and hence they are not reported. Fig. 1 shows the mean errors on absolute motions (rotation errors are measured in degrees while translation errors are commensurate with the simulated data) obtained by all the analysed techniques, as a function of the standard deviation of noise. The worst resilience to noise is achieved by Sharp *et al.* while LRS decomposition techniques and the remaining algorithms return good estimates of absolute motions. A possible explanation of such behaviour is that in Sharp *et al.*'s method the error is distributed among the motions but it is not reduced. The best accuracy is achieved by non robust methods: DIFFUSION, NULL-SPACE and Govindu's. In
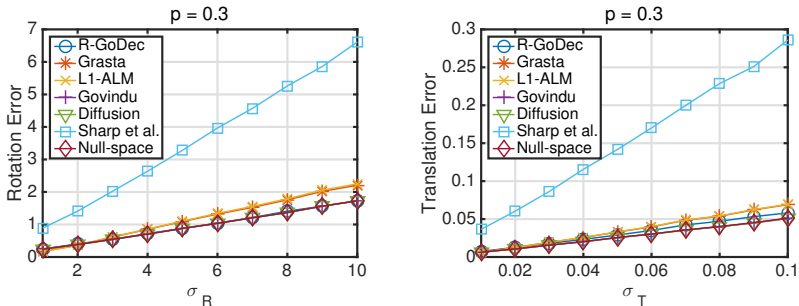
Fig. 1: Mean errors on absolute motions as a function of σ_R and σ_T.
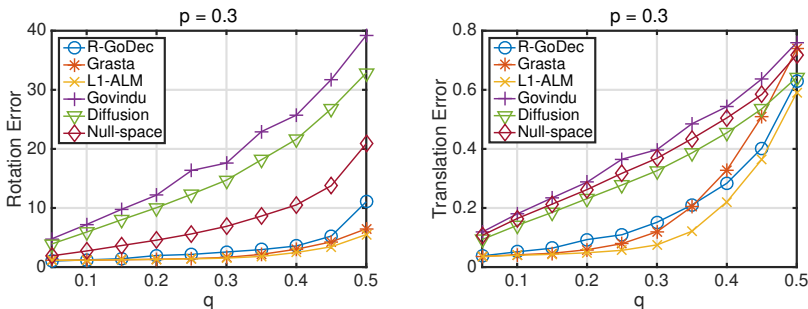

Fig. 2: Mean errors on absolute motions as a function of q.

general, robust techniques are not optimal with respect to noise, since they trade robustness for statistical efficiency.

Outliers. In this experiment we study the robustness to outliers of our approach. Each edge $(i, j) \in \mathcal{E}$ was designated as an outlier with uniform probability $q \in [0, 1]$, independently of all other edges. Outlier edges were assigned random elements of $SE(3)$. We considered $n = 100$ absolute motions sampled as before, we chose $p = 0.3$ to define the density of the measurement graph, and we introduced a fixed level of noise on relative motions ($\sigma_T = 0.05, \sigma_R = 5°$). The probability q that an edge is outlier ranges from 0.05 to 0.5, which correspond to about 5% and 50% of effective outliers. Figure 2 shows the mean errors on absolute motions as a function of q obtained by our approach, DIFFUSION, NULL-SPACE and Govindu's. The errors obtained by Sharp *et al.* are not reported in Fig. 2 so as to better visualize differences between the remaining algorithms (the method by Sharp *et al.* yields an average rotation error of 20° for $q = 0.05$ and 100° for $q = 0.5$). Figure 2 confirms that DIFFUSION, NULL-SPACE and the method by Govindu are not robust, and it clearly shows the resilience to outliers gained by R-GODEC, GRASTA and L1-ALM. In particular, the errors obtained by LRS decomposition techniques remain almost unchanged until $q = 0.4$ for rotations and $q = 0.3$ for translations.
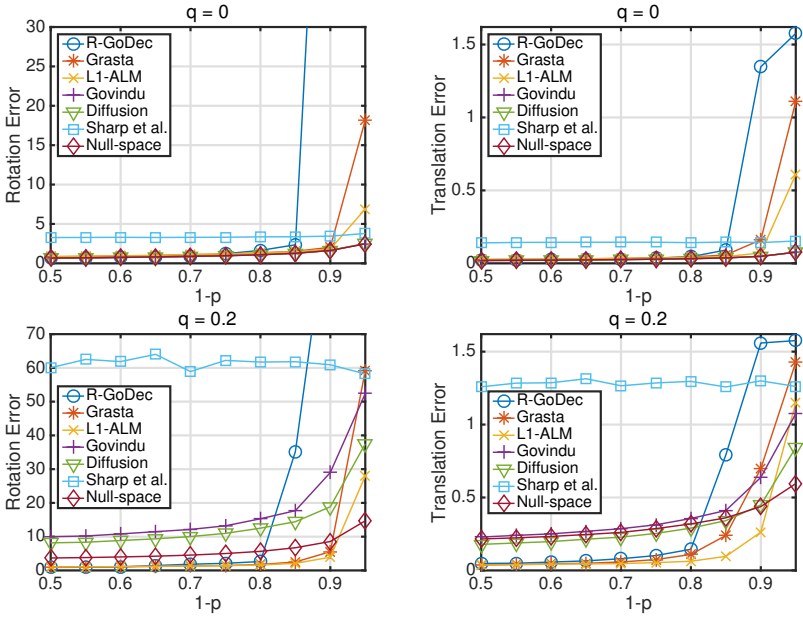
Fig. 3: Mean errors on absolute motions as a function of $(1 - p)$, with $q = 0$ (top) and $q = 0.2$ (bottom). In the left sub-figures, the average rotation errors of R-GoDec are approximately $90°$ for $(1 - p) = 0.9$ and $120°$ for $(1 - p) = 0.95$.
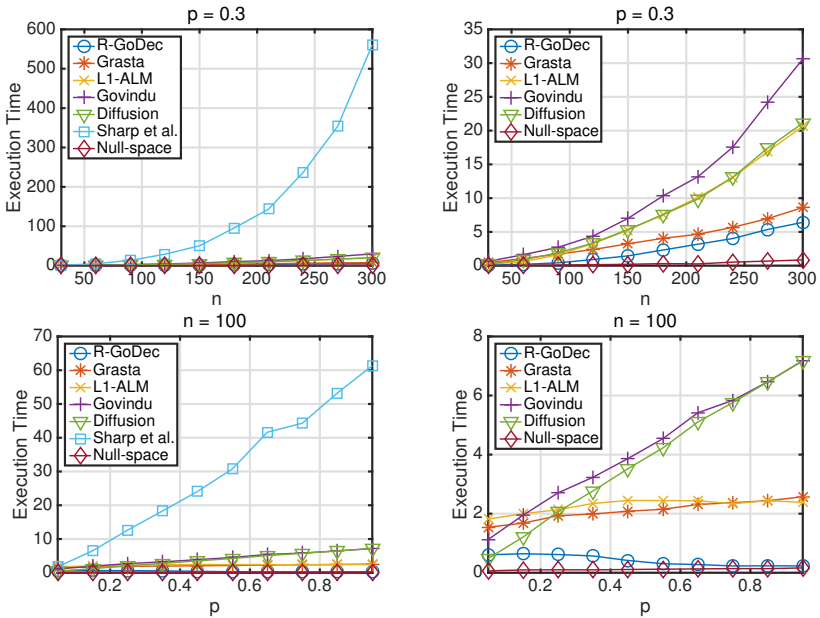

Fig. 4: Execution times (seconds) as a function of n (top) and p (bottom). The right figures are a magnification of the left ones.

Missing Data. In this experiment we study how missing data influence the performances of our approach. We considered $n = 100$ absolute motions sampled as before and we introduced a fixed level of noise on relative motions ($\sigma_T = 0.05, \sigma_R = 5°$). The sparsity parameter $(1 - p)$ ranges from 0.5 to 0.95, which correspond to about 50% and 95% of missing pairs. Results with lower values of $(1-p)$ yield the same behaviour as $(1-p) = 0.5$, and hence they are not reported. We considered both the ideal case where outliers are absent ($q = 0$) and a more realistic situation in which a fixed percentage of outliers is introduced ($q = 0.2$). Results are reported in Fig. 3, which shows the mean errors on absolute motions as a function of the sparsity parameter $(1 - p)$. The errors obtained by Sharp *et al.* remain constant as $(1 - p)$ increases, showing no sensitivity to missing data. The same holds for the method by Govindu, DIFFUSION and NULL-SPACE, if outliers are not present. As for our approach, GRASTA and L1-ALM can tolerate up to 90% of missing pairs in the case $q = 0.2$, whereas R-GODEC breaks down with 80% of missing pairs. If there are no outliers ($q = 0$), all the LRS methods can tolerate an extra 5% of missing data.

Execution Time. In this experiment we assess the computational efficiency of all the methods in two scenarios. First, we kept the density level of the measurement graph fixed ($p = 0.3$) and let n vary between 30 and 300. Then, we kept the number of absolute motions fixed ($n = 100$) and let p vary between 0.05 (about 95% of missing data) and 0.95 (about 5% of missing data). In both cases we introduced a fixed level of noise and outliers on relative motions ($\sigma_T = 0.05, \sigma_R = 5°$, $q = 0.2$). DIFFUSION is implemented in C++ (by the authors), while the remaining algorithms are implemented in MATLAB. Results are reported in Fig. 4, showing that the method by Sharp *et al.* is remarkably slower than the other techniques. In particular, L1-ALM is comparable to DIFFUSION and faster than Govindu's, while both R-GODEC and GRASTA are slower than NULL-SPACE but faster than the other methods. The bottom row in Fig. 4 shows that the execution time of matrix decomposition techniques and NULL-SPACE do not change significantly when p varies, whereas the other techniques require more time as the measurement graph gets denser.

The rundown of these tests is that, collectively, motion synchronization methods based on LRS decomposition qualify among the fastest solutions and provide a good trade-off between statistical efficiency and resilience to outliers. However, they are more affected than the other methods by the sparsity of the graph.

5.2   Real data

In this section we report the outcome of tests on real datasets of range images. Relative motion estimates were produced thanks to the MATLAB implementation of ICP (`pcregrigid`). The measurements graph was defined by discarding all the pairs with registration error higher than a threshold. This produced a redundant set of relative motions which were compensated by solving a motion synchronization problem, returning the transformations that align the original point

Table 1: Mean errors (rotations in degrees, translations in millimetres) on absolute motions for the Stanford repository. The number of point sets and the percentage of missing pairs are also reported.

Dataset	n	% miss.	R-GoDec		Grasta		L1-ALM		Govindu		Diffusion		Sharp et al.		Null-space	
			rot.	tra.	rot.	tra.	rot.	tra.	rot.	tra.	rot.	tra.	rot.	tra.	rot.	tra.
Bunny	10	0	0.82	2.9	0.84	1.9	0.78	1.6	1.07	3.7	1.07	3.7	1.07	4.5	1.07	3.7
Buddha	15	0	0.85	0.3	0.79	0.4	0.94	0.4	1.28	0.4	1.28	0.4	2.22	0.6	1.27	0.4
Dragon	15	0	0.79	0.3	0.91	0.4	0.77	0.4	1.52	0.4	1.52	0.4	1.45	0.6	1.51	0.4

Table 2: Execution times (seconds) of motion synchronization. The number of point sets and the percentage of missing pairs are also reported.

Dataset	n	% miss.	R-GoDec	Grasta	L1-ALM	Govindu	Diffusion	Sharp et al.	Null-space
Gargoyle	27	40	0.05	0.47	0.15	0.13	0.35	1.94	0.03
Capital	100	71	0.67	1.62	1.66	0.98	2.53	25.27	0.09

sets. These estimates could have been improved by alternating motion synchronization and computing relative motions, as suggested in [23, 24]; however, such a refinement was not applied in these experiments, *i.e.* we performed motion synchronization only once. Experimentally we observed that LRS decompositions perform better when translation components have values comparable to rotations, namely in the range $[-1, 1]$. For this reason, before performing motion synchronization, we divided all the relative translations by the maximum of the translations norm (and eventually multiplied the absolute translations by such a scale). This normalization also improves the results of the other algorithms.

From the Stanford 3D Scanning Repository [39] we used the Bunny, Happy Buddha (standing) and Dragon (standing) datasets, which contain 10, 15 and 15 point sets, respectively. As for the initialization of the ICP algorithm, we perturbed the available ground-truth motions by a rotation with random axis and angle uniformly distributed over $[0, 2°]$, similarly to the experiments carried out in [24]. Since ground-truth motions are available for these datasets, we evaluated quantitatively the results by reporting the mean errors in Tab. 1. Differences in execution time are meaningless for such relatively small datasets and are not reported. The errors obtained by R-GoDec, Grasta and L1-Alm are always lower than the other techniques, highlighting the benefit of robustness.

In another experiment we considered two datasets, named Gargoyle and Capital, which contain 27 and 100 point sets respectively. Since there is no information about the scans, we simply initialized the ICP algorithm with identity matrices. Execution times are reported in Tab. 2. They are referred to the motion synchronization step, *i.e.* computing absolute motions from relative motions, and they do not include the time for computing relative motions, which is the same for all the techniques. R-GoDec is slower than Null-space but faster than the other solutions, the method by Sharp *et al.* is the slowest tech-

Table 3: Cross-sections of registered point-sets.

Dataset	R-GoDec	Grasta	L1-ALM	Govindu	Diffusion	Sharp et al.	Null-space
Bunny							
Buddha							
Dragon							
Gargoyle							
Capital							

nique, while GRASTA and L1-ALM are faster than DIFFUSION but slower than
NULL-SPACE and Govindu's method.

The different registration techniques can be appraised qualitatively from the
cross-sections of output 3D models reported Tab. 3, as it is customary in the reg-
istration literature. The cross-sections obtained by our approach are crisper than
the others, proving the effectiveness of LRS decomposition in handling measure-
ment errors in the context of multiple point-set registration. In particular, the
best visual accuracy is achieved by L1-ALM and GRASTA, while R-GoDEC get
slightly worse results, yet better than the remaining methods. There is no signif-
icant difference between the cross-sections obtained by DIFFUSION, NULL-SPACE
and Govindu's, while the misalignment produced by Sharp *et al.* is evident, es-
pecially for the Gargoyle dataset. Figure 5 shows the 3D models produced by
L1-ALM with different colours for each point cloud.

In summary, these experiments with real data confirms the conclusions drawn
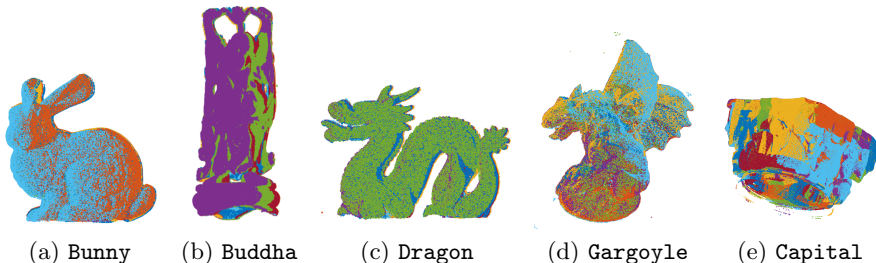from the simulations.

(a) Bunny   (b) Buddha   (c) Dragon   (d) Gargoyle   (e) Capital

Fig. 5: 3D models obtained with L1-ALM. Different point sets are colour coded.

6   Conclusions

For the first time in the literature we formulated frame-space registration as a
low-rank and sparse decomposition problem that neatly caters for missing-data,
outliers and noise, and it benefits from a wealth of available decomposition al-
gorithms that can be seamlessly used as alternatives. Experimental results show
that this approach is efficient and provides a good trade-off between statistical
efficiency and resilience to outliers. However, it is more affected than the other
methods by the sparsity of the measurement graph. It must be said, though,
that the goal of synchronization is to exploit redundancy: if the measures are
barely sufficient the problem looses significance.

References

1. Besl, P., McKay, N.: A method for registration of 3-D shapes. IEEE Transactions on Pattern Analysis and Machine Intelligence **14**(2) (February 1992) 239–256
2. Chen, Y., Medioni, G.: Object modeling by registration of multiple range images. In: Proceedings of the IEEE International Conference on Robotics and Automation. (1991) 2724 – 2729
3. Rusinkiewicz, S., Levoy, M.: Efficient variants of the ICP algorithm. In: Proceedings of the International Conference on 3-D Digital Imaging and Modeling. (2001) 145–152
4. Pulli., K.: Multiview registration for large data sets. In: Proceedings of the International Conference on 3-D Digital Imaging and Modeling. (1999) 160–168
5. Arrigoni, F., Fusiello, A., Rossi, B.: Spectral motion synchronization in SE(3). ArXiv e-prints (1506.08765) (2015)
6. Govindu, V.M.: Lie-algebraic averaging for globally consistent motion estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2004) 684–691
7. Williams, J., Bennamoun, M.: Simultaneous registration of multiple corresponding point sets. Computer Vision and Image Understanding **81**(1) (2001) 117–141
8. Pennec, X.: Multiple registration and mean rigid shape: Applications to the 3D case. In: 16th Leeds Annual Statistical Workshop. (1996) 178–185
9. Benjemaa, R., Schmitt, F.: A solution for the registration of multiple 3D point sets using unit quaternions. In: Proceedings of the European Conference on Computer Vision. (1998) 34–50
10. Bergevin, R., Soucy, M., Gagnon, H., Laurendeau, D.: Towards a general multiview registration technique. IEEE Transactions on Pattern Analysis and Machine Intelligence **18**(5) (May 1996) 540–547
11. Krishnan, S., Lee, P.Y., Moore, J.B., Venkatasubramanian, S.: Optimisation-on-a-manifold for global registration of multiple 3D point sets. International Journal of Intelligent Systems Technologies and Applications **3**(3/4) (2007) 319–340
12. Bonarrigo, F., Signoroni, A.: An enhanced 'optimization-on-a-manifold' framework for global registration of 3D range data. Proceedings of the Joint 3DIM/3DPVT Conference: 3D Imaging, Modeling, Processing, Visualization and Transmission (2011) 350 – 357
13. Mateo, X., Orriols, X., Binefa, X.: Bayesian perspective for the registration of multiple 3D views. Computer Vision and Image Underst. **118** (2014) 84 – 96
14. Chaudhury, K.N., Khoo, Y., Singer, A.: Global registration of multiple point clouds using semidefinite programming. SIAM Journal on Optim. **25**(1) (2015) 468 – 501
15. Raghuramu, A.C.: Robust multiview registration of 3D surfaces via ℓ_1-norm minimization. In: British Machine Vision Conference. (2015)
16. Fantoni, S., Castellani, U., Fusiello, A.: Accurate and automatic alignment of range surfaces. In: Second Joint 3DIM/3DPVT Conference: 3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT). (2012) 73 – 80
17. Fitzgibbon, A.: Robust registration of 2D and 3D point sets. Image and Vision Computing **21**(13-14) (2003) 1145 – 1153
18. Toldo, R., Beinat, A., Crosilla, F.: Global registration of multiple point clouds embedding the generalized procrustes analysis into an ICP framework. In: 3D Data Processing Visualization and Transmission Conference. (2010)
19. Beinat, A., Crosilla, F.: Generalized procrustes analysis for size and shape 3D object reconstruction. In: Optical 3-D Measurement Techniques. Wichmann Verlag (2001) 345–353

20. Thomas, D., Matsushita, Y., Sugimoto, A.: Robust simultaneous 3D registration via rank minimization. Proceedings of the Joint 3DIM/3DPVT Conference: 3D Imaging, Modeling, Processing, Visualization and Transmission (2012) 33–40
21. Sharp, G.C., Lee, S.W., Wehe, D.K.: Multiview registration of 3D scenes by minimizing error between coordinate frames. In: Proceedings of the European Conference on Computer Vision. (2002) 587–597
22. Fusiello, A., Castellani, U., Ronchetti, L., Murino, V.: Model acquisition by registration of multiple acoustic range views. In: Proceedings of the European Conference on Computer Vision. (2002) 805–819
23. Torsello, A., Rodola, E., Albarelli, A.: Multiview registration via graph diffusion of dual quaternions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2011) 2441 – 2448
24. Govindu, V.M., Pooja, A.: On averaging multiview relations for 3D scan registration. IEEE Transactions on Image Processing **23**(3) (2014) 1289–1302
25. Bernard, F., Thunberg, J., Gemmar, P., Hertel, F., Husch, A., Goncalves, J.: A solution for multi-alignment by transformation synchronisation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2015)
26. Shih, S., Chuang, Y., Yu, T.: An efficient and accurate method for the relaxation of multiview registration error. IEEE Trans. Image Proc. **17**(6) (2008) 968 – 981
27. Arrigoni, F., Magri, L., Rossi, B., Fragneto, P., Fusiello, A.: Robust absolute rotation estimation via low-rank and sparse matrix decomposition. In: Proceedings of the International Conference on 3D Vision (3DV). (2014) 491–498
28. He, J., Balzano, L., Szlam, A.: Incremental gradient on the Grassmannian for online foreground and background separation in subsampled video. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2012) 1568 – 1575
29. Zheng, Y., Liu, G., Sugimoto, S., Yan, S., Okutomi, M.: Practical low-rank matrix approximation under robust L_1-norm. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2012) 1410–1417
30. Zhou, X., Yang, C., Zhao, H., Yu, W.: Low-rank modeling and its applications in image analysis. ACM Computing Surveys **47**(2) (2014) 36:1–36:33
31. Candès, E.J., Li, X., Ma, Y., Wright, J.: Robust principal component analysis? Journal of the ACM **58**(3) (June 2011) 11:1–11:37
32. Candès, E.J., Recht, B.: Exact matrix completion via convex optimization. Foundations of Computational Mathematics **9**(6) (2009) 717–772
33. Candès, E.J., Tao, T.: The power of convex relaxation: near-optimal matrix completion. IEEE Transactions on Information Theory **56**(5) (2010) 2053–2080
34. Zhou, T., Tao, D.: GoDec: Randomized low-rank & sparse matrix decomposition in noisy case. In: Proceedings of the 28th International Conference on Machine Learning (ICML11), ACM (June 2011) 33–40
35. Beck, A., Teboulle, M.: A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM Journal on Imaging Sciences **2**(1) (March 2009) 183–202
36. Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J.: Distributed optimization and statistical learning via the alternating direction method of multipliers. Foundations and Trends in Machine Learning **3**(1) (January 2011) 1–122
37. Bertsekas, D.P.: Constrained Optimization and Lagrange Multiplier Methods. Academic Press (1982)
38. Hartley, R., Aftab, K., Trumpf, J.: L1 rotation averaging using the Weiszfeld algorithm. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2011) 3041–3048
39. Stanford University: 3D scanning repository http://graphics.stanford.edu/data/3Dscanrep/.