Computer Aided Geometric Design 35–36 (2015) 16–26

Contents lists available at ScienceDirect

Computer Aided Geometric Design

www.elsevier.com/locate/cagd

Efficient Sparse ICP

Pavlos Mavridis *, Anthousis Andreadis, Georgios Papaioannou

Department of Informatics, Athens University of Economics and Business, 76 Patission Str., 10434, Athens, Greece

ARTICLE INFO

Article history:
Available online 7 May 2015

Keywords:
Surface registration
3D scan alignment
Digital geometry processing
ℓ_p minimization

ABSTRACT

The registration of two geometric surfaces is typically addressed using variants of the Iterative Closest Point (ICP) algorithm. The Sparse ICP method formulates the problem using sparsity-inducing norms, significantly improving the resilience of the registration process to large amounts of noise and outliers, but introduces a significant performance degradation. In this paper we first identify the reasons for this performance degradation and propose a hybrid optimization system that combines a Simulated Annealing search along with the standard Sparse ICP, in order to solve the underlying optimization problem more efficiently. We also provide several insights on how to further improve the overall efficiency by using a combination of approximate distance queries, parallel execution and uniform subsampling. The resulting method provides cumulative performance gain of more than one order of magnitude, as demonstrated through the registration of partially overlapping scans with various degrees of noise and outliers.

1. Introduction

Geometric registration methods search for the transformation that optimally aligns two or more surfaces. In this paper, we focus on the *pairwise* registration of a *source* surface to a *target* one, using a *rigid* transformation that consists of a rotation and a translation.

Solving this registration problem is essential to many applications: the limited range of 3D object acquisition devices in conjunction with the self-occlusion of physical objects, requires the reconstruction of digital surfaces from individual overlapping partial scans, initially available at different coordinate systems. Furthermore, in quality inspection of manufactured objects, a cloud of measured points must be aligned with the reference CAD model from which the physical object has been manufactured. With this procedure, manufacturing errors can be detected, measured and visualized. Registration algorithms can also be used for the reassembly of fractured objects. In this case, two matching fragments often share a common *contact surface*. Therefore, finding the transformation that aligns these fragments can be seen as a registration task that optimally aligns the contact surfaces (Papaioannou et al., 2001; Huang et al., 2006). In robotics, registration algorithms can be used to determine the exact location and orientation of a mobile autonomous agent in the environment (Segal et al., 2009). Finally, rigid registration methods are often used as an initialization step for non-rigid registration techniques (Pauly et al., 2005; Cheng et al., 2010) and can be also used to detect symmetries in geometric objects (Gelfand et al., 2005; Jiang et al., 2013).

Pairwise rigid registration is often handled with the well-known *Iterative Closest Point* (ICP) algorithm (Besl and McKay, 1992; Chen and Medioni, 1992) or one of its variants. This algorithm effectively performs a *local* search for the

* Corresponding author.
E-mail addresses: pmavridis@gmail.com (P. Mavridis), anthousis@aueb.gr (A. Andreadis), gepap@aueb.gr (G. Papaioannou).

http://dx.doi.org/10.1016/j.cagd.2015.03.022
0167-8396/

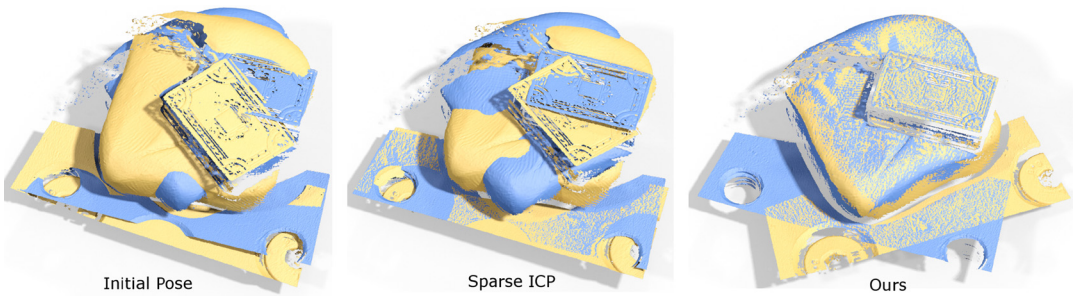

Fig. 1. Rigid registration of two partially overlapping scans of a cultural heritage object. In this challenging registration problem, outliers inherently occur due to the partial overlap of the scans, especially at the edges of the scanner's field of view. Furthermore, reflected light on polished surfaces is misread by the sensor, introducing additional outliers. Left: The input scans in their initial pose. Middle: Pairwise registration using Sparse ICP. The optimizer in this case is trapped in a local minimum, failing to align the book at the middle of the scene. Right: Pairwise registration using our method provides the desired alignment. The 3D dataset was provided by Breuckmann GmbH.

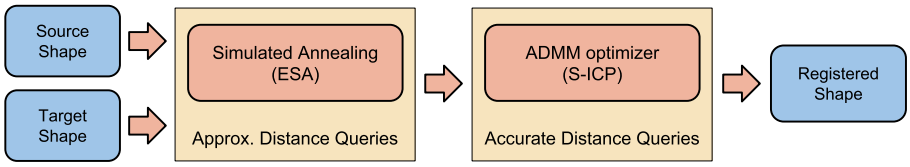

Fig. 2. Overview of our pairwise rigid registration method. We combine Simulated Annealing that can quickly approach near the solution, with an ADMM optimizer that is used to guarantee convergence to an optimal solution. To further improve the efficiency, our hybrid coarse-to-fine approach uses approximate distance queries at the first stage.

optimal alignment, and is often combined with *global* registration methods (Gelfand et al., 2005; Aiger et al., 2008; Mellado et al., 2014), which provide a *coarse* approximation of the globally optimal solution. The problem becomes significantly more challenging, when noise and outliers are introduced to the input data set, due to the partial overlap of the registered surfaces, missing regions of data due to self-occlusion or imprecise measurements, when low-accuracy 3D acquisition methods are used. Furthermore, even in the case of high-accuracy measurements, outliers occur when scanning reflective surfaces, such as metals, or at the edges of the scanner's field of view, as shown in Fig. 1, creating the need for robust alignment algorithms.

Bouaziz et al. (2013) demonstrated that the registration problem can be formulated as an ℓ_p-norm minimization of a vector of residual errors, for $p \in [0, 1]$, and proposed the *Sparse ICP* method to find a *locally* optimal solution to this problem. The parameter p controls the influence of outliers to the solution, and as it tends towards zero, the robustness of the method increases. However, at the same time the efficiency of the method is negatively affected, because the *Alternating Direction Method of Multipliers* (ADMM) optimizer, which is used to solve the underlying optimization problem, is forced to take smaller steps in the parameter space and the convergence rate of the method is significantly reduced. The design of a more efficient optimizer for this non-convex problem is the main contribution of our work.

1.1. Overview and contributions

In this paper, after identifying the reasons for the low efficiency of the original Sparse ICP approach (Section 4), we propose a registration pipeline that, similarly to Sparse ICP, formulates the problem using robust sparsity-inducing norms and significantly improves the convergence rate of the method by using a more efficient hybrid optimization strategy (Section 5). In particular, our method combines a randomized Simulated Annealing search, which allows the optimizer to take large steps in the parameter space and quickly approach the solution, with an ADMM-based optimizer that guarantees the convergence to the final optimal solution. Furthermore, we provide several insights on how to further increase the efficiency of our approach by using a combination of approximate distance queries, subsampling and parallel execution (Section 6). Fig. 2 provides an overview of the proposed alignment scheme.

We evaluate our method on the registration of partially overlapping scans with various degrees of noise and outliers and we demonstrate a performance improvement of more than one order of magnitude over the original Sparse ICP method for sufficiently low values of the parameter p. Furthermore, our experiments indicate a significant improvement in the robustness of the method to noise and outliers.

2. Related work

Since surface registration is fundamental in many domains, various methods have been proposed in the bibliography to address this problem. In this section we review the ones that are mostly related to our work. For a more comprehensive

overview of both rigid and non-rigid registration methods, the interested reader is referred to a recent survey by Tam et al. (2013).

Registration methods can be roughly classified into two general categories, voting schemes and correspondence determination methods. Voting methods take advantage of the low dimensionality (6DoF) of the rigid registration problem and directly search the global parameter space for the optimal transformation. This transformation can be defined by two sets of matching points with cardinality of at least three. Then a voting scheme is used to find the sets of points that define the optimal alignment, transforming the continuous surface alignment problem into a discrete one. Popular methods include the *Generalized Hough Transform* (Hecker and Bolle, 1994), *Geometric Hashing* (Wolfson and Rigoutsos, 1997) and several RANSAC-based methods (Irani and Raghavan, 1999; Chen et al., 1999; Aiger et al., 2008; Mellado et al., 2014). The methods in this category search for a globally optimal alignment, but due to the discretization of the search space, the computed solution is often coarse and requires a local refinement in order to provide a highly accurate alignment.

The second category of registration methods is based on the observation that computing the optimal alignment between the source and the target surface is equivalent to finding, for each point on the source surface, the corresponding point at the target one. Given this set of corresponding points, closed-form solutions exist to compute the optimal alignment (Umeyama, 1991; Eggert et al., 1997). The *Iterative Closest Point* (ICP) algorithm (Besl and McKay, 1992; Chen and Medioni, 1992) assumes a closest-point correspondence between the points of the source and the target surface. The algorithm iteratively alternates between a first step, consisting of closest-point calculations that establish correspondences, and a second step that computes and applies the rigid transformation that optimally aligns the corresponding points that were selected in the first step. However, the method is effective only when the initial pose of the input shapes is close to the globally optimal alignment, as the method converges to a local minimum, which is not necessarily the global one.

Many variants of this algorithm have been proposed in the bibliography. Rusinkiewicz and Levoy (2001) present an excellent survey of ICP variants with comparisons. In the remainder of this section we present some later advancements in the field that are related to our work. Notably, Mitra et al. (2004) use quadratic approximations of the squared distance function that are precomputed and stored in a hierarchical data structure. GoICP (Yang et al., 2013) combines the local search of ICP with a branch-and-bound search of the global parameter space to guarantee a globally optimal alignment under the l_2-norm, regardless of the initial pose of the two models. However the exhaustive nature of the search significantly increases the computation time and makes the method impractical for large datasets. All these ICP variants are based on the squared Euclidean norm, which is highly sensitive to outliers. Re-weighting (Trucco et al., 1999; Fitzgibbon, 2003), pruning or trimming (Chetverikov et al., 2002) heuristics can be used to provide a reliable registration for data with outliers, but these heuristics or weighting functions are often difficult to tune, parameters often depend on the dataset and the specific application. Furthermore, pruning heuristics can also drastically increase the number of unwanted local minima in the solution space, as demonstrated by Bouaziz et al. (2013).

Instead of relying on the ℓ_2-norm and several heuristics to improve the robustness, it is more preferable to use a norm that is known to be inherently more robust, such as the ℓ_1-norm (Flöry and Hofer, 2010) or the general ℓ_p-norm (Bouaziz et al., 2013), as mentioned in the introduction. Both of these variants significantly improve the resilience of the ICP method to large amounts of noise and outliers, without using any difficult-to-tune reweighting heuristics and without introducing unwanted local minima in the solution space, as is the case with pruning and trimming heuristics.

The robustness of these methods increases as the parameter p of the ℓ_p-norm approaches zero. However, at the same time the problem becomes increasingly non-smooth and non-convex, making the design of an efficient optimizer difficult and negatively affects the performance of the Sparse ICP method. In this paper we propose a more efficient optimization approach for this problem.

Aside from the ICP method, general optimization algorithms, such as *Levenberg–Marquardt (LM-ICP)* (Fitzgibbon, 2003), *Simulated Annealing* (Blais and Levine, 1995) and *Particle Swarm Optimization* (Oikonomidis et al., 2012), have been used to solve alignment and tracking problems. However, unlike our approach, these techniques do not use an inherently robust ℓ_p-norm formulation, making the optimization sensitive to noise and outliers or requiring the use of difficult to tune robustness weights.

3. Mathematical formulation

Most registration algorithms try to minimize a distance metric between two surfaces $\mathcal{X}$, $\mathcal{Y}$. This problem can be formulated as

$$\arg\min_{\mathbf{R},\mathbf{t}} \int_{\mathbf{x}\in\mathcal{X}} \phi(\mathbf{R}\mathbf{x}+\mathbf{t},\mathcal{Y}) + I_{SO(3)}(\mathbf{R}), \tag{1}$$

where $\mathbf{R}\in\mathbb{R}^{3\times3}$ denotes the rotation matrix and $\mathbf{t}\in\mathbb{R}^3$ the translation vector of the alignment transformation. The matrix $\mathbf{R}$ is restricted to the special orthogonal group $SO(3)$ using the indicator function $I_{SO(3)}(\mathbf{X})$, which evaluates to zero when $\mathbf{X}\in SO(3)$ and to $+\infty$ otherwise, enforcing the rigidity of the transformation. The function $\phi(\mathbf{x},\mathcal{Y})$ measures the distance of an arbitrary point $\mathbf{x}\in\mathbb{R}^3$ to the surface $\mathcal{Y}$ and is defined as

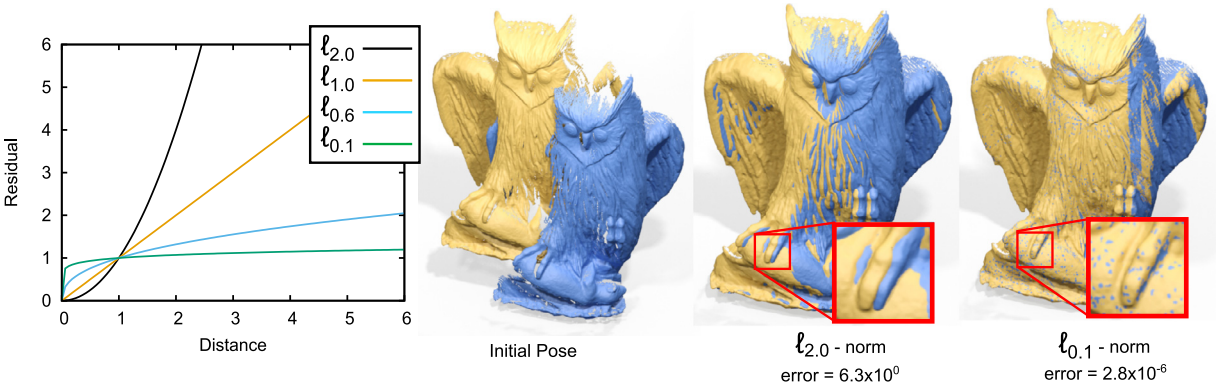

Fig. 3. Left: ℓ_p-norms reduce the penalty associated with large distances (outliers). Right: The accuracy of our registration procedure for different ℓ_p-norms on a synthetic registration problem with a well-defined ground truth solution. When the ℓ_2-norm is used, the optimizer will skew the solution in order to avoid the large penalty associated with outliers, as shown in the blown-up area. ℓ_p-norms avoid this problem for sufficiently small values of the parameter p $(0 \leq p \leq 1)$. For $p = 0.1$ the overlapping parts of the two models match so closely that we observe z-fighting during the visualization. The error is measured as the root mean square distance in Euclidean space from the ground truth alignment. The error for the $\ell_{1.0}$ and $\ell_{0.6}$-norms is 3.2×10^{-3} and 1.3×10^{-4} respectively. The Owl model is from Bouaziz et al. (2013).

$$\phi(\mathbf{x}, \mathcal{Y}) = \min_{\mathbf{y} \in \mathcal{Y}} \phi(\mathbf{x}, \mathbf{y}), \tag{2}$$

where the metric $\phi(\mathbf{x}, \mathbf{y})$ measures the distance between two points in space. This equation is often referred to as the *distance transform* or the *distance field* of surface $\mathcal{Y}$.

In most applications we often deal with discretized surfaces, represented by a set of points $\mathbf{x}_i \in \mathcal{X}$, $i = 1 \ldots n$. In this case, we can define the vector of *residual distances* $\mathbf{z} \in \mathbb{R}^n$, whose elements are defined as:

$$z_i(\mathbf{R}, \mathbf{t}) = \phi(\mathbf{R}\mathbf{x}_i + \mathbf{t}, \mathcal{Y}) + I_{SO(3)}(\mathbf{R}), \ i = 1, \ldots, n. \tag{3}$$

Using this definition, Eq. (1) can be rewritten as:

$$\arg\min_{\mathbf{R},\mathbf{t}} \sum_{i=1}^{n} z_i(\mathbf{R}, \mathbf{t}). \tag{4}$$

Many methods use the squared Euclidean norm as the distance metric and optimize Eq. (4) in a least squares manner. However, in the presence of outliers, an optimizer will skew the solution in order to reduce the large penalty associated with distant points. To avoid this problem, similar to recent work in the field (Bouaziz et al., 2013), we define the distance metric in Eq. (2) as:

$$\phi(\mathbf{x}, \mathbf{y}) = \mu_p(\|\mathbf{x} - \mathbf{y}\|_2), \quad \mu_p(x) = |x|^p. \tag{5}$$

This sparsity inducing ℓ_p-norm formulation of the problem, for $p \in [0, 1]$, imposes a lower penalty to large distances than other norms, making the optimizer resilient to outliers, as shown in Fig. 3. In practice, this metric aims to maximize the number of zero distances between the corresponding points, thus maximizing the contact area of the two surfaces. Please note that the distances z_i in the residual vector $\mathbf{z}$ are computed using the Euclidean norm. The ℓ_p-norm is applied on vector $\mathbf{z}$ in order to compute a single scalar residual value from these distances.

Eq. (5) defines the well-known *point-to-point* distance metric in ℓ_p space. However, this metric is not very accurate in the surface's near-field, as it corresponds to a zeroth-order Taylor expansion of the surfaces distance-field. We can obtain a better approximation using a first-order Taylor expansion. In this case, we define the metric $\phi(\mathbf{x}, \mathbf{y}) = \mu_p(\mathbf{n}^T(\mathbf{x} - \mathbf{y}))$, where $\mathbf{n}$ is the normal of the point y. This equation defines the well-known *point-to-plane* distance metric in ℓ_p space.

4. Sparse ICP convergence properties

One method to solve the optimization problem that we have described in the previous section is to use the Sparse ICP algorithm. At each iteration of ICP, the total energy (residual distance) weakly decreases, thus the algorithm provides a guaranteed convergence to a local minima (Besl and McKay, 1992). However, the actual convergence rate depends of the norm used during the optimization. As shown in Fig. 3, problems with many outliers require the use of ℓ_p norms with a sufficiently low value of p. In these cases, since the penalty induced by the outliers becomes smaller, the slope of the objective function is reduced and the optimizer is forced to perform smaller steps. This is clearly shown in Fig. 5, where the convergence rate of both Sparse ICP and our method is shown. In the case of Sparse ICP, we observe that decreasing the value of p significantly reduces the convergence rate of the method.

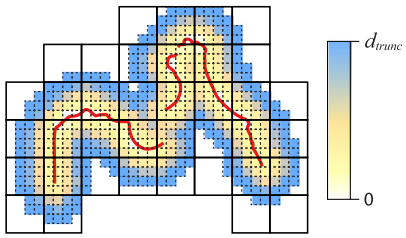

Fig. 4. The distance field of the target surface (shown as a red line) is discretized over a 3D grid, truncated at a distance d_{trunc} and stored using the hierarchical VDB data structure. The color variations represent the distance from the surface. For clarity only the last two levels of the hierarchy are shown. The target surface can be either a triangle mesh or a point cloud. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

5. Hybrid optimizer

After the analysis of the previous section, it is clear that in order to improve the overall efficiency of the method, we should aim to improve the convergence rate for low values of p. To this end, we propose a two-level hybrid optimizer, as shown in Fig. 2. The optimization begins with a general *Simulated Annealing* search that is able to perform large random jumps in the parameter space and thus can approach faster to the solution, but with low accuracy. When the optimizer is close enough to the solution, we switch to the ADMM-based optimizer of the Sparse ICP, in order to guarantee a convergence to the optimal solution.

For the Simulated Annealing process, the alignment problem is parameterized with three continuous variables for the translation and three continuous variables (Euler angles) for the rotation. Many simulated annealing variants require that the range of these continuous variables to be bounded. To this end, two of the Euler angles are bounded in the $[-\pi, \pi]$ range and the third one in the $[-\pi/2, \pi/2]$. The translation range is bounded by the extends of the target point cloud.

Our implementation is based on the *Enhanced Simulated Annealing* (Siarry et al., 1997) method. In this Simulated Annealing variant, for a problem with n continuous variables, the original Metropolis iterative random search, which takes place in $\mathbb{R}^n$ space, is replaced by another similar exploration, performed within a succession of $\mathbb{R}^p$ spaces, with $p \ll n$. In our experiments, we have found that using single dimensional sub-spaces ($p = 1$) provides faster and more reliable convergence for our problem. This configuration also outperformed a number of other Simulated Annealing variants that we have tested, but not necessarily with a large margin.

The annealing process at each iteration generates a candidate transform and then directly evaluates the objective function (Eq. (1)). Unlike ICP variants, this process does not involve any closest-point correspondence determination. After we switch to the ADMM-based optimization, we alternate between a closest-point correspondence determination step and a rigid motion determination step, similar to the original Sparse ICP method.

It is worth noting that the annealing is a stochastic process, and sometimes it gets trapped to local minima. When no improvement is detected after several iterations, we restart the stochastic search with a higher temperature, which allows larger jumps, in order to escape from the local minima. The maximum number of restarts is limited by a user parameter, which was fixed to five in our experiments.

6. Additional efficiency improvements

The overall efficiency of the proposed method can be further increased by reducing the cost of each Simulated Annealing iteration. This cost is governed by the evaluation of the objective function. To reduce this cost, we use a combination of approximate distance queries, parallel evaluation and uniform subsampling of the source mesh.

6.1. Approximate distance queries

The evaluation of the objective function requires multiple closest point queries (Eq. (2)). To minimize the cost associated with this operation, we use a data structure that allows fast approximate distance queries. To this end, we use a discretized version of the target surface's distance field. In particular, the distance function is discretely sampled on a 3D grid that extends over the narrow band of the target surface $\mathcal{Y}$ (see Fig. 4) and is stored using a sparse hierarchical volumetric data structure, the VDB (Museth, 2013). This *truncated distance field* directly encodes the approximate distance of a point in space $\mathbf{x}$ to the closest point on a surface $\mathcal{Y}$. For points beyond the narrow band of the surface, the maximum encoded distance d_{trunc} (truncation threshold) is returned. This truncation can be beneficial, since it reduces the influence of very distant outliers, but most importantly, also reduces the memory and processing requirements for the distance field data structure. After creating the discretized distance field, subsequent distance queries involve fetching a precomputed value from the data structure, thus the efficiency is increased. The initial precomputation cost for the discretization of distance field, which is higher than typical alternative approaches, can be amortized over multiple alignment queries in problems that involve

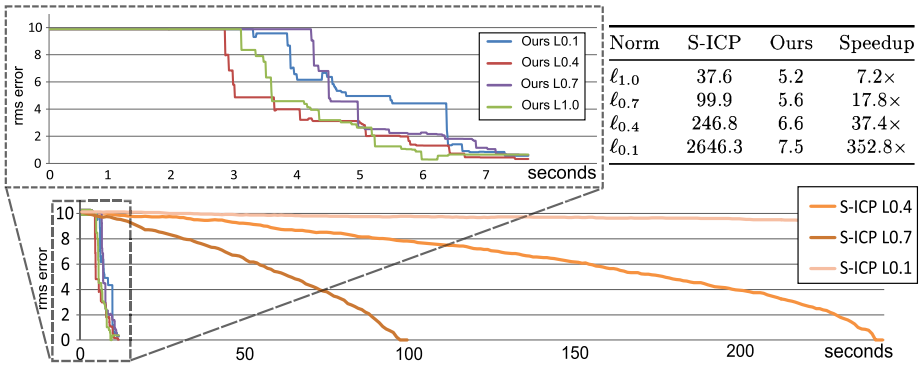

Norm	S-ICP	Ours	Speedup
$\ell_{1.0}$	37.6	5.2	7.2×
$\ell_{0.7}$	99.9	5.6	17.8×
$\ell_{0.4}$	246.8	6.6	37.4×
$\ell_{0.1}$	2646.3	7.5	352.8×

Fig. 5. The table on the upper-right shows the total computation time in seconds for our registration approach compared to Sparse ICP (S-ICP) for similar accuracy, for various ℓ_p-norms. The lower diagram shows the convergence rate of the algorithms. Since the convergence of S-ICP is much slower than our method, the upper diagram shows a zoomed-in area of the lower one, focusing only on the first seconds of the optimization process. The RMS error is measured in Euclidean space for all the norms against a ground truth solution. The error in some iterations increases, since the optimizers minimize the error in ℓ_p space, but here we measure it in Euclidean space, in order to perform meaningful comparisons. In contrast to Sparce ICP, the convergence rate of our method is widely unaffected by the value of the parameter p. Registration problems that involve partial overlap or large amounts of outliers require sufficiently small values of p, making our approach preferable in these cases. The preprocessing time for VDB or ANN data structures can be seen in Table 1 and is not included in these measurements.

multiple surfaces. When this is not the case, our system uses a data structure for *Approximate Nearest Neighbor* (ANN) queries (Arya et al., 1998).

It is worth noting that both the ANN and VDB structures encode distances in the Euclidean space. Our method works in ℓ_p space, but since the $\mu_p(x)$ function in Eq. (5) is non-decreasing and both $\mu(\|.\|_2)$ and $\|.\|_2$ achieve a minimum value at the same points, we can use existing well-known data structures and algorithms for the computation of distance fields in the Euclidean space, and then transform the values to ℓ_p space, by applying the function $\mu_p(x)$. These approximate queries are only used during the Simulated Annealing search. The ADMM-based ICP optimizer always uses accurate queries, to avoid compromising the quality of the final registration.

6.2. Parallel evaluation

From mobile devices to desktop workstations, multicore processors are omnipresent in the modern computing landscape. For this reason, algorithms should be designed to be parallelizable, in order to take advantage of all the available computational resources. In our method, the evaluation of the objective function is trivially parallelizable over N threads, by dividing the sum in Eq. (1) into N equal-sized workloads that are computed independently by each thread. Then, a parallel reduction combines the results. It is worth noting that the ADMM optimizer in Sparse ICP is based on a series of linear algebra operations that are more difficult to parallelize efficiently. The simulated annealing approach is better suited for parallel computation.

An alternative approach would be to divide the global parameter space to N equal sized blocks and perform N independent annealing searches in parallel. However, this approach is suboptimal, since it will waste a lot of computational effort searching on blocks of the parameter space that do not contain an optimal solution. In contrast, our approach can quickly focus on the parts of the parameter space that are near to a solution.

6.3. Uniform subsampling

Digitally acquired surfaces typically consist of very dense point clouds. Large efficiency improvements can be obtained by using only a uniformly sampled subset of these points during the evaluation of the objective function. The final ADMM-based optimization uses the complete set of points, to avoid compromising the quality of the final registration. In all cases, only the source point cloud is subsampled and not the target one.

7. Evaluation

In this section we evaluate the robustness and the efficiency of our method. The number of possible experiments that can be performed is very large. Here, we focus on datasets with partial overlap and large amounts of noise and outliers, since these are the most challenging alignment problems, and we limit our analysis on a set of experiments that better demonstrate the unique properties of our approach.

Throughout the paper, unless otherwise noted, we have employed the $\ell_{0.4}$ norm and the point-to-plane distance metric is used when normals are available or can be easily computed. On noisy points clouds (Fig. 6), where normals cannot be reliably estimated, the point-to-point metric was used. Unless otherwise specified, our algorithm uses the VDB data structure with a voxel size of three units. This results in hierarchical volumes that consume roughly from 30 MB to 90 MBytes

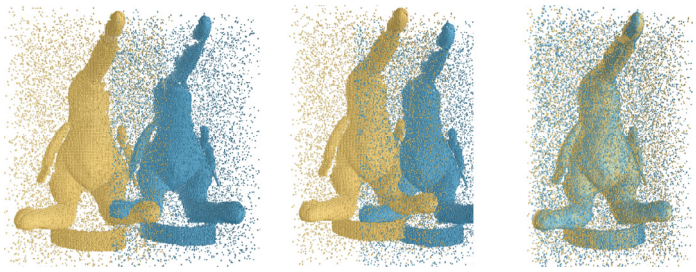

Fig. 6. Registration of two point clouds with a large number of synthetically generated outliers. Both the source and the target point clouds are contaminated with outliers. Left: Initial position of the two surfaces. Middle: The original Sparse ICP method is known to be robust only to source outliers and in the presence of destination outliers the optimizer gets stuck to a suboptimal local minimum solution. Right: Our approach successfully aligns the two datasets, avoiding undesired local minima. The Coati model is from Weise et al. (2007).

of memory, depending on the physical size of the object. All experiments were performed on an *Intel Core i7-3820* CPU at 3.6 GHz with 4 cores and when possible, all algorithms use 8 threads to take advantage of hyper-threading. For our comparisons with the original Sparse ICP we have used the implementation from Bouaziz et al. (2013). For the comparison with Super4PCS we have used the implementation from Mellado et al. (2014).

Our method is designed to work on general unstructured point clouds, without requiring the existence of any triangle connectivity information. However, to better visualize the resulting alignments, we triangulate the test datasets after performing the registration. On datasets with large amounts of outliers, such as the one in Fig. 6, we directly visualize the point cloud, since the computation of triangle connectivity is unreliable and requires further filtering algorithms.

Fig. 5 compares our method against Sparse ICP in terms of total computation time and convergence rate for various values of the parameter p. This parameter controls the tradeoff between robustness and performance. Problems with large amounts of outliers require sufficiently small values of p and in these cases, our method significantly outperforms Sparse ICP. In fact, the performance advantage gets larger as the parameter p gets lower and for $p = 0.1$, our method is more than two orders of magnitude faster. These measurements indicate that our approach enables the use of very low values for the parameter p, something that was not practical with previous methods. Therefore, the overall robustness of the registration is increased, while the computational cost is reduced.

For small values of the parameter p, the alignment problem becomes increasingly non-convex and non-smooth, negatively affecting the performance of the ADMM optimizer in Sparse ICP. In fact, any gradient-based optimizer would be forced to make smaller steps in the parameter space, as the slope of the objective function will decrease, negatively affecting the convergence rate. For this reason, in our approach we have opted for an optimization method that is not based on the gradient of the objective function, and thus, the convergence rate of the method is not significantly affected by the value of the parameter p, as the measurements in Fig. 5 indicate.

In Fig. 3 we evaluate the accuracy of the alignment with relation to the parameter p, when aligning two synthetically generated, partially overlapping scans of an Owl model. For this experiment, we have used synthetic scans from a previously digitized model, since in real-world alignment scenarios, it is often difficult or impossible to properly define a ground truth solution. As expected, the accuracy of the registration, measured as the Root Mean Square Error with relation to the ground truth, increases as the parameter p decreases. The difference with the Sparse ICP is that in our case, due to the improved efficiency, the use of small values for p is now viable. It is worth noting that the parameter p does not require any tuning, a value of p as close to zero as possible should be used, depending on the performance requirements of an application. The RMS error is always measured and reported in the Euclidean space, although the residual distance during the optimization is always measured in ℓ_p space.

Fig. 6 demonstrates the alignment of two data sets with large amounts of synthetically generated outliers. In this experiment, the outliers exist in both the source and the target data set. When outliers exist in the target data set, the optimization in Sparse ICP is very easily driven towards a bad local minimum, since the whole registration algorithm is based on closest-point correspondences, which are very unreliable in this case, as also noted by Bouaziz et al. (2013). In contrast, the alignment in the initial stage of our algorithm is not based on closest-point correspondences and our search strategy better explores the global parameter space, using the Simulated Annealing method, and successfully aligns the two surfaces, avoiding undesired local minima.

Our measurements in Table 1 indicate that the use of a precomputed distance field offers large efficiency gains over the more traditional kd-tree or *Approximate Nearest Neighbor* (ANN) data structures. In this test, ANN was configured to provide roughly the same accuracy as the VDB structure. As expected, the gain from the use of a more efficient data structure becomes larger as the number of source surface points increases. For the largest dataset in this test, VDB is more than 3 times faster than ANN and 30 times faster than a k-D tree. In practice, when subsampling is used, the actual speed-up is lower (close to 33%), since even for dense point clouds only a subset of uniform points is used for the objective function evaluation. Still, in all cases the VDB has a performance advantage. It is worth noting that one of the fastest available ANN implementations was used in our comparisons (nanoflann). However, the time required for the creation of the discretized

Table 1
Average alignment time when using the discretized distance field (VDB), Approximate Nearest Neighbor (ANN) and a standard k-D tree, with relation to the number of source surface points. T_p: precomputation time of the corresponding data structure for the largest dataset tested (155K target points). All times are in seconds.

Data structure	T_p	154K	77K	38K	9K
VDB	1.9	25.5	9.3	4.1	1.8
ANN	0.01	78.4	22.8	9.1	2.4
k-D tree	0.01	890.6	236.8	67.6	8.5

Table 2
The accuracy of the registration results with relation to the accuracy of the distance queries (voxel size) before (RMS_1) and after (RMS_2) the ADMM-based refinement for the dataset of Fig. 3. The final alignment remains always highly accurate, even for large voxel sizes. This is to be expected, since the final stage of the registration always uses accurate distance queries.

	3.0	2.0	1.5	1.0
RMS_1	7.8×10^{-1}	5.0×10^{-1}	6.1×10^{-1}	2.6×10^{-1}
RMS_2	3.5×10^{-5}	2.7×10^{-5}	6.2×10^{-5}	3.6×10^{-5}

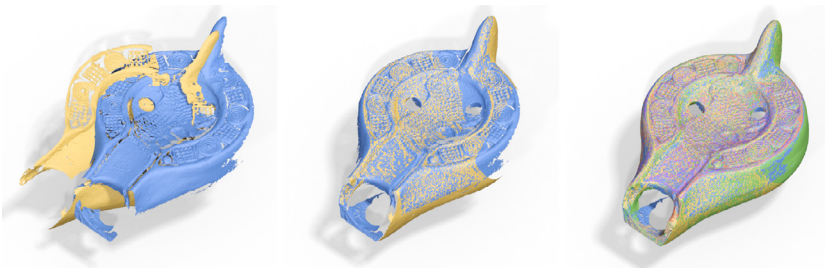

Fig. 7. Reconstruction of an ancient oil lamp from partially overlapping scans. Left: Two of the input scans in their initial pose. Middle: Registration results with our method for the two scans on the left. In this dataset, our method computes the alignment in 11 seconds, which is 31 times faster than Sparse ICP for similar accuracy on the same dataset. Right. Reconstruction of the oil lamp model from multiple scans. The 3D dataset was provided by Breuckmann GmbH.

distance field is higher than other data structures. This cost can be amortized over many alignment queries, which is typical for problems involving more than two surfaces. When this is not the case, the ANN data structure might be preferable.

Aside from the main parameter p that controls the ℓ_p-norm, our method introduces a new parameter that controls the accuracy of the distance queries in the Simulated Annealing method. In Table 2 we examine how the accuracy of the registration is affected when changing the accuracy of the distance queries during the Simulated Annealing search. For this experiment we have used the synthetic dataset of Fig. 3 and we compare the results with the ground truth alignment. Our measurements indicate that the final alignment is not significantly influenced by the accuracy of the distance queries. This is to be expected, as the final ADMM-based S-ICP optimization is always performed using accurate distance queries. Therefore, there is no need to tune this additional parameter and it is reasonable to keep it constant in all experiments. The measurements in Table 2 were performed with the VDB data structure, while measurements with the ANN data structure lead to exactly the same conclusions.

Figs. 1 and 7 demonstrate the reconstruction of two cultural heritage objects from partially overlapping scans. In these two cases, the initial coarse alignment is computed with a few iterations of the RANSAC-based 4-PCS method (Aiger et al., 2008). The partial scans in these examples are densely sampled and consist of 450K to 808K points. The dataset of Fig. 1 poses several challenges to an alignment algorithm. First, the metallic chain (beside the book) reflects the incoming light from the scanner and creates a number of outliers. Furthermore, the background of this scene was not removed and many outliers are inherently created at the edges of the scanner's field of view. Even in this challenging case, our algorithm provides a very precise alignment. In contrast, the optimizer in the original Sparse ICP method is trapped in a local minimum.

To better demonstrate our method in realistic real-world applications, Fig. 7 also includes alignment of multiple partial scans. For the multi-part registration, similar to Huber (2002), we construct a graph, where each fragment corresponds to a node, and each pairwise match to an edge between two nodes. The optimal set of pairwise connections is computed by finding the *Minimum Spanning Tree* of this graph. The multi-part results do not include any global error relaxation (bundle adjustment), in order to better demonstrate the quality of the pairwise registration. Fig. 8 demonstrates additional alignment results in challenging registration problems that involve small overlap, holes or many local minima. In all cases our method provides a highly reliable registration.

The coarse Simulated Annealing search in our method can be replaced with other global registration methods. A potential candidate is Super4PCS (Mellado et al., 2014), a RANSAC variant which is known to be optimal in terms of computational complexity for a single core. In Table 3 we compare the performance of our approach with Super4PCS for similar registration

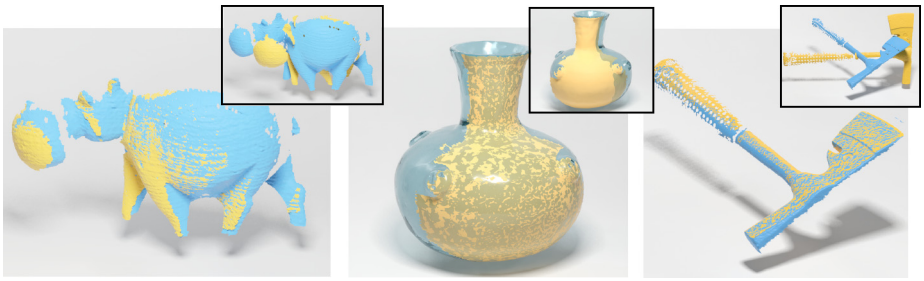

Fig. 8. Rigid registration of various datasets computed with our method. Left: Registration of two partial scans with small overlap. Dataset is taken from Mellado et al. (2014). Middle: A registration problem with many local minima, due to the rotational symmetry of the pottery object. Right: Registration of two partially overlapping scans of an axe. The insets show the initial position of the objects. The axe 3D dataset is courtesy of Breuckmann GmbH.

Table 3
Comparison of our method with Super4PCS for similar registration accuracy on various datasets shown in this paper. All times are in seconds and do not include any local ICP refinement. While Super4PCS performs a very smart exploration of the global search space, our algorithm is based on a more efficient data structure (VDB over ANN) and can take advantage of multiple processing cores (here 4), making our approach largely more efficient in practice.

	Oil lamp	Owl	Coati	Hippo	Pot	Axe
# points	367K	154K	36.7K	21K	7.9K	84.6K
Super4PCS	36.6	173.1	58.6	5.4	2.5	59.1
Ours – 1 core	18.2	9.1	22.2	10.6	4.0	35.8
Ours – 4 cores	4.7	2.6	6.0	4.5	1.3	13.9

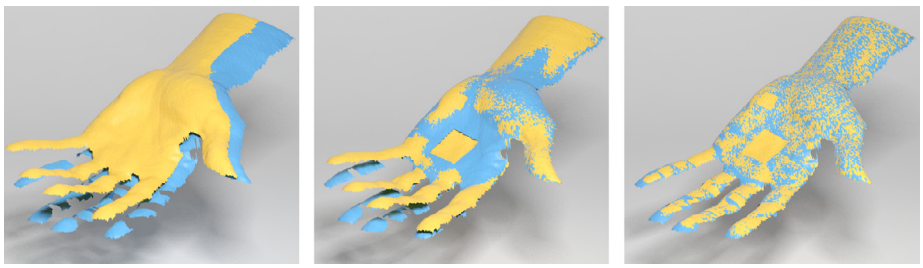

Fig. 9. A challenging alignment problem that involves a closing palm. Since the palm is deforming, rigid registration cannot fully align the two models. The result of the rigid registration, shown at the middle, is used for the initialization of a non-rigid deformation that successfully aligns the two meshes. The holes in the target point cloud, which appear as uniform yellow regions in the resulting non-rigid registration, were synthetically generated to make the alignment problem more challenging. Dataset is taken from Weise et al. (2007). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

accuracy on various datasets. Since both methods require a local ICP refinement, we only measure the global registration time in this test. The measurements indicate that our method is largely more efficient. While Super4PCS explores the global parameter space in a very efficient way, by aligning congruent sets of four points from the two point clouds, the large performance advantage of our method can be attributed to the use of a more efficient data structure for distance queries (VDB over ANN), which is very important for large datasets, as shown in Table 1. Furthermore, Super4PCS is designed for a single core, making this algorithm suboptimal for multi-core processors. In particular, the authors of this algorithm note that a parallel implementation might be possible, but it is left as a future work. In contrast, parallelization of our approach is rather trivial (Section 6.2). Finally, it is worth noting that our approach does not require an a priori estimation of the amount of overlap in the two surfaces, as required by Super4PCS or Trimmed-ICP (Chetverikov et al., 2002).

While the main focus of this work is the rigid registration problem, many practical alignment problems require non-rigid registration. This is the case when the object being scanned deforms between subsequent partial scans, such as the closing palm in Fig. 9. In such challenging problems, our rigid registration approach can be extended to include a non-rigid deformation step after the rigid alignment. In this case, the quality of the rigid registration is still very important, since it is used as the initial pose for the subsequent non-rigid deformation. To better evaluate our method in this scenario, we have extended our algorithm with the non-rigid deformation model described by Pauly et al. (2005). The deformation is based on two metrics. The first metric measures the Euclidean distance between the two models and the second one measures distortion of the deforming point cloud. The two measures are linearly combined on a single objective function with the *rigidity* parameter and this function is minimized using a least-squares optimizer. In order to improve the quality of the closest-point correspondences and to avoid unwanted local minima, we use an iterative rigidity relaxation scheme, where

the desired deformation is computed over multiple iterations, starting with a high rigidity value, which is relaxed in the subsequent iterations.

8. Conclusions and future work

We have presented an efficient variation of the Sparse ICP algorithm that is based on a new hybrid optimization which starts with a general Simulated Annealing search and then switches to ADMM-based ICP, to ensure convergence to an optimal solution. We have also provided several insights on how to further improve the efficiency of our method using a combination of approximate distance queries, parallel execution and uniform subsampling. The cumulative performance increase over the original Sparse ICP is more than one order of magnitude when tested on the registration of partially overlapping scans. At the same time, we have demonstrated that the hybrid optimization approach in our method avoids undesired local minima, increasing the robustness of the registration process in very challenging alignment problems that involve a large number of outliers and partially overlapping surfaces.

An interesting direction of research for the future is the adaptation of our hybrid optimization strategy and the underlying data structures to GPUs and similar massively parallel architectures, in order to achieve further performance improvements. Additionally, it would be interesting to explore possible extensions of our method that use salient features in order to further improve the reliability of the alignment.

Acknowledgements

The authors would like to thank the anonymous reviewers for their insightful comments and suggestions. We would also like to thank Dirk Rieke-Zapp and Michael Hermstein from Breuckmann GmbH for providing the datasets in Figs. 1 and 7 and for having several valuable discussions on the topics of digitization and 3D data acquisition. The pottery dataset in Fig. 8 was kindly provided by Ioannis Pratikakis from Athena Research Center. Finally, we would like to thank the authors of SparseICP and Super4PCS algorithms for freely providing an implementation of their methods. This work was supported by EC FP7 STREP project “PRESIOUS”, grant No. 600533.

References

Aiger, D., Mitra, N.J., Cohen-Or, D., 2008. 4-points congruent sets for robust pairwise surface registration. In: ACM SIGGRAPH 2008 Papers. SIGGRAPH '08. ACM, New York, NY, USA, pp. 85:1–85:10.

Arya, S., Mount, D.M., Netanyahu, N.S., Silverman, R., Wu, A.Y., 1998. An optimal algorithm for approximate nearest neighbor searching fixed dimensions. J. ACM 45 (6), 891–923.

Besl, P.J., McKay, N.D., 1992. Method for registration of 3-D shapes. In: Robotics-DL Tentative. International Society for Optics and Photonics, pp. 586–606.

Blais, G., Levine, M.D., 1995. Registering multiview range data to create 3D computer objects. IEEE Trans. Pattern Anal. Mach. Intell. 17 (8), 820–824.

Bouaziz, S., Tagliasacchi, A., Pauly, M., 2013. Sparse iterative closest point. In: Symposium on Geometry Processing. Comput. Graph. Forum 32 (5), 1–11.

Chen, C.-S., Hung, Y.-P., Cheng, J.-B., 1999. RANSAC-based DARCES: a new approach to fast automatic registration of partially overlapping range images. IEEE Trans. Pattern Anal. Mach. Intell. 21 (11).

Chen, Y., Medioni, G., 1992. Object modelling by registration of multiple range images. Image Vis. Comput. 10 (3), 145–155.

Cheng, Z.-Q., Jiang, W., Dang, G., Martin, R., Li, J., Li, H., Chen, Y., Wang, Y., Li, B., Xu, K., Jin, S., 2010. Non-rigid registration in 3D implicit vector space. In: Shape Modeling International Conference (SMI). 2010, pp. 37–46.

Chetverikov, D., Svirko, D., Stepanov, D., Krsek, P., 2002. The trimmed iterative closest point algorithm. In: Proceedings of the 16th International Conference on Pattern Recognition. 2002, vol. 3. IEEE, pp. 545–548.

Eggert, D.W., Lorusso, A., Fisher, R.B., 1997. Estimating 3-D rigid body transformations: a comparison of four major algorithms. Mach. Vis. Appl. 9 (5–6), 272–290.

Fitzgibbon, A.W., 2003. Robust registration of 2D and 3D point sets. Image Vis. Comput. 21 (13), 1145–1153.

Flöry, S., Hofer, M., 2010. Surface fitting and registration of point clouds using approximations of the unsigned distance function. Comput. Aided Geom. Des. 27 (1), 60–77.

Gelfand, N., Mitra, N.J., Guibas, L.J., Pottmann, H., 2005. Robust global registration. In: Proceedings of the Third Eurographics Symposium on Geometry Processing. SGP '05. Eurographics Association, Aire-la-Ville, Switzerland.

Hecker, Y.C., Bolle, R.M., 1994. On geometric hashing and the generalized hough transform. IEEE Trans. Syst. Man Cybern. 24 (9), 1328–1338.

Huang, Q.-X., Flöry, S., Gelfand, N., Hofer, M., Pottmann, H., 2006. Reassembling fractured objects by geometric matching. In: ACM SIGGRAPH 2006 Papers. SIGGRAPH '06. ACM, New York, NY, USA, pp. 569–578.

Huber, D.F., 2002. Automatic three-dimensional modeling from reality. PhD thesis. Carnegie Mellon University.

Irani, S., Raghavan, P., 1999. Combinatorial and experimental results for randomized point matching algorithms. Comput. Geom. 12 (1), 17–31.

Jiang, W., Xu, K., Cheng, Z.-Q., Zhang, H., 2013. Skeleton-based intrinsic symmetry detection on point clouds. Graph. Models 75 (4), 177–188.

Mellado, N., Aiger, D., Mitra, N.J., 2014. Super 4pcs fast global pointcloud registration via smart indexing. Comput. Graph. Forum 33 (5), 205–215.

Mitra, N.J., Gelfand, N., Pottmann, H., Guibas, L., 2004. Registration of point cloud data from a geometric optimization perspective. In: Proceedings of the 2004 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing. SGP '04. ACM, New York, NY, USA, pp. 22–31.

Museth, K., 2013. Vdb: high-resolution sparse volumes with dynamic topology. ACM Trans. Graph. 32 (3).

Oikonomidis, I., Kyriazis, N., Argyros, A.A., 2012. Tracking the articulated motion of two strongly interacting hands. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, pp. 1862–1869.

Papaioannou, G., Karabassi, E.-A., Theoharis, T., 2001. Virtual archaeologist: assembling the past. IEEE Comput. Graph. Appl. 21, 53–59.

Pauly, M., Mitra, N.J., Giesen, J., Gross, M., Guibas, L.J., 2005. Example-based 3D scan completion. In: Proceedings of the Third Eurographics Symposium on Geometry Processing. SGP '05. Eurographics Association, Aire-la-Ville, Switzerland.

Rusinkiewicz, S., Levoy, M., 2001. Efficient variants of the icp algorithm. In: Proceedings of the Third International Conference on 3-D Digital Imaging and Modeling. 2001. IEEE, pp. 145–152.

Segal, A., Haehnel, D., Thrun, S., 2009. Generalized-icp. In: Robotics: Science and Systems, vol. 2, p. 4.

Siarry, P., Berthiau, G., Durdin, F., Haussy, J., 1997. Enhanced simulated annealing for globally minimizing functions of many-continuous variables. ACM Trans. Math. Softw. 23 (2), 209–228.

Tam, G., Cheng, Z.-Q., Lai, Y.-K., Langbein, F., Liu, Y., Marshall, D., Martin, R., Sun, X.-F., Rosin, P., 2013. Registration of 3D point clouds and meshes: a survey from rigid to nonrigid. IEEE Trans. Vis. Comput. Graph. 19 (7), 1199–1217.

Trucco, E., Fusiello, A., Roberto, V., 1999. Robust motion and correspondence of noisy 3-D point sets with missing data. Pattern Recognit. Lett. 20 (9), 889–898.

Umeyama, S., 1991. Least-squares estimation of transformation parameters between two point patterns. IEEE Trans. Pattern Anal. Mach. Intell. 13 (4), 376–380.

Weise, T., Leibe, B., Van Gool, L., 2007. Fast 3D scanning with automatic motion compensation. In: IEEE Conference on Computer Vision and Pattern Recognition, 2007. CVPR '07, pp. 1–8.

Wolfson, H.J., Rigoutsos, I., 1997. Geometric hashing: an overview. Comput. Sci. Eng. 4 (4), 10–21.

Yang, J., Li, H., Jia, Y., 2013. Go-ICP: solving 3D registration efficiently and globally optimally. In: Proceedings of the 14th International Conference on Computer Vision (ICCV).