Multiview Registration via Graph Diffusion of Dual Quaternions

Andrea Torsello, Emanuele Rodolà, and Andrea Albarelli

torsello@dsi.unive.it rodola@dsi.unive.it albarelli@unive.it

Dipartimento di Scienze Ambientali, Informatica e Statistica - Università Ca' Foscari Venezia

Abstract

Surface registration is a fundamental step in the reconstruction of three-dimensional objects. While there are several fast and reliable methods to align two surfaces, the tools available to align multiple surfaces are relatively limited. In this paper we propose a novel multiview registration algorithm that projects several pairwise alignments onto a common reference frame. The projection is performed by representing the motions as dual quaternions, an algebraic structure that is related to the group of 3D rigid transformations, and by performing a diffusion along the graph of adjacent (i.e., pairwise alignable) views. The approach allows for a completely generic topology with which the pairwise motions are diffused. An extensive set of experiments shows that the proposed approach is both orders of magnitude faster than the state of the art, and more robust to extreme positional noise and outliers. The dramatic speedup of the approach allows it to be alternated with pairwise alignment resulting in a smoother energy profile, reducing the risk of getting stuck at local minima.

1. Introduction

Surface registration is a fundamental step in the reconstruction of three-dimensional objects. This is typically a two step process where all the views are first registered against each other, and then all the pairwise transformations are lowered to a common coordinate frame through a process commonly referred to as multiview registration.

The literature on pairwise registration is quite ample, with modifications to the original ICP proposed by Zhang [23] and Besl and McKay [4] taking the lion's share. ICP-based methods start from an initial pose estimate and iteratively refine it by minimizing a distance function measured between pairs of selected neighboring points. The variants generally differ in the strategies used to sample points from the surfaces, reject incompatible pairs, or measure error. In general, the precision and convergence speed of these techniques is highly data-dependent and very sensitive to the fine-tuning of the model parameters. Sev-

eral approaches that combine these variants have been proposed in the literature in order to overcome these limitations (see [19] for a comparative review). Some recent variants avoid hard culling by assigning a probability to each candidate pair by means of evolutionary techniques [15] or Expectation Maximization [10]. ICP variants, being iterative algorithms based on local, step-by-step decisions, are very susceptible to the presence of local minima. Other fine registration methods include the well-known approach by Chen [6] and signed distance fields matching [16].

By contrast, the literature of multiview registration is more diverse. In [6] Chen and Medioni propose to iteratively merge new views into a single metaview: The registration of a new view against the metaview is obtained with a common pairwise registration technique, such as ICP, and then the points of the new registered view are merged to the metaview; the approach is iterated until all the range images are merged. This metaview approach has problems since registration errors are accumulated rather than mediated. To solve this problem Bergevin *et al.* [3] match points in every view with all the views overlapping with it, and calculate a transformation that registers the first view using all the mating points. This process is iterated to convergence, thus diffusing the errors among all views. This is implicitly a diffusion process where the random walk in the transformation space is governed by the constraints offered by nearby views in the view-graph; however, convergence toward the steady-state is extremely slow and computationally demanding. Eggert *et al.* [9] constrain the pairings so that the points of each scan map with exactly one other point and then minimize the total distance between the paired points. This speeds up convergence, but can prevent the algorithm from converging to a correct solution as the views may cluster into groups that are well registered, without improving inter-group registration. With these iterative algorithms based on global point correspondences, how and when to apply the transformation remains an open issue: For example, Bergevin *et al.* [3] calculate a transformation for each view separately and then apply them simultaneously before the next round of matchings, while Benjemaa and Schmitt [2] apply the new transformations

independently as soon as they are calculated, and Eggert *et al.* [9] solve for the update by simulating a spring model. An alternative was explored in [11] where an approximate surface model is created and the view are registered against the model. The surface model is then iteratively refined using the new registrations.

In [18] Pulli takes a simplifying view that pairwise registrations are "as good as it gets" and that the role of multiview registration is only to project the transformation into a common reference frame in such a way as to limit the accumulation of registration errors. To this end, he proposes a greedy approach that tries to limit the difference between the position of point sets as positioned in two frames and transformed by the pairwise registration of the two frames. More formally, he tries to keep the distortion $\mathcal{D}(S)$ of the points from a set S within a given tolerance ϵ, where

$$\mathcal{D}(S) = \sum_{s \in S} \sum_{(i,j) \in \mathcal{V}} ||P_i(s) - T_{ij}(P_j(s))||^2.$$

Here P_i is the transformation that maps a point into the coordinate system of view i, T_{ij} is the transformation that maps the coordinate frame j into the coordinate frame i obtained through pairwise registration, and $\mathcal{V}$ is the set of pairs of neighboring views for which pairwise registration is performed. Pulli suggests to sample the set of fiduciary points S from the surface of the object. Interestingly, by working only on the space of transformations, this approach limits the memory requirements since it does not need to keep all the points from all the views in memory at once. Note, however, that the approach cannot guarantee that an optimal solution will be found, nor that any solution within the given tolerance will be found.

More recently, Williams and Bennamoun [22] adopted a similar view, posing the problem as the minimization of the distortion on a set of fiduciary points and computing the minimization by an iterative approach optimizing each rotation via singular value decomposition.

In this paper we propose a novel multiview registration algorithm where the poses are estimated through a diffusion process on the view-adjacency graph. The diffusion process is over dual quaternions [7], a non-commutative and non-associative algebraic structure that is related to the group $SE(3)$ of 3D rigid transformations, leading to an approach that is both orders of magnitude faster than the state of the art, and more robust to extreme positional noise and outliers.

2. Dual Quaternions and 3D Transformations

Quaternions have been a popular geometrical tool for more than 20 years as they represent 3D rotations in a way that is arguably more efficient and robust than 3×3 rotation matrices [20]. Quaternions are an algebraic extension of complex numbers with 3 imaginary bases i, j,

and k, thus a quaternion is a number of the form $q = a + ix + jy + kz$. The multiplication of two quaternions is defined through the following multiplication rules for the three imaginary bases: $i^2 = j^2 = k^2 = -1$, $ij = k = -ji$, $jk = i = -kj$, $ki = j = -ik$. The *conjugate* of a quaternion $q = a + ix + jy + kz$ is the quaternion $q^* = a - ix - jy - kz$, while the *norm* of a quaternion is the quantity $||q|| = \sqrt{qq^*} = \sqrt{q^*q} = \sqrt{a^2 + x^2 + y^2 + z^2}$. Quaternions with unitary norm are called *unit quaternions*. In the following we will use the vectorial representation of quaternions: Let $\mathbf{i} = (i, j, k)$ be the row-vector of the imaginary bases, we can write the quaternion q as $a + \mathbf{iv}$, where $\mathbf{v} = (x, y, z)^T$ is a 3D vector. A right-handed 3D rotation of angle θ around the axis of unit vector $\mathbf{v}$ is in relation with the unit quaternion $q = \cos(\theta/2) + \sin(\theta/2)\mathbf{iv}$. In fact, let $\mathbf{p} = (p_x, p_y, p_z)^T$ be a 3D point and $\mathbf{p}_r$ its rotation, we have $\mathbf{ip}_r = q(\mathbf{ip})q^*$. The ring of quaternions, however, is a dual cover of the group $SO(3)$ of 3D rotations, as q and $-q$ represent the same rotation. Quaternions are particularly interesting since they allow for optimal interpolation between rotations. The famous Spherical Linear Interpolation (SLERP) algorithm [20] interpolates quaternions on the unit hypersphere and exhibits the following useful properties:

- **Shortest path:** the motion between the initial rotation R_0 and the final rotation R_1 is a rotation about a fixed axis with the smallest angle.

- **Constant speed:** the angle of the interpolated rotation varies linearly with respect to parameter t.

- **Coordinate system invariance:** the interpolation path does not change if we change the coordinate system.

On the other hand, linear interpolation followed by reprojection onto the unit hypersphere guarantees the first and third properties, but exhibits changes in speed, in particular it accelerates around the middle of the interpolation. This implies that a linear averaging of quaternions does not minimize the squared geodesic distance in the unit quaternion manifold in the same way that the mean of a set of points minimizes the squared Euclidean distances to the points. Unfortunately, SLERP does not generalize to the blending of several rotations. Buss and Fillmore [5] provide an iterative algorithm to find the proper (weighted) mean in the unit-quaternion manifold, while Kavan and Žára [14] show that the difference between spherical and linear interpolation is always less than 0.071 radians.

Dual quaternions are less known than quaternions, but their ability to efficiently represent rigid transformations has been successfully adopted in 3D animation and skinning [12], robot control and registration [1, 8], and have been used in theoretical kinematics for a long time [17]. Dual quaternions are an algebraic extension of quaternions

Table 1. Multiplicative table of dual quaternions.

	1	i	j	k	ϵ	ϵi	ϵj	ϵk
1	1	i	j	k	ϵ	ϵi	ϵj	ϵk
i	i	-1	k	$-j$	ϵi	$-\epsilon$	ϵk	$-\epsilon j$
j	j	$-k$	-1	i	ϵj	$-\epsilon k$	$-\epsilon$	ϵi
k	k	j	$-i$	-1	ϵk	ϵj	$-\epsilon i$	$-\epsilon$
ϵ	ϵ	ϵi	ϵj	ϵk	0	0	0	0
ϵi	ϵi	$-\epsilon$	ϵk	$-\epsilon j$	0	0	0	0
ϵj	ϵj	$-\epsilon k$	$-\epsilon$	ϵi	0	0	0	0
ϵk	ϵk	ϵj	$-\epsilon i$	$-\epsilon$	0	0	0	0

much like complex numbers are an extension of the reals. They are defined in terms of a dual basis ϵ that commutes with the imaginary bases; thus, a dual quaternion is a number of the form $q + \epsilon r$, where q and r are quaternions, and the product follows the multiplicative rule $\epsilon^2 = 0$, yielding

$$(q + \epsilon r)(s + \epsilon t) = qs + \epsilon(qt + rs).$$

This results in the multiplicative table shown in Table 1. Dual quaternions have three different conjugates:

$$(q + \epsilon r)^* = q^* + \epsilon r^* \quad (q + \epsilon r)^\dagger = q^* - \epsilon r^* \quad (q + \epsilon r)^+ = q - \epsilon r$$

The norm of a dual quaternion dq is $||dq|| = \sqrt{dq^* dq} = \sqrt{dq\, dq^*} = ||dq^*||$ and the inverse of a dual quaternion dq is $dq^{-1} = \frac{dq^*}{||dq||^2}$. The rigid transformation obtained by a rotation defined by the unit quaternion r and then a translation by $\mathbf{t} = (t_x, t_y, t_z)^T$, is represented by the dual quaternion

$$dq = r + \frac{1}{2}\epsilon \mathbf{i} \mathbf{t}\, r.$$

In fact, if we represent a 3D point $\mathbf{p}$ as $1 + \epsilon \mathbf{i} \mathbf{p}$, the following holds:

$$(r + \frac{1}{2}\epsilon \mathbf{i} \mathbf{t} r t)(1 + \epsilon \mathbf{i} \mathbf{p})(r + \frac{1}{2}\epsilon \mathbf{i} \mathbf{t} r)^\dagger =$$
$$(r + \frac{1}{2}\epsilon \mathbf{i} \mathbf{t} r + \epsilon r \mathbf{i} \mathbf{p})(r^* - \frac{1}{2}\epsilon(\mathbf{i} \mathbf{t} r)^*) =$$
$$(r + \frac{1}{2}\epsilon \mathbf{i} \mathbf{t} r + \epsilon r \mathbf{i} \mathbf{p})(r^* + \frac{1}{2}\epsilon r^* \mathbf{i} \mathbf{t}) = 1 + \epsilon(r \mathbf{i} \mathbf{p} r^* + \mathbf{i} \mathbf{t}).$$

Thus, $1 + \epsilon \mathbf{i} \mathbf{p}$, i.e., the dual quaternion representation of the point $\mathbf{p}$, gets mapped into the dual quaternion $1 + \epsilon(r \mathbf{i} \mathbf{p} r^* + \mathbf{i} \mathbf{t})$, which is the representation of $\mathbf{p}$ after the rotation and translation have been applied. In fact, $r \mathbf{i} \mathbf{p} r^*$ represents the rotation by r with the usual quaternion notation, while the addition of $\mathbf{i} \mathbf{t}$ takes care of the subsequent translation. Further, any dual quaternion $q + \epsilon r$ with $||q|| = 1$ and $q \cdot r = 0$ represents a rigid transformation. Here $\cdot$ represents the standard dot product in the quaternion viewed as a four-dimensional vector space over $\mathbb{R}$. However, the dual quaternion representation is not unique since, as with the normal quaternions, dq and $-dq$ represent the same transformation.

In [13, 12] the authors present ScLERP, a generalization of the SLERP interpolation algorithm for dual quaternions.

Let $\alpha(t)$, $\mathbf{a}(t)$, $\delta(t)$, and $\mathbf{d}(t)$ be respectively the rotation angle and axis, and the translation magnitude and direction, then ScLERP was shown to have the following properties: a) $\mathbf{a}(t)$, and $\mathbf{d}(t)$ are constant and $\alpha(t) \in [-\pi; \pi]$ (shortest path); b) $\frac{d}{dt}\alpha(t) = 0$ and $\frac{d}{dt}\delta(t) = 0$ (constant speed). Further, it is invariant to changes in the coordinate system. In [13] was presented an iterative algorithm called Dual quaternion Iterative Blending (DIB) for averaging dual quaternions in a way that minimizes the (weighted) squared geodesic distances between the target mean and the input quaternions in the Riemannian manifold of unit dual quaternions, and it was shown that the variation between the proper geodesic average and a linear blending followed by a reprojection has an upper bound of 0.143 radians in rotation and a relative variation of 15% in translation, while in general the differences remain much smaller. In particular, if the set of dual quaternions we want to blend has small variance, the linear average and the geodesic average converge rapidly, since the difference between the geodesic and Euclidean distance is $O(\theta)$ where θ is the angle of rotation. In the following we will use the notation $\mathrm{ScAVG}(q_1, \ldots, q_n)$ to refer to geodesic mean of the quaternions $(q_1, \ldots, q_n)$ as obtained by applying DIB with uniform weights.

3. View-Graph Diffusion

We cast the multiview registration problem into a diffusion of rigid transformations over the view-graph, i.e., a graph in which nodes correspond to the range images and the edges reflect the adjacency relation between views, or, equivalently, the existence of an overlap between the scans. Let V_i be a transformation taking the coordinate frame of view i into a global coordinate frame, and T_{ij} the result of the pairwise registration taking the coordinate frame of view j into the frame of view i. Then, if the pairwise registration was noise-free, we would have $T_{ji} * V_i = V_j$ for all adjacent views i and j. In this setting the problem of multiview registration is that of finding a set of rigid motions from each view to a common frame of reference, say that of view 0, such that a measure of distortion between the position V_i of view i and the position $T_{ij}V_j$ obtained from the composition of position V_j and the pairwise registration T_{ij} is minimized, i.e., we seek to minimize the functional

$$D = \sum_i \sum_{j \in N(i)} d(T_{ij}V_j, V_i)$$

for an appropriate distortion function d. Here $N(i)$ is the set of neighbors of view i. This is in spirit similar to the approach taken by Pulli [18], where the distortion function is the (squared) Euclidean distance between the final position of a set of points S transformed with motions V_i and $T_{ij}V_j$:

$$D_P = \sum_i \sum_{j \in N(i)} \sum_{\mathbf{p} \in S} ||T_{ij}V_j\mathbf{p} - V_i\mathbf{p}||^2.$$
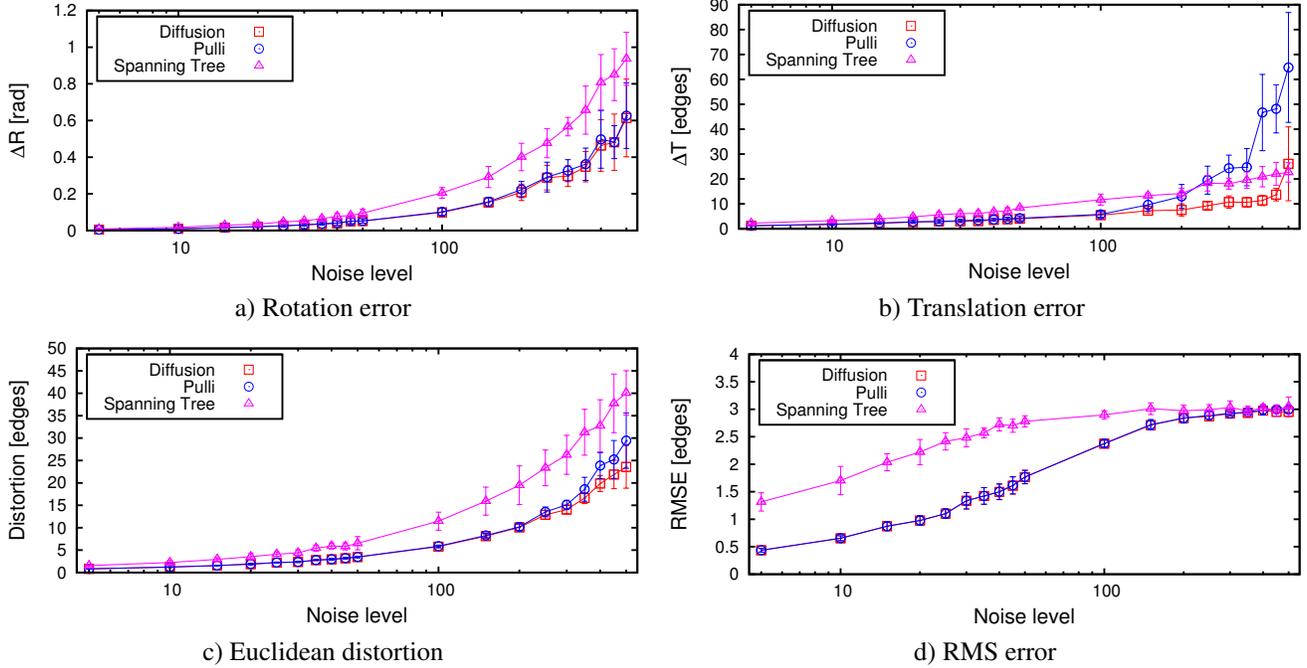
Figure 1. Comparison of the different methods in the synthetic experiments at various levels of noise.

We adopt a different measure of distortion that derives from the fact that any rigid transformation is in fact a screw motion, i.e., a rotation around an axis placed anywhere in the 3D space, and a translation along the direction of the axis. We define the *screw distance* $d_{SC}(q_i, q_2, \mathbf{p})$ as the length of the screw path of the point $\mathbf{p}$ along the transformation $\hat{q} = q_i^\dagger q_2$, i.e., $d_{SC}(q_i, q_2, \mathbf{p}) = \sqrt{t^2 + \alpha^2 r_{\mathbf{p}}^2}$ where t is the length of the translation along the axis of $\hat{q}$, α is the rotation angle, and $r_{\mathbf{p}}$ is the distance of $\mathbf{p}$ from the rotation axis. The *screw distortion* is then defined as

$$D_{SC} = \sum_i \sum_{j \in N(i)} \sum_{\mathbf{p} \in S} d_{SC}(T_{ij}V_j, V_i, \mathbf{p})^2 \,.$$

A direct consequence of the fact that ScLERP interpolation is both shortest path and constant speed is that, given a set of dual quaternions $Q = q_i, \ldots, q_n$, their screw average $\hat{q} = \text{ScAVG}(q_i, \ldots, q_n)$ minimizes the sum of squared screw distances $\sum_i d_{SC}(q_i, \hat{q}, \mathbf{p})$ for any point $\mathbf{p} \in \mathbb{R}^3$ [21]. That is, by measuring along the curved screw path, the transformation that minimizes the distortion does not depend on the points selected, which was arguably the most problematic aspect with the Euclidean distortion adopted by Pulli. Further, when the variation in orientation among the dual quaternions $q_i, \ldots, q_n$ is very small, we have that $d_{SC}(T_{ij}V_j, V_i, \mathbf{p}) \approx ||T_{ij}V_j\mathbf{p} - V_i\mathbf{p}||$ for any point $\mathbf{p} \in \mathbb{R}^3$. In fact, we have $d_{SC}(T_{ij}V_j, V_i, \mathbf{p})^2 = d_E^{\parallel \, 2} + \frac{\theta/2}{\sin(\theta/2)} d_E^{\perp \, 2}$ where θ is the rotation angle, and $d_E^{\parallel}$ and $d_E^{\perp}$ are respectively the components of the Euclidean distance parallel and orthogonal to the axis.

The optimal multiview alignment is thus obtained by computing the steady-state of the following process

$$V_i^{t+1} = \underset{j \in N(i)}{\text{ScAVG}}(\pm T_{ij}V_j^t)$$

where the sign uncertainty is a consequence of the sign uncertainty in the dual quaternion representation, and is chosen so that $\pm T_{ij}V_j^t \cdot V_i^t > 0$. Further, since for small and moderate rotational variability in the vectors the linear average approximates well the screw average, while being much faster, in all our experiments we are using linear averages.

Finally, the proposed approach is a refinement method that requires initial motion estimates, but these can be computed by simple composition of the transformations along adjacent views with a breadth-first visit starting from the view 0. A demo application and code is available at http://www.dsi.unive.it/~rodola/.

4. Experimental Evaluation

Multiview registration techniques have both a sparse and diverse coverage in literature, and as such they suffer from the lack of a robust and fair methodology for performance assessment and comparison. Specifically, in real scenarios, where ground-truth data is not available, it can be very hard to evaluate and quantify the results of a global alignment and settle for a solution, without resorting to a thorough and time-consuming analysis of the registered views.
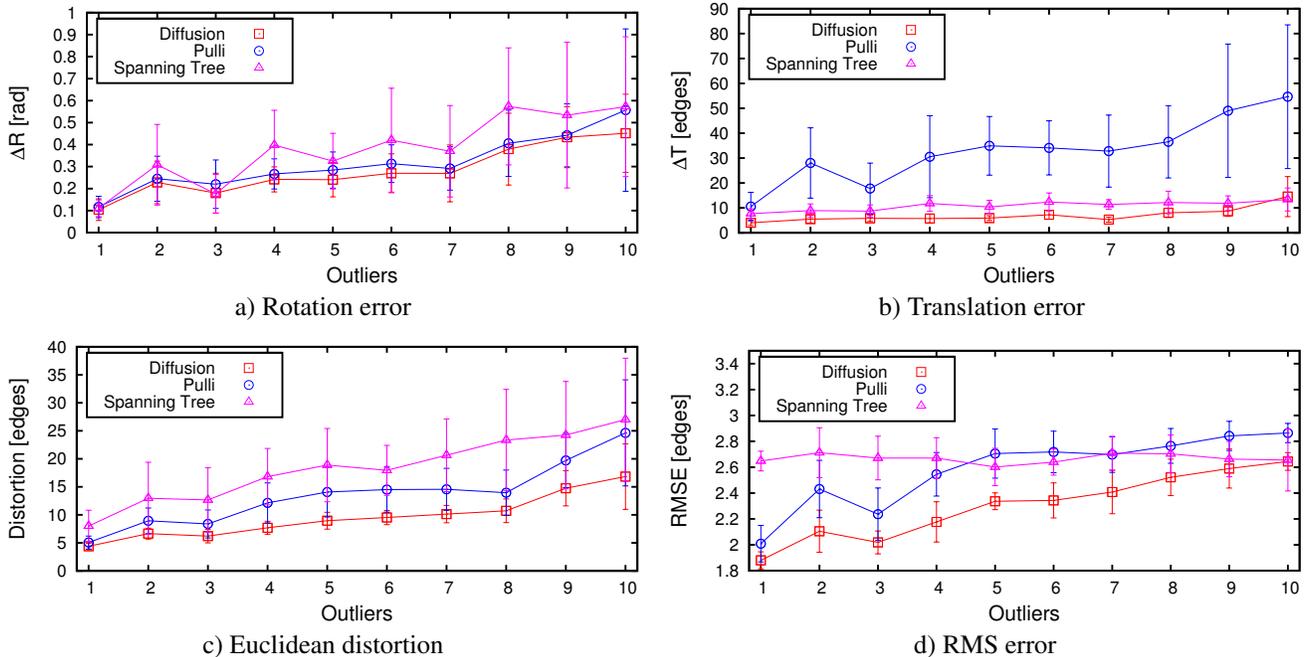
Figure 2. Comparison of the different methods in the synthetic experiments with different numbers of outliers.

To this end, we performed a wide range of experiments with both synthetic and real-world data. For each complete 3D model (from Georgia Tech's large geometric models archive[1]), a total of 36 orthographic snapshots were taken, each at a different angle of view; these, together with the ground-truth rigid motions used to produce the range images, constitute the dataset over which synthetic experiments were performed. In all the experiments we compare our method against Pulli's algorithm (as implemented by the author in the Scanalyze[2] software package), currently the method of choice in many applications. We evaluated the performance of the two algorithms, together with the initialization results obtained through a breadth-first coverage of the view graph (indicated as *Spanning Tree*), under different noise conditions and connectivity levels. For all the experiments we show the relative displacement with respect to ground-truth motion (with ΔR being the angle between the two unit quaternions, and ΔT the translation error expressed in median edge length), the RMS error among all the ranges and range 0 (point pairs were obtained through normal-shooting in both directions), and the Euclidean distortion metric adopted by Pulli (indicated as *Distortion*).

In Fig. 1 we show the results at different levels of initial displacement, where every view in the graph is connected with the next two in a ring topology. *Noise level* refers to a quantity which is proportional to the amount of Gaussian noise applied to the ground-truth pairwise mo-

tions, and ranges from a few units of edges and radians to tens of units; for each noise level, 10 independent runs of each method were performed. Both the Diffusion and Pulli methods compensate well rotational errors and are comparably good at low levels of noise, whereas the latter is outperformed when positional noise increases, both in terms of translation error and Euclidean distortion. It is worth noting that when we perturbed either the rotational or translational part of the rigid motion, keeping the other fixed, Spanning Tree and Pulli's methods performed a joint optimization modifying both components, while our method never changes the already optimal part of the motion, yielding better results at all levels of noise.

In Fig. 2 we assess the resilience of the tested methods to the presence of outliers: starting from a close-to-optimal initialization, we introduced strong pairwise misalignments so as to simulate a realistic scenario in which pairwise registrations get stuck at local minima. Connectivity is the same as in the previous experiments. In these graphs, the x-axis grows with the number of such mis-registrations; here, 20 runs were performed at each level, and for each run random pairs were picked from a uniform distribution and perturbed strongly. It can be seen that all the methods handle well rotational errors, with the Diffusion method giving particularly good performance constantly, even when the number of outlying pairs becomes large. Fig. 2b) and Fig. 2d) interestingly show the inability of Pulli's method to deliver good results in such situations: this is an inherent weakness of the method, since it acts in such a way to minimize all

---

[1] http://www.cc.gatech.edu/projects/large_models
[2] http://www.graphics.stanford.edu/software/scanalyze

a) Rotation error

b) Translation error

c) Euclidean distortion

d) RMS error
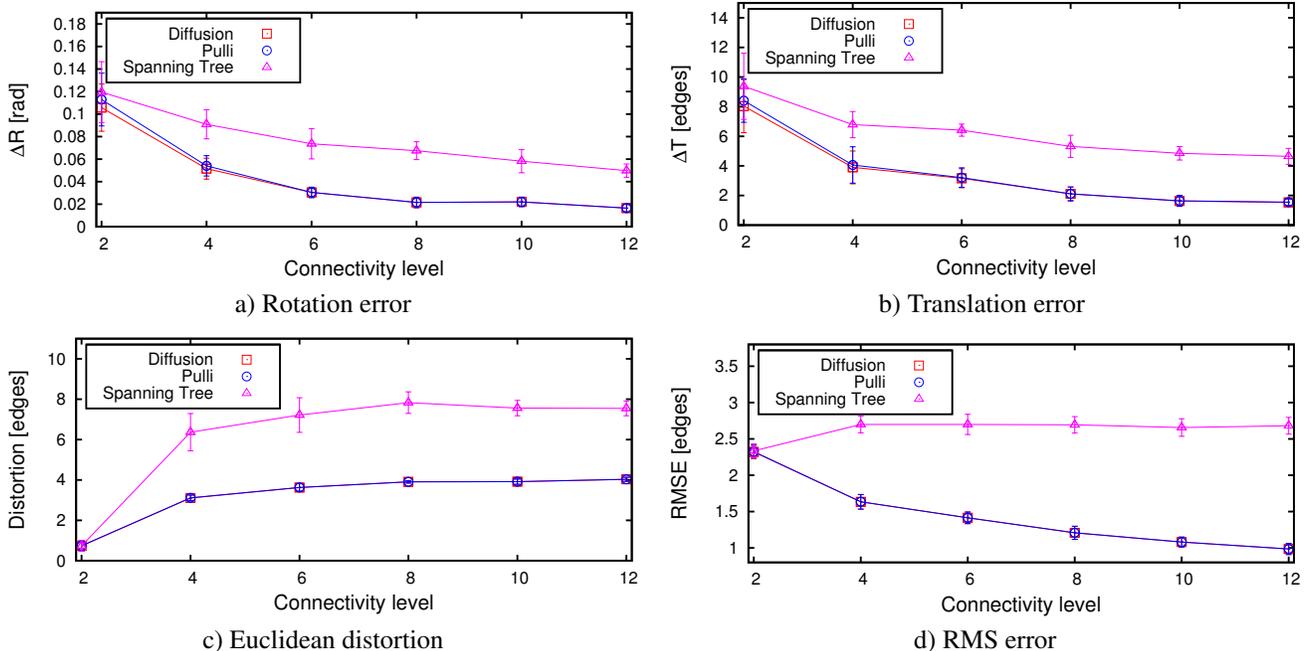
Figure 3. Comparison of the different methods in the synthetic experiments at various connectivity levels.
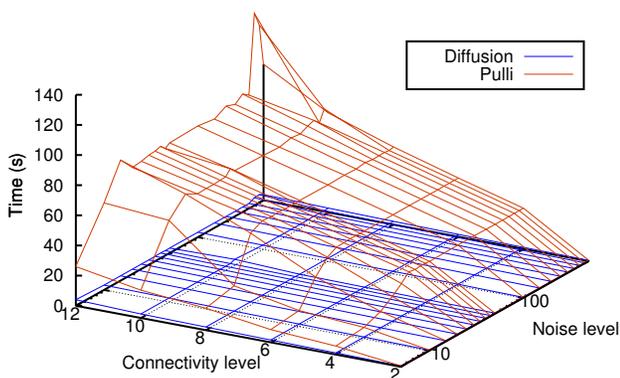
Figure 4. Computation times versus noise and connectivity levels.

motion, relying on the assumption that pairwise alignments are nearly perfect.

The next set of experiments (Fig. 3) is aimed at studying the effect of view-graph connectivity on the registration results. In these figures, *Connectivity level* refers to the number of links per view. As it can be readily seen, performance tends to increase with the number of edges in the view graph, as all three methods greatly benefit from more structure being brought in. It must be noted, however, that connectivity augmentation dramatically increases the time requirements of Pulli's method, bringing up convergence times by orders of magnitude (see Fig. 4).

Figure 4 compares the average convergence time of the proposed approach with the time required by Pulli's method. The times are shown as a function of noise level and connectivity. We can see that, with the exception of

extreme values, the times required by Pulli's approach are independent from the level of noise, but grow linearly with the connectivity level and are in almost all tested situations in the order of 10 to 100 seconds. The proposed approach exhibits the same independence with respect to noise and linear growth with respect to connectivity level, but it is always around 2 to 3 orders of magnitude faster.

In order to simulate a more realistic type of noise, we also tested the following setup: we perturbed the initial pairwise motion as for the first set of experiments, and then applied ICP to refine the alignment. This setup simulates a normal registration process with increasingly bad coarse registration, with the possibility that ICP gets stuck on local minima providing more structured outliers than the previous experiments. The results of this set of experiments can be seen in Fig. 5. Note that our approach and Pulli's method yield very similar results in all the metrics except for ΔT. This can be justified by the fact that, when caught in local minima, ICP slides the surface one over the other, resulting in strong translational outliers which Pulli's method cannot deal with. The very low RMSE derives from the fact that sliding along the surface each point still finds close-by mates on the other mesh. The proposed method, on the other hand, manages to smooth all the outliers effectively, yielding low errors in all the metrics.

Figure 6 shows an example of a real set of range images acquired with a scanner and aligned using Pulli's method and the proposed approach. While at a large scale the overall alignment appears similar, by examining closeups of various sections of the glasses we see that Pulli's method pro-
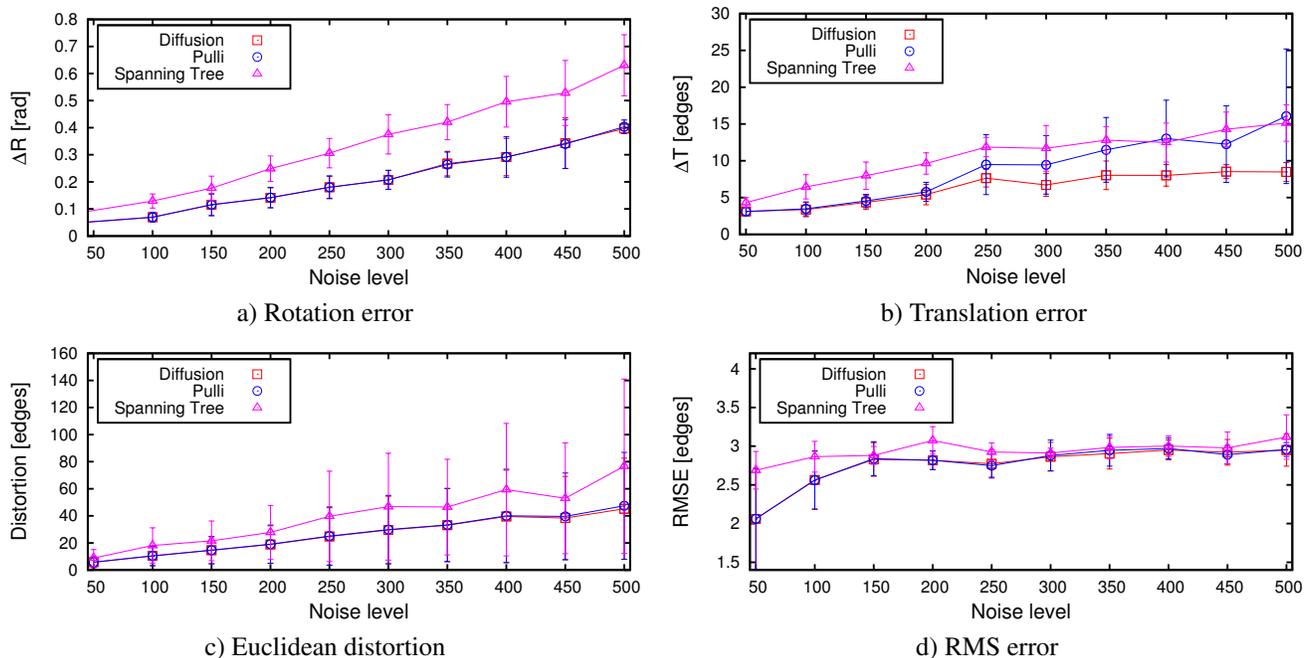
Figure 5. Comparison of the different methods in the experiments with motion refined by ICP at various levels of noise.

Pulli's method   Dual quaternion diffusion
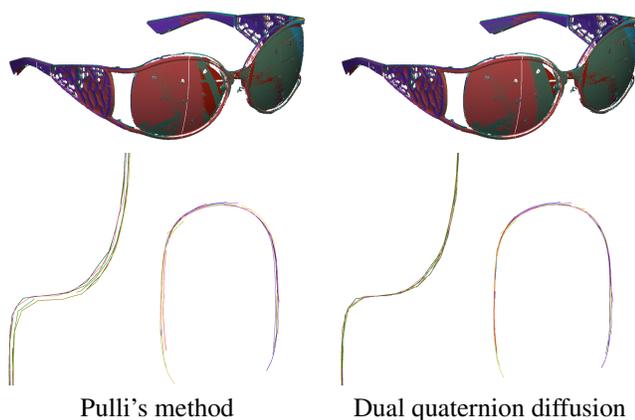
Figure 6. Global registration and closeup of slices of Pulli's method and our approach.

vides a slightly worse motion estimate, resulting in a wider stratification of the meshes.

It is worth noting that the orders of magnitude speedup provided by our approach makes it possible to run it several times combining it with pairwise registration. The idea is to alternate a few steps of ICP performed on all adjacent views with the diffusion. This way the diffusion process can be seen as a projection operator taking the incremental pairwise motion onto a set of consistent motion estimates. The advantage of this projection is that the constrained motion space smooths the energy profile of the resulting "global" ICP, reducing the risk of getting stuck in local minima. Figure 7 shows two examples of alignments obtained by per-

forming ICP from bad initial motion estimates and then performing diffusion at the end (Trailing diffusion) and alternating between 10 steps of ICP and a diffusion process until convergence (Alternating diffusion). Clearly alternating pairwise registration allows to avoid local minima in these examples without incurring in any noticeable penalty in running times.

5. Conclusions

In this paper we proposed a novel multiview registration algorithm that projects several pairwise alignments onto a common reference frame. The projection is performed by representing the motions as dual quaternions which are then diffused along the graph of adjacent (i.e., pairwise alignable) views. The approach is general allowing for any topology of the view-adjacency graph.

An extensive set of experiments have shown that the proposed approach is both orders of magnitude faster than the state of the art, and more robust to extreme positional noise and outliers. Finally, the dramatic speedup of the approach allows it to be alternated with the pairwise alignment process resulting in a "global" ICP that exhibits a smoother energy profile, reducing the risk of getting stuck at local minima.

Acknowledgments

262 sec.	233 sec.	115 sec.	117 sec.
Trailing diffusion	Alternating diffusion	Trailing diffusion	Alternating diffusion
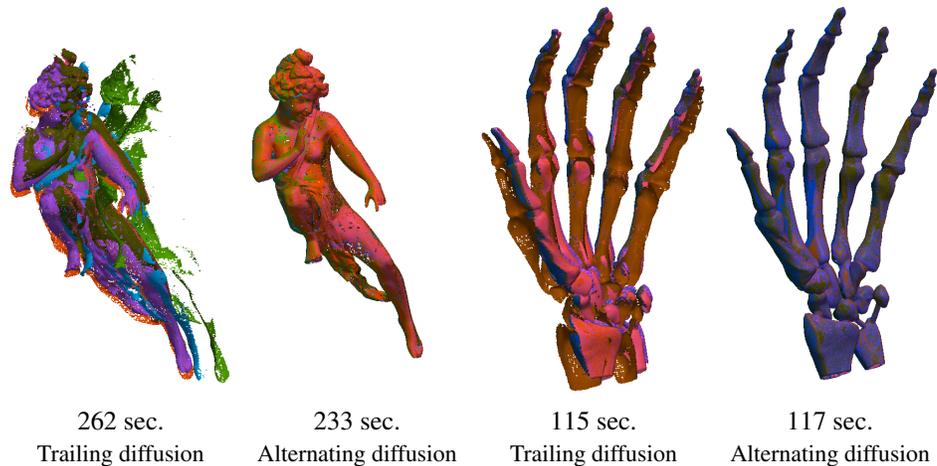
Figure 7. Alignments obtained with the trailing and alternating approaches, and respective timings.

References

[1] M. Agrawal. A lie algebraic approach for consistent pose registration for motion estimation. In *Proc. Int. Robotics Symposium*, 2006. 2442

[2] R. Benjemaa and F Schmitt. Fast global registration of 3d sampled surfaces using a multi-z-buffer technique. In *Proc. Int. Conf. on Recent Advances in 3-D Digital Imaging and Modeling*, pages 113–120, 1997. 2441

[3] R. Bergevin et al. Towards a general multi-view registration technique. *IEEE Trans. Patt. Anal. Machine Intell.*, 18(5):540–547, 1996. 2441

[4] P. J. Besl and N. D. McKay. A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(2):239–256, 1992. 2441

[5] S. R. Buss and J. P. Fillmore. Spherical averages and applications to spherical splines and interpolation. In *ACM Transactions on Graphics*, volume 20, pages 95–126, 2001. 2442

[6] Y. Chen and G. Medioni. Object modeling by registration of multiple range images. *Image and Vision Computing*, pages 145–155, 1992. 2441

[7] W. Clifford. *Mathematical Papers*. Macmillan, London, 1882. 2442

[8] K. Daniilidis. Hand-eye calibration using dual quaternions. *Int. J. of Robotics Research*, 18:286–298, 1999. 2442

[9] D. W. Eggert, A. W. Fitzgibbon, and Fisher R. B. Simultaneous registration of multiple range views for use in reverse engineering. Technical Report 804, Dept. of Artificial Intelligence, University of Edinburgh, 1996. 2441, 2442

[10] S. Granger, X. Pennec, and A. Roche. Rigid point-surface registration using an em variant of icp for computer guided oral implantology. In *MICCAI*, pages 752–761, London, UK, 2001. Springer-Verlag. 2441

[11] H. Jin, T. Duchamp, H. Hoppe, J. A. McDonald, K. Pulli, and W. Stuetzle. Surface reconstruction from misregistered data. In *Proc. SPIE vol. 2573: Vision Geometry IV*, pages 32–328, 1995. 2442

[12] L. Kavan, S. Collins, J. J. Žára, and C. O'Sullivan. Geometric skinning with approximate dual quaternion blending. *ACM Trans. Graph.*, 27(4):105, 2008. 2442, 2443

[13] L. Kavan, S. Collins, C. O'Sullivan, and J. Žára. Dual quaternions for rigid transformation blending. Technical Report TCD-CS-2006-46, Trinity College Dublin, 2006. 2443

[14] L. Kavan and J. Žára. Spherical blend skinning: a real-time deformation of articulated models. In *Proc. 2005 Symp. on Interactive 3D Graph. and Games*, pages 9–16, 2005. 2442

[15] Y. Liu. Replicator dynamics in the iterative process for accurate range image matching. *Int. J. Comput. Vision*, 83(1):30–56, 2009. 2441

[16] T. Masuda. Registration and integration of multiple range images by matching signed distance fields for object shape modeling. *Comput. Vis. Image Underst.*, 87(1-3):51–65, 2002. 2441

[17] J. M. McCarthy. *An Introduction to Theoretical Kinematics*. MIT Press, 1990. 2442

[18] K. Pulli. Multiview registration for large data sets. *Int. conf. on 3D Digital Imaging and Modeling*, pages 160–168, 1999. 2442, 2443

[19] S. Rusinkiewicz and M. Levoy. Efficient variants of the icp algorithm. In *Proceedings of the Third Intl. Conf. on 3D Digital Imaging and Modeling*, pages 145–152, 2001. 2441

[20] K. Shoemake. Animating rotation with quaternion curves. In *Proc. of the 12th annual conference on Computer graphics and interactive techniques*, pages 245–254, 1985. 2442

[21] A. Torsello. Point invariance of the screw tension minimizer. Technical Report DAIS-2011-2 http://www.dsi.unive.it/~atorsell/DAIS-2011-2.pdf, DAIS, Università Ca' Foscari Venezia, 2011. 2444

[22] J. Williams and M. Bennamoun. Simultaneous registration of multiple corresponding point sets. *Comput. Vis. Image Underst.*, 81(1):117–142, 2001. 2442

[23] Z. Zhang. Iterative point matching for registration of free-form curves and surfaces. *Int. J. Comput. Vision*, 13(2):119–152, 1994. 2441