

Global Registration of Large Collections of Range Images with an Improved Optimization-on-a-Manifold approach

Francesco Bonarrigo, Alberto Signoroni*

Information Engineering Dept., DII

University of Brescia, Brescia, Italy

**Email: alberto.signoroni@ing.unibs.it*

Abstract

Concurrently obtaining an accurate, robust and fast global registration of multiple 3D scans is still an open issue for modern 3D modeling pipelines, especially when high metric precision as well as easy usage of high-end devices (structured-light or laser scanners) are required. Various solutions have been proposed (either heuristic, iterative and/or closed form solutions) which present some compromise concerning the fulfillment of the above contrasting requirements. Our purpose here, compared to existing reference solutions, is to go a step further in this perspective by presenting a new technique able to provide improved alignment performance, even on large datasets (both in terms of number of views and/or point density) of range images. Relying on the ‘Optimization-on-a-Manifold’ (OOM) approach, originally proposed by Krishnan *et al.*, we propose a set of methodological and computational upgrades that produce an operative impact on both accuracy, robustness and computational performance compared to the original solution. In particular, always basing on an unconstrained error minimization over the manifold of rotations, instead of relying on a static set of point correspondences, our algorithm updates the optimization iterations with a dynamically modified set of correspondences in a computationally effective way, leading to substantial improvements in terms of registration accuracy and convergence trend. Other proposed improvements are directed to a substantial reduction of the computational load without sacrificing the alignment performance. Stress tests with increasing views misalignment allowed to appreciate the convergence robustness of the proposed solution. Eventually, we demonstrate that

for very large datasets a further computational speedup can be reached by the adoption of an hybrid (local heuristic followed by global optimization) registration approach.

Keywords: Global Registration, 3D Scanning, Range Images, Correspondence Selection, Newton-type Optimization, Differential Geometry

1. Introduction

Objects and scene modeling from 3D scan data requires an accurate alignment of the acquired multiple views. In a conventional alignment pipeline [1], two consecutive alignment steps, usually called coarse (or initial) and fine (or refined), are performed on all the overlapping views. Either these steps are carried out in a pairwise fashion (with possible accumulation of alignment errors) or according to some multi-view heuristics, high quality 3D modeling requires a final optimized global registration of the whole set of views in order to further reduce and evenly distribute the residual alignment errors. Other than guaranteeing the highest alignment *accuracy*, this last step should also provide other desirable properties. First of all *robustness*, because it is not guaranteed that all views are already well aligned, e.g. when loop closure problems are to be solved. Secondly *computational efficiency*, where computation time becomes an issue especially for dense view and large view collections. Moreover, several advanced specialized and professional application fields (e.g. biomedicine, orthopedics and orthodontia, cultural heritage, industrial design and reverse engineering) tend to generate large data sets (both in terms of number of views, spatial point density, as well as metric accuracy) so that the overall *usability* of the acquisition and processing pipeline emerges as another primary requirement.

Pretty fast heuristic methods have been proposed to handle the global registration problem; the most popular one, and also a recognized reference in terms of performance, is the method proposed by Pulli [2]. However, such kind of methods can not guarantee (as we also verify in this work) to actually minimize the alignment error, so that the views' position remains suboptimal while, as stated, maximized metric accuracy of the views alignment is a sine qua non for the aforesaid demanding application fields. For this reason, more connatural methods for approaching the global registration problem are those that adopt some numerical optimization solutions. However, 1) the

global nature of the problem can easily generate computational issues, and 2) optimization methods are to be well designed and conditioned in order to reduce the risk to be stuck in local minima or to behave incorrectly. Our proposed solution relies on the optimization over the differentiable manifold of 3D rotation matrices (i.e. the orthogonal Lie group $SO(3)$) [3, 4], a functional theoretical framework which, however, in its early proposed global registration implementation [5, 6], reveals itself not fully adequate to address the above problems 1) and 2). Our main contribution here consists in a new re-designed algorithmic approach which significantly improves the performance of the original one, demonstrating to be suitable for the most demanding modern applications in terms of convergence, accuracy, and computational performance. A hybrid approach, which combines a first heuristic-based alignment stage followed by the proposed solution, is also tested in order to further reduce the global registration time for the most demanding cases. Since the proposed solutions are designed to be used in the context of modern and high-quality 3D object modeling chains, a gentle introduction to the problem we tackle is firstly provided in Sec.1.1 (which can be possibly skipped by the informed reader). Subsequently, a literature review is presented in Sec.1.2. Finally, an outline of the paper concludes this introduction (Sec.1.3).

1.1. Problem description

The possibility to create accurate 3D models of real objects is a primal asset of several application domains such as industrial reverse engineering [7] and visual inspection [8], cultural heritage preservation [9], robot localization and navigation [10], biological and medical imaging [11]. This is increasingly influencing a number of related fields such as computer animation and augmented reality, environmental remote sensing, construction site management, e-commerce, entertainment, fashion industry, and more. When metric accuracy is a priority, the acquisition is usually made by active 3D optical devices such as laser or structured light 3D scanners which today guarantee high quality and dense point sets captured from their field of view (FoV). At the same time, the physical FoV limitation of the scanner and the spatial morphology of the objects to acquire, require the collection of several (dozen or even hundreds) views in order to build a complete 3D model of the object or the scene of interest. Starting from a collection of independent views (lying in their own local reference systems), and under proper view overlapping and object surface coverage hypothesis, the construction of a 3D

model is obtained through a pipeline consisting in a series of (geometric) data processing steps, where the main objective is the registration (or alignment) in a common reference system of all views toward a unique digital model, under application-driven requirements that are usually related to accuracy, robustness, computational speed and automatism.

[Figure 1 about here.]

Due to the intrinsic complexity nature of the acquired data (caused by the presence of occlusions, noise, outliers, residual distortions and variable view overlapping area) it is not viable to conceive and formalize the alignment problem according to a unique single-step, closed-form solving approach. A main distinction regards the concepts of ‘local’ and ‘global’. It is clear how the alignment problem would have a global nature as a whole, however it must be recognized that the alignment of a single view is quite a local problem, and that in most cases the acquisition of an object or the navigation of an environment follows a certain path, either decided on-the-fly or planned in advance, and this is usually linked to local decisions as well. Therefore the alignment problem can be conceived as a local-to-global approach, in that each view is registered with respect to its neighboring ones, possibly according to a specific path, with the aim of arriving to a global registration. The role of global registration is twofold: to perform an equal redistribution of residual alignment errors due to data imperfections, and to solve the closure problems that are likely to happen whenever the acquisition path turns around the object, possibly several times and through crossing paths, in order to obtain a complete coverage of the object. The typical strategy employed by well-defined, state-of-art approaches is to apply a local coarse alignment technique followed by a local fine alignment step, which is in turn used to initialize a global registration phase. Usually, each phase defines different problems that require distinct solving approaches. In the present work we focus on the final step of the alignment problem, that is, the global registration phase. We initially assume that we can rely on a fairly good initial alignment of the views, as the one that can ordinarily be provided by reference fine alignment techniques, e.g. those based on pairwise Iterative Closest Point solutions, being the fine alignment usually applied in a sequential pairwise order following the object scanning path. We will also strain our system by gradually releasing the strictness of the fine alignment assumption, while concurrently performing an enlarged exploration of the basin of convergence of the proposed optimization solution. Overall, our method will prove to be

effective, robust and versatile to guarantee a highly accurate global alignment for a wide class of challenging (large and dense) experimental datasets (here anticipated in Fig. 1) and initial alignment conditions.

1.2. The related work

Due to the complexity of the alignment task and the strict interdependence of the phases of fine alignment and global registration, we now focus on an essential literature overview of the above steps of the 3D modeling pipeline. For a state-of-the-art of *coarse alignment* techniques, the interested reader can start off with our recent work [12].

Fine alignment These methods are usually employed to improve the alignment obtained after the application of a coarse alignment technique. Many approaches to this problem are related to the Iterative Closest Point (ICP) technique, first introduced in [13] and [14], and to its efficient variants introduced afterwards [15], [16], [17], [18]. The ICP technique is an iterative approach for which, at every iteration, a set of correspondences is established between two views, and used to estimate the rototranslation matrix that brings these views closer to alignment. Due to problems induced by sample noise and partial occlusions in the superposition zones, some methods try to avoid the computation of the correspondence set and define the registration problem using surface distance functions [19], [20]. Some work has also been done on ‘non-rigid’ fine alignment as a form of compensation for nonlinearities related to the acquisition device, such as optical distortions and calibration errors [21, 22]. These techniques can be used to compensate for discrepancies that appear in the overlapping areas of different scans, which cannot be removed by the application of rigid transformations, since each scan is subject to noise and distortion introduced during the acquisition process. To this end, also some classes of surface reconstruction algorithms can effectively cope with this kind of residual misalignment [23].

Global registration After an initialization of the presentation of the different scans of an object provided by a multiview coarse alignment, which is commonly followed by a sequence of multiple pairwise fine alignments, there is the need to perform a simultaneous ‘adjustment’ (commonly addressed as ‘global registration’) of the positions of all the views. In fact, a mere chained application of pairwise fine alignments (as proposed in early works [13], [24]) is usually unable to solve the typical problems that arise in multi-view alignment, such as error propagation and loop closure issues. A global registration can be done in several ways and it usually leads to formalizations which com-

prise non-convex minimization problems to solve with constraints related to the rigid transformations of point sets belonging to the views. Graph based representations of the overlapping views can help to define heuristics or algorithms to better condition the problem while keeping the computational load under control. In general, although various methods exhibit similarities, it can be very difficult to compare them due to several factors: the diversity of used metrics and experimental conditions, different HW and SW environments and, not least, the large heterogeneity existing in the dataset we may want to consider (e.g. due to creation/acquisition tools/equipments, synthetic/real(noisy) data, point/vertexes density, mesh/point cloud format, number of views,...). In general, convergence trend studies are valuable but should be carefully verified on real-world data and, as we will see, it is not at all guaranteed that a better rate of convergence is also towards a better minimum. The point correspondence theme is also very important. Some works are based on out-of-core computed correspondence sets, others foresee to update the correspondences in an iterative manner, others try to avoid using them at all.

Among representative early works on global registration, Blais and Levine [25] minimize an Euclidean distance cost function, calculated on sets of control points, by simulated annealing. Later, Silva *et al.* [26] adopt a similar approach exploiting genetic algorithms with a surface interpenetration measure. Bergevin *et al.* [27] organize the pair of views in a network structure to simultaneously and iteratively minimize the alignment error. Benjema and Schmitt [28] propose a randomized ICP over a multi z-buffer structure capable of representing overlapping portions of the views and accelerate operations on them. Another generalized multiview ICP is proposed by Williams and Bennamoun [29]. Pennec [30] propose a method which alternates between the calculation and refinement of a mean shape and the alignment of the views on it, while Stoddart and Hilton [31] map the registration problem on a physically inspired model where a minimum of potential energy is found with an iterative numerical method based on gradient descent. An improved force-based optimization method is also proposed by Eggert *et al.* [32]. Huber and Hebert [33] use a global consistency measure on a graph of pairwise matching looking for globally connected subgraphs on which to solve a multiview point-to-plane distance minimization problem according to the approach proposed by Neugebauer [34]. Pulli in [2] uses correspondence-based pairwise alignments and matching heuristics as constraints for its effective multi-view error distribution among the aligned views. The method

is still considered a valid approach thanks to its effective error distribution heuristic, computational efficiency, outlier rejection rules and the capability to handle large datasets. However, as we also observed in our experiments, some criticality remain when handling multiple closure problems (from the acquisition of complex objects), as well as when the number of views to align increases up to the point that the underlying heuristic fails to converge to the global minimum of the error function.

Quaternion representation of rigid rototranslation transformations [35] is exploited in several global registration works [36], [37], [38]. Benjema and Schmidt [36] demonstrate that the optimal translation can be decoupled and solved independently from the optimal rotation. Their approach is based on an iterative method where rotation solutions are found based on the movement of one view at a time while keeping the others fixed. The same decoupling is exploited in the work of Sharp *et al.* [37] where optimization over the graph of neighboring views in a quaternion space is done and closed form solutions are obtained on the cycles of a graph decomposition. The method does not require the computation of point correspondences and can be combined with any pairwise alignment algorithm to generate the estimates of relative motion between each pair of views. Torsello *et al.* [38] adopt a dual quaternion formulation to project pairwise alignments on the same reference frame and to perform a diffusion along the graph of adjacent views. Also in this case an alternation of the diffusion method and ICP pairwise alignment is suggested for real-world data where alignment performance similar to those of the Pulli method are obtained.

Pottman *et al.* [39] developed a method based on a first order kinematical analysis that exploits local quadratic approximants of the squared distance function associated to the surfaces to be aligned. This is further investigated in [40], where the global registration problem from a geometric optimization point of view is studied, trying to provide a theoretical framework for better understanding of the empirical results reported in former literature (including ICP-based methods which exhibit linear convergence), and to explore various constrained nonlinear least squares approaches based on Newton type descent algorithms which can lead to faster (locally quadratic) convergence. Krishnan *et al.* [5], [6] proposed an unconstrained optimization procedure which exploits translation and rotation decoupling and solve for the vector of all view rotations through an optimization method that explicitly works on the constrained manifold of rotations $SO(3)$. Breitenreicher and Schnörr [41] further show (even if limited to pairwise alignment) that taking the intrinsic

geometry of the underlying manifold into account for registration, significantly increases robustness with respect to poor initialization. The global registration method described in [6] guarantees a closed form solution in the noiseless case (assuming the knowledge of a set of ‘exact’ correspondences), while for real-world data the use of an iterative Gauss-Newton scheme obtained locally quadratic convergence and computational costs independent of the number of points in the views. In fact, to perform global registration, a set of correspondences between the overlapping pairs of views has to be provided, which is kept fixed during minimization. In the noisy case this ‘known correspondences’ assumption is motivated by the fact that methods like ICP already deliver reliable correspondences at a pairwise level and that global error distribution can rely on that. This is the same line of reasoning we can find in Pulli’s work and in others that see the possibility of performing global error distribution without or with minimal correspondence updating. This is also considered an enabling factor to attain computationally affordable techniques. Nevertheless, with regard to computational performance, it has been observed by some comparison studies [42], [43] that, in general, accuracy of the alignment comes at the price of heavier computational cost. Therefore, methods that reach a good compromise between accuracy and speed are of great value for practical applications. Our effort here is directed to advance some steps in this direction by proposing and demonstrating that a substantial accuracy improvement and benefits in terms of computational cost can be achieved with respect to [6] with a solution based on iterative correspondence updating and other adjuvant technical solutions.

1.3. Paper outline

Following, we focus on rigid global registration solutions which are particularly suited for good quality (low noise and distortion) and high-resolution (computational demanding) range scans produced by modern optical range scanners, such as structured light optical devices. We propose some operative modifications to the original ‘Optimization-on-a-Manifold’ (OOM) framework proposed in [6] (here recapped in Sec.2.1 and 2.2) which are intended to obtain a superior alignment accuracy as well as a boost in computational speed. In Sec.2.3 we present the main ideas at the base of our ‘Improved Optimization-on-a-Manifold’ (IOM) technique. They consist in the introduction of a fast technique for updating correspondences at each OOM iteration, which was not foreseen in the original algorithm and which can give significant improvements in terms of convergence rate and alignment

accuracy. Moreover, we devise a significant computational speedup with respect to the original algorithm, and a different correspondence matching strategy that contributes to improve the convergence rate. Other optimized implementation solutions are described in Sec.3, which allow us to increase the execution speed of the IOM method. Experimental results and discussions are presented in Sec.4, where robustness (with respect to input misalignments), accuracy and execution time are evaluated and compared with respect to reference methods. The present work is an updated extension of our conference paper [44], where results are presented on a substantially expanded dataset. In particular, the paper has been reorganized, improved and partially rewritten in all parts (with expanded introductory sections), a new stress test has been introduced aiming to establish the basin of convergence of the proposed method, as well as an enriched comparison with the Global Registration method of Pulli [2], with also new experiments about its combined usage with IOM, aimed to explore the possibility to better manage larger datasets.

2. Material and Methods

2.1. Mathematical framework

In this section we concisely give an overview of the underlying mathematical framework for both the original OOM and our modified algorithm described later. We essentially maintained the notation of [6] and [4] to which the reader can refer for additional details. An early reference about minimization on manifolds for computer vision applications is the work by Taylor and Kriegman [3].

Let’s start with an idealization of the global registration problem. We will see how, even if initially quite far from a real acquisition setup, this framework allows to derive an optimization core that, with due expedients, can be successfully deployed for the solution of real problems.

2.1.1. Algebraic derivation of the error function

Consider a 3D object as a set of K points lying in a ‘world’ reference system:

$$W = \{w^k \in \mathbb{R}^3, k = [1, K]\} \quad (1)$$

Now imagine to acquire N views of the object, each taken from a different viewpoint. Then, view i only see a subset $W_i \subseteq W$, that is:

$$W_i = \{w_i^k \in \mathbb{R}^3, k = [1, K_i]\} \subseteq W \quad (2)$$

Viewpoint variation also causes each point subset to lie within a different reference frame:

$$V_i = \{v_i^k \in \mathbb{R}^3, k = [1, K_i]\} \quad (3)$$

For each view i , however, it is possible to define the rotation and translation matrices (R_i, t_i) that bring the viewpoint reference frame into the ‘world’ reference frame, so that the following holds:

$$w_i^k = R_i v_i^k + t_i \quad (4)$$

In order to globally register the N views, each must have an overlapping area with respect to the others. Given a pair of views, say i and j , we address $W_{ij} = W_i \cap W_j$ as the set of K_{ij} 3D points that appear in both views. Clearly, $W_{ij} = W_{ji}$, and $K_{ij} = K_{ji}$. With respect to the acquisition reference frames we have:

$$\begin{aligned} V_{ij} &= \{v_{ij}^k \in \mathbb{R}^3, k = [1, K_{ij}]\} \\ V_{ji} &= \{v_{ji}^k \in \mathbb{R}^3, k = [1, K_{ij}]\} \end{aligned} \quad (5)$$

such that, by applying (4), we have:

$$R_i v_{ij}^k + t_i = w_{ij}^k = R_j v_{ji}^k + t_j \quad (6)$$

So, we can reinterpret what we defined above as an ‘ideal’ global registration problem, to be solved with the minimization of an error function:

$$g = \sum_{i=1}^N \sum_{j=i+1}^N \sum_{k=1}^{K_{ij}} \|(R_i v_{ij}^k + t_i) - (R_j v_{ji}^k + t_j)\|^2 \quad (7)$$

After some algebraic operations that we minimally trace here, error function (7) can be rewritten¹ into a matrix form:

$$\begin{aligned} g(\mathcal{R}, \mathcal{T}) &= tr(\mathcal{R} \mathcal{A} \mathcal{R}^T + 2\mathcal{R} \mathcal{B} \mathcal{T}^T + \mathcal{T} \mathcal{C} \mathcal{T}^T) \\ &= tr(\mathcal{R} \mathcal{A} \mathcal{R}^T) + 2vec^T(\mathcal{T}) vec(\mathcal{R} \mathcal{B}) + \\ &\quad + vec^T(\mathcal{T}) (\mathcal{C} \otimes I_3) vec(\mathcal{T}), \end{aligned} \quad (8)$$

with $\mathcal{R} := [R_1 \ R_2 \ \dots \ R_N] \in \mathbb{R}^{3 \times 3N}$
and $\mathcal{T} := [t_1 \ t_2 \ \dots \ t_N] \in \mathbb{R}^{3 \times N}$

¹Given a $r \times c$ matrix A , we recall the definition of the vectorization operator $vec(\cdot)$, which generates a $rc \times 1$ vector: $vec(A) = [a_{1,1}, \dots, a_{r,1}, a_{1,2}, \dots, a_{r,2}, \dots, a_{1,c}, \dots, a_{r,c}]^T$. The trace operator $tr(\cdot)$ of a square matrix B is such that $tr(B) = \sum_i b(i, i)$. The tensor (or Kronecker) product of two matrices A and B , $A \otimes B$, is a block matrix where each elements $a_{i,j}$ of A is substituted with the block element $a_{i,j} B$.

Matrices A, B and C are constructed as follows:

$$\begin{bmatrix} A & B \\ B^T & C \end{bmatrix} = \sum_{i=1}^N \sum_{j=i+1}^N \sum_{k=1}^{K_{ij}} \begin{bmatrix} a_{ij}^k \\ e_{ij} \end{bmatrix} [a_{ij}^{kT} \quad e_{ij}^T] \geq 0$$

where e_i is the i_{th} column of an $N \times N$ identity matrix I_N , $e_{ij} = e_i - e_j$ and $a_{ij}^k = (e_i \otimes I_3) v_{ij}^k - (e_j \otimes I_3) v_{ji}^k$.

It is possible to remove the dependency from \mathcal{T} by first minimizing (8) with respect to $vec(\mathcal{T})$, which gets $\mathcal{T}(\mathcal{R}) = -\mathcal{R}BC^\dagger$. Now substituting back $\mathcal{T}(\mathcal{R})$ in (8), a new error function can be obtained:

$$f(\mathcal{R}) = tr(\mathcal{R}\mathcal{M}\mathcal{R}^T) = vec^T(\mathcal{R}^T)(I_3 \otimes M)vec(\mathcal{R}^T) \quad (9)$$

with \mathcal{M} equal to $A - BC^\dagger B^T$, where C^\dagger is the pseudoinverse of C , since matrix C is singular.

2.1.2. The product manifold of $SO(3)^N$ and its parametrization

Now we concisely look at the geometry of the special orthogonal group $SO(3)$ (the group of 3×3 orthogonal matrices with determinant $+1$, $R_i \in SO(3)$ for $i = [1, N]$) which is a Lie group. Its group operator is matrix multiplication. Its associated Lie algebra \mathfrak{so}_3 is the set of 3×3 skew symmetric matrices of the form:

$$\Omega = \begin{bmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & -\omega_x \\ -\omega_y & \omega_x & 0 \end{bmatrix} \quad (10)$$

Now, considering a vector $\omega = [\omega_x \ \omega_y \ \omega_z] \in \mathbb{R}^3$ this can be transformed to the matrix $\Omega \in \mathfrak{so}_3$ by the map $\Omega = \Omega(\omega) = Q_x \omega_x + Q_y \omega_y + Q_z \omega_z$, with Q_x , Q_y and Q_z equal to

$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

respectively.

This denotes a known isomorphism which exists from the Lie algebra (\mathbb{R}^3, \times) to the Lie algebra $(\mathfrak{so}_3, [.,.])$, with \times the cross product and $[.,.]$ the matrix

commutator. Of interest here is the N -fold product manifold of $\text{SO}(3)$ which is a smooth manifold of dimension $3N$, given by

$$\text{SO}(3)^N = \text{SO}(3) \times \dots \times \text{SO}(3) \text{ (} N \text{ times)} \quad (11)$$

Recall that the tangent space of $\text{SO}(3)$ at R_i is given as $T_{R_i}\text{SO}(3) = \{R_i\Omega_i \mid \Omega_i \in \mathfrak{so}_3\}$ and the affine tangent space is $T_{R_i}^{\text{aff}}\text{SO}(3) = \{R_i + R_i\Omega_i \mid \Omega_i \in \mathfrak{so}_3\}$. Define the block diagonal matrix

$$\tilde{\Omega} = \Omega_1 \oplus \Omega_2 \oplus \dots \oplus \Omega_N, \quad \Omega_i \in \mathfrak{so}_3 \quad (12)$$

Due to isomorphism, the tangent space of $\text{SO}(3)^N$ at $\mathcal{R} \in \text{SO}(3)^N$ can be identified as $T_{\mathcal{R}}\text{SO}(3)^N = \mathcal{R}\tilde{\Omega}$ and the affine tangent space is $T_{\mathcal{R}}^{\text{aff}}\text{SO}(3)^N = \mathcal{R} + \mathcal{R}\tilde{\Omega}$. Let $\mathcal{N}(0) \subset \mathbb{R}^3$ denote a sufficiently small open neighborhood of the origin in \mathbb{R}^3 , and let $R_i \in \text{SO}(3)$. Then the exponential mapping

$$\mu : \mathcal{N}(0) \subset \mathbb{R}^3 \rightarrow \text{SO}(3), \quad \omega_i \mapsto R_i e^{\Omega_i(\omega_i)} \quad (13)$$

is a local diffeomorphism from $\mathcal{N}(0)$ onto a neighborhood of R_i in $\text{SO}(3)$. Due to isomorphism, the product manifold $\text{SO}(3)^N$ at $\mathcal{R} \in \text{SO}(3)^N$ can be locally parametrized by

$$\begin{aligned} \varphi : \mathbb{R}^{3N} &\rightarrow \text{SO}(3)^N, \\ \omega &\mapsto \mathcal{R} e^{\tilde{\Omega}(\omega)} \end{aligned} \quad (14)$$

with $\omega = [\omega_1 \dots \omega_N]^T \in \mathcal{N}(0) \times \dots \times \mathcal{N}(0) \subset \mathbb{R}^{3N}$, and $\mathcal{R} e^{\tilde{\Omega}(\omega)} = \mathcal{R} (e^{\Omega(\omega_1)} \oplus \dots \oplus e^{\Omega(\omega_N)})$.

2.1.3. Constructing a local approximation

A local approximation of the error function f (9) is now constructed using a second order Taylor expansion. Instead of differentiating f , the local parametrization of $\text{SO}(3)$ will be exploited by performing the approximation on the function $f \circ \varphi$, whose domain is \mathbb{R}^{3N} . Intuitively, the use of the local parametrization φ ensures to always remain on the manifold. The cost function f at $\mathcal{R} \in \text{SO}(3)^N$ expressed in local parameter space using a smooth local parametrization $\varphi_{\mathcal{R}} : \mathbb{R}^{3N} \rightarrow \text{SO}(3), \omega_i \mapsto R_i e^{\Omega(\omega_i)}$ is given by:

$$f \circ \varphi_{\mathcal{R}}(\omega) = \text{tr} \left(\mathcal{R} e^{\tilde{\Omega}(\omega)} M e^{\tilde{\Omega}(\omega)^T} \mathcal{R}^T \right) \quad (15)$$

Now we express

$$\begin{aligned}
& j_0^{(2)}(f \circ \varphi_{\mathcal{R}}) : \mathbb{R}^{3N} \rightarrow \mathbb{R}, \\
\omega \mapsto & \left((f \circ \varphi_{\mathcal{R}})(t\omega) + \frac{d}{dt}(f \circ \varphi_{\mathcal{R}})(t\omega) + \right. \\
& \left. + \frac{1}{2} \frac{d^2}{dt^2}(f \circ \varphi_{\mathcal{R}})(t\omega) \right) \Big|_{t=0}
\end{aligned} \tag{16}$$

as the second order Taylor approximation of $f \circ \varphi_{\mathcal{R}}$ about $0 \in \mathbb{R}^{3N}$ in direction ω , where:

$$\begin{aligned}
(f \circ \varphi_{\mathcal{R}})(t\omega) \Big|_{t=0} &= \text{tr}(\mathcal{R}M\mathcal{R}^T) \\
\frac{d}{dt}(f \circ \varphi_{\mathcal{R}})(t\omega) \Big|_{t=0} &= 2 \cdot \text{tr}(\mathcal{R}\tilde{\Omega}M\mathcal{R}^T) \\
&= 2\omega^T \nabla_{f \circ \varphi_{\mathcal{R}}}(0) \\
\frac{1}{2} \frac{d^2}{dt^2}(f \circ \varphi_{\mathcal{R}})(t\omega) \Big|_{t=0} &= \omega^T H_{f \circ \varphi_{\mathcal{R}}}(0) \omega
\end{aligned} \tag{17}$$

Since $\text{vec}(\Omega^T) = Q\omega$, we can compute the gradient as:

$$\nabla_{f \circ \varphi_{\mathcal{R}}}(0) = J^T \text{vec}(M\mathcal{R}^T) \tag{18}$$

where

$$\begin{aligned}
J &= (\mathcal{R} \otimes I_{3N}) \tilde{Q} \\
\tilde{Q} &= Q_{e1} \oplus Q_{e2} \oplus \dots \oplus Q_{eN} \quad Q_{ei} = \begin{bmatrix} e_i \otimes Q_x \\ e_i \otimes Q_y \\ e_i \otimes Q_z \end{bmatrix}
\end{aligned} \tag{19}$$

By defining $H = \hat{H} + \tilde{H}$, the quadratic expansion term of (17) can be decomposed in a sum of two terms

$$\omega^T H_{f \circ \varphi_{\mathcal{R}}}(0) \omega = \omega^T \hat{H}_{f \circ \varphi_{\mathcal{R}}}(0) \omega + \omega^T \tilde{H}_{f \circ \varphi_{\mathcal{R}}}(0) \omega \tag{20}$$

where:

$$\begin{aligned}
\hat{H}_{f \circ \varphi_{\mathcal{R}}}(0) &= J^T (I_3 \otimes M) J \geq 0 \\
\tilde{H}_{f \circ \varphi_{\mathcal{R}}}(0) &= -\tilde{Q}^T (I_{3N} \otimes M\mathcal{R}^T\mathcal{R}) \tilde{Q}
\end{aligned} \tag{21}$$

2.2. The original algorithm (OOM)

It is now possible to describe the algorithm proposed in [6], which consists in the following iteration:

$$s = \pi_2 \circ \pi_1 : \text{SO}(3)^N \rightarrow \text{SO}(3)^N \tag{22}$$

where π_1 maps a point $\mathcal{R} \in \text{SO}(3)^N$ to an element in the affine tangent space $T_{\mathcal{R}}^{\text{aff}}\text{SO}(3)^N$ that minimizes $j_0^{(2)}(f \circ \varphi_{\mathcal{R}})(0)$ and π_2 projects that element back to $\text{SO}(3)^N$ by means of the parametrization $\varphi_{\mathcal{R}}$. In turn, optimization in local parameter space π_1 consists of two steps: first, calculate a suitable descent direction and then search for a step length that ensures reduction in cost function, as described by the mapping:

$$\pi_1 = \pi_1^b \circ \pi_1^a : \text{SO}(3)^N \rightarrow T^{\text{aff}}\text{SO}(3)^N \quad (23)$$

Here, π_1^a is used to obtain a descent direction

$$\pi_1^a : \text{SO}(3)^N \rightarrow T^{\text{aff}}\text{SO}(3)^N, \quad \mathcal{R} \mapsto \mathcal{R} + \mathcal{R}\tilde{\Omega}(\omega_{\text{opt}}(\mathcal{R}))$$

where $\omega_{\text{opt}} \in \mathbb{R}^{3N}$ as a function of $\mathcal{R} = \varphi_{\mathcal{R}}(0)$ can be given by the Newton direction when $H_{f \circ \varphi_{\mathcal{R}}}(0) > 0$ as

$$\omega_{\text{opt}}^{\text{Newton}}(\mathcal{R}) = -[H_{f \circ \varphi_{\mathcal{R}}}(0)]^{-1} \nabla_{f \circ \varphi_{\mathcal{R}}}(0) \quad (24)$$

otherwise, Gauss direction is calculated:

$$\omega_{\text{opt}}^{\text{Gauss}}(\mathcal{R}) = -\left[\hat{H}_{f \circ \varphi_{\mathcal{R}}}(0)\right]^{-1} \nabla_{f \circ \varphi_{\mathcal{R}}}(0) \quad (25)$$

Once an optimal direction is computed, an approximate one-dimensional line search is carried out in that direction. The search is performed over the scalar value $\lambda > 0$ such that the cost function $f \circ \varphi_{\mathcal{R}}(0)$ is reduced at every step, giving rise to the mappings:

$$\begin{aligned} \pi_1^b : T^{\text{aff}}\text{SO}(3)^N &\rightarrow T^{\text{aff}}\text{SO}(3)^N \\ \mathcal{R} + \mathcal{R}\tilde{\Omega}(\omega_{\text{opt}}(\mathcal{R})) &\mapsto \mathcal{R} + \mathcal{R}\tilde{\Omega}(\lambda_{\text{opt}}\omega_{\text{opt}}(\mathcal{R})) \end{aligned} \quad (26)$$

where λ_{opt} is the step length that reduces the cost function in direction ω_{opt} , and is found using the simple backtracking line search.

Eventually, the OOM algorithm pseudocode can be summarized as follows:

```

initialize ( $\mathcal{R}, \mathcal{T}$ );
provide the set of correspondences;
while (! convergence) do
    apply ( $\mathcal{R}, \mathcal{T}$ ) to the corresponding points;
    compute matrices  $A, B$ ;

```

```

compute matrices  $\mathcal{M}, J$ ;
compute gradient  $\nabla_{f \circ \varphi_{\mathcal{R}}}(0)$ ;
compute  $\hat{H}_{f \circ \varphi_{\mathcal{R}}}(0)$  and  $\tilde{H}_{f \circ \varphi_{\mathcal{R}}}(0)$ ;
if  $(\hat{H}_{f \circ \varphi_{\mathcal{R}}}(0) + \tilde{H}_{f \circ \varphi_{\mathcal{R}}}(0) > 0)$  then
    compute Newton step  $\omega_{\text{opt}}^{\text{Newton}}(\mathcal{R})$ ;
else
    compute Gauss step  $\omega_{\text{opt}}^{\text{Gauss}}(\mathcal{R})$ ;
end if
compute step size  $\lambda_{\text{opt}}$ ;
projection step:  $\mathcal{R} = \mathcal{R}e^{\tilde{\Omega}(\lambda_{\text{opt}}\omega_{\text{opt}}(\mathcal{R}))}$ ;
compute matrix  $\mathcal{T}$  from  $\mathcal{R}$ ;
end while
apply  $(\mathcal{R}, \mathcal{T})$  to the range data;

```

The two initialization steps are left quite general on purpose because there are significant differences depending on whether the algorithm is used in an ideal context or in a real context, as will be discussed later.

2.3. Algorithm modifications

Despite the original OOM solution [6] can be seen as a definitely favorable mathematical framework to the problem of global registration, it presents the limitation of remaining stuck on the same set of correspondences, without updating them during the iterative minimization procedure. If, on one hand, this produces a fast convergence to a solution, at the same time this represents a great limitation (as we will clearly see) to the potentialities of the method, especially in terms of accuracy, i.e. the main requirement we have. On the other hand, updating the correspondence at each iteration appears discouraging because, even not considering that the OOM already presents some computational burden, this would come with a great increase in computation load, especially for large datasets, i.e. our most challenging target. In the present work, having recognized the importance of the correspondence update, we conceive and experiment methodological and algorithmic solutions to realize it efficiently, producing at the same time relevant improvements in terms of alignment accuracy and a performance boost in terms of robustness and convergence behavior, especially for challenging datasets. The main adopted solutions are presented later in this section where, to give an immediate idea of their effect, they are tested on representative cases. Other implementation solutions which determine further computational gains are

presented in Sec.3, while more exhaustive experimental comparisons on the considered set of data are documented in Sec.4. In order to allow a direct and fair comparison of the original OOM and the proposed IOM we reimplemented the OOM algorithm as described in [6], and use the same software modules for our IOM implementation.

[Figure 2 about here.]

[Figure 3 about here.]

2.3.1. \mathcal{R} and \mathcal{T} initialization

The first substantial difference we introduce with respect to the algorithm presented in [5, 6] is about the initialization of matrices \mathcal{R} and \mathcal{T} . The original work presents a closed-form solution that is capable of perfectly aligning, in a single step, a set of views in the ideal case of absence of noise and prior knowledge of exact correspondences. Of course, this cannot be assumed for real registration and modeling scenarios. What the authors of [5, 6] propose in these cases is to use a modified version of the closed-form solution to initialize \mathcal{R} and \mathcal{T} at the beginning of the iterative OOM algorithm. However, under our assumptions of good initial alignment, such an initialization is redundant thus it can be skipped, to the benefit of the computation time. Thus, in a real scenario, one can initialize \mathcal{R} and \mathcal{T} to zero motion, that is to identity R_i 's and zero valued t_i 's, and let the first optimization step do the hardest job. This will be clearer in Sec.3.1.2 where we will take full advantage of this zero initialization assumption to realize a major reduction of computation time, while in Sec.4.2 the application extent and robustness of our initialization related assumptions is evaluated.

2.3.2. Correspondence set update

As stated, a major limitation of the framework proposed in [6] is related to the fact that it requires an ‘exact’ correspondence set as an input. However, in such a case, a closed form solution such as [35] would suffice in recovering the correct alignment. In real cases, e.g. under the reasonable assumption we made in Sec.1.1 of a good approximate initial alignment between the views, the initial correspondences that can be obtained are still far from being exact. In order to get rid of such an implausible correctness requirement, rather than exploiting the same correspondence set throughout the optimization process, we constantly update the correspondence set at each iteration by randomly selecting points on one scan, and determining the matching points on each

overlapping scan. This way, provided that the initial alignment between the views is close enough to the optimal solution, at each iteration the views will be brought closer to the final solution, which in turn improves the correctness of the next correspondence set, until convergence is reached. Updating the correspondence set, however, comes at a price, as it would require for each iteration applying the estimated roto-translation matrices at each scan in the dataset to allow the correspondence update procedure. This would signify a substantial increase in computational cost, proportional to both the number and resolution of the scans, which would render our modified technique not at all appealing for practical applications. However, such computational burden can be considerably reduced by employing a smart correspondence update procedure (described in Sec.3.1.1).

One factor that can negatively affect the success of any correspondence-based alignment technique is the potential presence of outliers: depending on its severity, such presence can slow down, or even prevent the attainment of a successful alignment. In our case (good initial alignment assumption) we can devise to use a simple outlier rejection policy, which should be effective in handling problems that occur on the scan borders and holes. To this purpose, we implemented a twofold correspondence filtering criteria: for each potential correspondence, both the maximum euclidean distance between candidate points and the angle formed between their normals are computed. In case one of such values exceeds a given threshold, the considered correspondence is discarded. Default values employed for such thresholds were 2 mm for the absolute distance, and 60 degrees for the normal angles. However, in some cases, the distance threshold needed to be adjusted to the field of view of the acquisition device.

Fig. 2 allows to appreciate the impact of the correspondence update, both in absolute terms and on the convergence behavior, by depicting the convergence rates obtained for the dataset ‘Capital’ (see Fig. 1 for a pictorial and Sec.4 for a quantitative description of the datasets) by the original algorithm (thus without updating the correspondences) with respect to a modified version of the OOM that instead updates the correspondence set at each iteration: while the first one only obtains a modest improvement of convergence in a single iteration and does not progress any further², the second one ex-

²This is consistent with what observed in [6], where only 2-3 iterations were enough to get convergence for datasets of similar size.

exploits the additional information to reach a definitely better alignment and convergence behavior.

2.3.3. Correspondence matching strategy

[Figure 4 about here.]

Another factor that influences the algorithm convergence rate is related to the correspondence matching strategy that is employed. Given the source (S) and destination (D) scans, and a point s_p with associated normal s_n on the source scan, the matching process tries to determine the corresponding point d_p on the destination scan that best matches the original point s_p (see Fig. 4). Rather than employing a straightforward closest point search (Fig. 4a), we compared two different variants of the normal shooting strategy, depicted respectively in Fig. 4b and Fig. 4c. Both the strategies proceed iteratively as follows: given the starting point s_p on S , at first the closest point to D surface d'_p is computed. Then, the difference vector $l' = d'_p - s_p$ is computed and projected onto the source normal $\pi' = l' \cdot s_n$. Such projection value is used to compute a new starting point $s'_p = s_p + \pi' \cdot s_n$ along the normal vector which draws closer to the D surface. From this point on, we iterate until the corresponding point to s_p^n (the new source point at iteration n) does not change with respect to the one found at s_p^{n-1} . Given the destination point d_p^n determined through this iterative procedure, the first variant assumes as putative correspondence to s_p the point $d_p = d_p^n$, while the second one determines a fictional point by performing a last projection of d_p^n onto the source normal, obtaining $d_p = s_p^n + \pi^{n+1} \cdot s_n$. This second variant is particularly useful when the two scans are close to alignment, since the norm associated with the distance $d_p - s_p$ tends to be consistently lower than the norm of vector l , and thus better convey the alignment condition (similarly to what happens for point-to-point versus point-to-plane error metric). We compared the convergence rates obtained for the dataset ‘Venus’ of the first variant [13] with respect to the second version we adopted, to guarantee a fair error comparison point-to-plane metric has been employed in both cases. As can be seen in Fig. 3, the second variant obtained a non-negligible error reduction³. These evaluations are consistent and lead, on a regular basis, to

³In the assessment of a global registration technique even an improvement of the order of 10 μm , as obtained here, can be significant. As we will see in the experimental section this improvement is of the same order of magnitude of the final obtained average view

an improvement we observed for all the considered datasets. Therefore from now on we adopt this modified normal shooting technique for correspondence matching as an additional feature of our improved manifold optimization.

[Figure 5 about here.]

2.3.4. Avoiding \tilde{H} computation

One of the most expensive operations foreseen in [6] is the computation of the matrix \tilde{H} , used to determine the optimal descent direction for a Gauss-Newton iteration (see (16)-(21)). The reason can be easily understood if one considers the dimension of the matrices involved in its calculation:

$$\underbrace{\tilde{H}_{f \circ \varphi_{\mathcal{R}}}(0)}_{[3N \times 3N]} = \underbrace{-\tilde{Q}^T}_{[3N \times 9N^2]} \underbrace{(I_{3N} \otimes M\mathcal{R}^T\mathcal{R})}_{[9N^2 \times 9N^2]} \underbrace{\tilde{Q}}_{[9N^2 \times 3N]}, \quad (27)$$

where N is the number of images constituting the dataset. In our implementation, for a small dataset of 23 views, \tilde{H} computation takes approximately half of the time required by a single iteration. Moreover, as stated in [5], there is no guarantee that it could be eventually employed. In fact $H = \hat{H} + \tilde{H}$ is in general not positive-definite, while as seen in (25), a trustworthy descent direction can be estimated from \hat{H} , which is positive-semidefinite (as reported in [5]). This is confirmed by a number of tests we run to check in which proportion matrix \tilde{H} contributed to the final convergence. Every performed test on all the considered datasets revealed that \tilde{H} is always discarded since H never passed the positive definiteness test, therefore we have decided to completely ignore the calculation of \tilde{H} , and estimate the optimal descent direction only through \hat{H} , which is expected to provide a reliable result due to its positive-semidefiniteness.

3. Computationally effective solutions

Given the formerly presented algorithmic modifications, additional implementation solutions are detailed in this section which lead to further computation speed up, without involving any accuracy degradation.

distance.

3.1. Boosting the computational speed

3.1.1. Correspondence update

In section 2.3.2 it has been shown how alignment accuracy and related convergence behavior can be improved by updating the correspondence set at each algorithm iteration. Such operation, however, can be very expensive since it would require applying the differential rototranslation matrices $e^{\hat{\Omega}(\lambda_{\text{opt}}\omega_{\text{opt}}(\mathcal{R}))}$ estimated after each iteration (see (13),(14) and the original algorithm in Sec.2.2), to every range data belonging to the dataset. Here we show how to avoid such steps by modifying the correspondence update process, thus greatly reducing its impact on the overall computational cost. Assume we have two range images, V_A^A and V_B^B , lying in their own reference systems (where the subscript letters identify the image, while the superscript letters indicate the reference system). Also, assume we have an estimation of the rototranslation matrices $[R_A, t_A]$ and $[R_B, t_B]$ that bring the images within a common reference system W , thus obtaining V_A^W and V_B^W . Our goal is to determine an updated set of correspondences C_{AB}^W . A standard correspondence update procedure would require to rototranslate images V_A^A and V_B^B , then select a subset of points $S_A^W \subseteq V_A^W$ (*src*) and determine the corresponding points (through, for example, point-to-plane correspondence search) on V_B^W (*dst*). Each found correspondence is put in the correspondence set C_{AB}^W for the current iteration. This procedure is also visually described in the left branch of Fig. 5. In order to avoid the time consuming burden of applying the rototranslation to all the points constituting V_A and V_B , we propose an alternate procedure for updating correspondences that consists in the following steps. We select the samples set S_A^A and rototranslate it in the reference system associated with image B , obtaining S_A^B . Then we determine the correspondence set C_{AB}^B , and finally rototranslate them within the common reference system W , obtaining C_{AB}^W . Fig. 5 shows the two alternative procedures. Since the proposed one only requires to perform the rototranslations over the sample sets S_A^A and C_{AB}^B rather than rototranslate the entire images V_A and V_B , it grants a tremendous improvement in terms of computational speed. Without the proposed expedient the correspondence update would have no practical use due to the preponderance of computational issues with respect to the obtainable benefits.

3.1.2. Matrix precomputing

In the previous section we have shown how to update a set of correspondence points given the scans and the related rototranslation matrices. Other

than being fast, another benefit of this method is that it implicitly applies the rototranslation matrices estimated by the algorithm up to that point: thanks to this fact, the algorithm can be run as if $(\mathcal{R}, \mathcal{T})$ were just initialized to identity (that is, each R_i composing \mathcal{R} is an identity, while each translation vector t_i is set to zero). In practice this means that we can employ a set of identity rototranslation matrices throughout all the steps of the algorithm, with the exception of the correspondence update step, where the overall rototranslation matrices are employed. As a consequence, matrix J (a sparse matrix of dimension $9N^2$ rows by $3N$ columns which depend on \mathcal{R}) can be precomputed and remains constant throughout all the iterations. This constitutes another significant computational reduction which does not influence accuracy.

3.1.3. Final modified algorithm (IOM)

Following, we show how the algorithm pseudocode varies with respect to the modifications proposed in Sec.2.3 and this section:

```

initialize  $(\mathcal{R}, \mathcal{T})$  to zero motion;
precompute  $J$  matrix;
while (! convergence) do
    compute the new correspondence matches;
    compute matrices  $A, B$ ;
    compute  $\mathcal{M}$  matrix;
    compute gradient  $\nabla_{f \circ \varphi_{\mathcal{R}}}(0)$ ;
    compute  $\hat{H}_{f \circ \varphi_{\mathcal{R}}}(0)$ ;
    compute Gauss step  $\omega_{\text{opt}}^{\text{Gauss}}(\mathcal{R})$ ;
    compute step size  $\lambda_{\text{opt}}$ ;
    projection step:  $\mathcal{R} = \mathcal{R}e^{\tilde{\Omega}(\lambda_{\text{opt}}\omega_{\text{opt}}(\mathcal{R}))}$ ;
    compute matrix  $\mathcal{T}$  from  $\mathcal{R}$ ;
end while
apply  $(\mathcal{R}, \mathcal{T})$  to the range data;

```

3.2. Convergence and stopping rules

The algorithm should iterate the registration until an absolute minimum for the alignment error is reached. However, setting an absolute threshold on the error is impractical, since such a threshold would be dataset-dependent and could not be determined in advance. As an alternative, we could evaluate the average alignment error difference $de_k = e_k - e_{k-1}$ between successive iterative steps, and exit when the difference goes below a given threshold.

Unfortunately, such an approach is again problematic due to the fact that the difference is influenced by two main factors: $de_k = de_{k,imp} + de_{k,var}$, where $de_{k,imp}$ is the error improvement obtained by the technique at every iteration, which will tend to decrease while the approach gets closer to convergence, while $de_{k,var}$ is the error variation caused by the fact that error analysis is performed on different sets of correspondences which are varied at every iteration. An efficient termination rule should therefore determine the value of k when the error dynamic begins to be only marginally influenced by $de_{k,imp}$, and is only influenced by $de_{k,var}$.

To do so, we consider a vector $eVec_k$ of the last error values e_k recorded during the previous five iterations, and estimate their standard deviation $\sigma(eVec_k)$. When close to alignment, such value should only be influenced by the same error variations that give rise to the factor $de_{k,var}$. Such variations are likely to follow a Gaussian distribution, due to the random selection of the correspondence sets. Therefore by setting a threshold $thresh_\sigma = 1.96 \cdot \sigma(eVec_k)$, if de_k is greater than $thresh_\sigma$, then it is highly probable that the $de_{k,imp}$ has not yet reached a negligible value, and the iterations should proceed.

We also found it useful to enforce a second constraint, $thresh_e$: if the error e_k of the current iteration results greater than $thresh_e$, it is assumed that convergence has not been reached yet and the iterative process continues. Such threshold is a fraction of the average point spacing on the incoming set of views (to be precise, $thresh_e = \frac{avgSpacing}{4}$), and is helpful to prevent the algorithm to quit prematurely in cases of particularly slow convergence rates. In conclusion, the termination rule can be described as a function, called at the end of every iteration, which accepts as input the current error estimate e , as follows:

```
bool DoAnotherIteration(double  $e$ )
```

```
    insert  $e$  as first element in  $eVec$ ;
```

```
    if (size of  $eVec$   $\geq$  5) then
```

```
        remove last element of  $eVec$ ;
```

```
    else
```

```
        return true;
```

```
    end if
```

```
    if ( $e < thresh_{err}$ ) then
```

```
         $de = eVec[4] - eVec[0]$ ;
```

```

 $e_{stdDev} = \text{stdDev}(\text{eVec});$ 
if ( $de < 1.96 \cdot e_{stdDev}$ ) then
    return false;
else
    return true;
end if
else
    return true;
end if

```

3.3. Hybrid global registration

As it will be confirmed in the next section, heuristic methods can be faster but their convergence is not guaranteed. On the other hand, although being definitely faster than the OOM, our IOM method remains slower than reference heuristics based methods, with a performance gap that tends to increase for bigger datasets (e.g. dozens or even hundreds of hi-res range images), where many correspondences must be evaluated and where matrix inversions suffer from the increase of N .

Despite our effort and achievements in finding technical solutions and expedients to reduce the computational burden of the original framework, other computational optimizations are still possible (optimized sparse matrix handling, multi-threading and parallel computed correspondence update) which are left to further works. What we propose and test here is the possibility to obtain additional computational benefits (always without sacrificing the algorithm’s accuracy) through a simple *hybrid* registration policy which consists in the cascade of an heuristic alignment method followed by an optimization based one, aiming at bringing the scans closer to the optimal solution in a first stage while reducing the subsequent optimization engine workload. In particular, we are interested in evaluating the performance of a hybrid scheme where Pulli’s algorithm [2] is executed prior to the proposed IOM. This will give us the chance to improve computational performance on bigger datasets and also to make some observations that could be useful for the design and use of such hybrid strategy.

4. Results and discussion

In this section we present a number of tests and comparisons involving the proposed IOM solution, the original OOM algorithm [6], the Global Regis-

tration technique proposed by Pulli (hereinafter referred as GRP)[2] and the hybrid approach (here called HYB) defined above. All the algorithms have been implemented in C++ and run on a PC Intel I5 M520 (2×2,4 GHz) with 4 GB of RAM. The implementations of OOM, IOM and HYB exploit the Newmat library package [45] for matrix computation, while the GRP approach does not require any matrix computation.

4.1. Datasets and settings

Comparisons were performed considering a heterogeneous collection of 14 datasets which are well representative of a wide variety of real life objects (Fig. 1). The considered datasets are also diversified in terms of number of views, as they are composed by a number of range scans comprised between 8 and 170 (for a total of around 740 scans). Dimensions of the acquired objects range from 10 cm (Denture) up to 1.5 meters (Neptune). Objects are representative of a variety of real high-quality and challenging acquisition scenarios and present very different surface and geometric properties. Range images have been acquired through a professional structured-light range scanner, and may contain up to 1,3 million points each. The datasets upon which the comparison is performed have been acquired and later processed by an experienced operator that supervised an alignment pipeline comprising a feature-based automated coarse alignment phase [46] followed by an ICP based fine alignment [16] along the acquisition path, and eventually a final global alignment performed through the Pulli’s algorithm [2]. From such a condition of supposed good global alignment, we executed a scrambling of each view by applying a random rototranslation matrix composed by a rotation of maximum ± 3.0 degrees and a translation of maximum ± 0.1 mm for each of the three axes, thus reproducing a situation we can find after a generic coarse alignment phase. This is then followed by a pairwise fine alignment performed over the scan sequence (again using the approach of [16]). This reproduces a realistic and general acquisition scenario in which the range images have been sequentially acquired and aligned in pairs, thus requiring a final global registration phase. This seemingly redundant double global alignment cycle is indeed suitable for a more neutral and flexible performance assessment. In fact, this way we are able to 1) realistically simulate a generic coarse alignment (i.e. without the constraint to be stuck with a specific method with the possible bias this could introduce with the consequence of potentially influencing the reliability and repeatability of our results) and 2) to modulate the gap from the desired optimal result (a very

useful and enabling feature for the stress test we executed).

In order to quantitatively assess the global registration quality of the considered methods we define an *average distance* over the dataset, which is computed by averaging error contributions referred to each scan of the dataset. The *scan error* is, in turn, computed as the average of the pairwise distances estimated between the current view and all other overlapping scans, the *pairwise distance* being calculated by first selecting a suitably large set of points on a view, finding its correspondences in the other scan, and finally calculating the average of their (point-to-plane) distances.

[Figure 6 about here.]

[Figure 7 about here.]

4.2. Robustness to initial misalignments (stress test)

With a preliminary test we want to assess and compare the convergence properties of the GRP and IOM algorithms with respect to an increasing initial view misalignment. To this end, we consider datasets which have been intentionally scrambled with an increasing strength, according to what is described in Sec.4.1 but this time without performing the subsequent fine alignment step. It is worth considering that such scenario does not fully satisfy our assumptions of good initial alignment, since we are stressing our IOM system on purpose, in order to widely assess its basin of convergence. This way, we also inherently verify the feasibility of the zero motion initialization described in Sec. 2.3.1 and further exploited with matrix J precomputation described in Sec. 3.1.2.

Here we consider two datasets, Hurricane and Cupid, and randomly scramble each of their range images by applying a bounded amount of angular and translation offset for each of the three axes. By increasing the maximum value of the offsets we created, for each dataset, 10 different misalignment scenarios which range linearly from ± 1.5 degrees and ± 0.1 mm (scenario 1) up to ± 15 degrees and ± 1.0 mm (scenario 10). The GRP and IOM techniques have been carried out for each of the misaligned datasets, giving the results summarized in Fig. 6, while Fig. 7 shows three different misalignment scenarios for the Cupid dataset (the misalignment entity for the worst case is evident). The obtained results suggest the fact that IOM is far more robust than GRP in avoiding local minima (it only slightly fails for the scenario number 10 of the Cupid dataset). The superior convergence properties of

IOM can be motivated by the fact that it optimally aligns all the views simultaneously at each iteration; on the contrary GRP tries to optimally align each view with respect to the rest, in a sequential way, thus being liable to error propagation and loop closure phenomena. Fig. 8 shows an example of the obtained global registration on the Hurricane dataset: Fig. 8a represents the misaligned dataset, Fig. 8b represents the alignment obtained through GRP (note the misaligned views corresponding to the statue hair, upper right), while Fig. 8c shows the alignment obtained through IOM.

The obtained robustness against misalignments, which we also observed in all the other datasets, reveals that the range of exploitation of our IOM approach is wider with respect to what is commonly intended and required as ‘good’ initial alignment conditions. Such working hypothesis should certainly be preserved in terms of ‘good practice’ rules, which however can be regarded as conservative in our framework. The obtained results are also a clear confirmation of the suitability of the algorithmic initializations we made and the observed stability of the alignment results, irrespective to the different set of randomly chosen correspondences, is an indirect confirmation of the fact that we are still moving within the basin of convergence of the method. In fact, although we cannot give a rigorous mathematical definition and assurance about the basin of convergence of the proposed approach (way too many possible factors can influence the convergence basin shape and dimension), the performed stress test at least gave an experimental evidence that the IOM algorithm can also be considered capable of handling up to moderate misalignments (i.e. more unfavorable conditions than what it can be normally expected to have as input of a global registration phase).

[Figure 8 about here.]

4.3. Global registration performance

4.3.1. Alignment accuracy

For the experimental assessment of the alignment accuracy we are interested in the registration error produced by every implemented algorithm (OOM, IOM, GRP, HYB) for each acquired dataset. The first three columns of Table 1 give quantitative information about the average scan densities (typical of state-of-the-art structured light scanners) as well as the number of views of each considered dataset (in ascending order with respect to dataset size). The fourth column reports the initial average distance among the views associated to their initial alignment condition. Global registration results are

then reported in terms of average and standard deviation of the measured view distances. Note that OOM could not be executed on big datasets, due to memory failures caused by the computation of \tilde{H} (Sec.2.3.4), which reached the maximum matrix dimension (as confirmed by the Newmat online documentation), and therefore generated a runtime exception. Nevertheless, as stated in Sec.2.3.2, OOM did not grant any further improvement at any iteration following the first one, therefore we would not expect it to behave differently when employed on bigger datasets.

4.3.2. Convergence trends

Fig. 9 shows the convergence behavior for a selection of datasets (note that iteration number, which is not informative of the actual computation time, is reported on the horizontal axis). From these results, a number of considerations can be drawn. First, IOM granted the lowest dataset error for all the datasets and, most importantly, the dynamic of its error function drifts regularly toward its minimum, following a linear convergence rate. On the contrary, no hints about the GRP convergence course can be deduced from its error dynamics, confirming the heuristic nature of this technique. While for small datasets (such as Denture, Dog and Capital) all the techniques converge to the optimal minimum, in some cases (such as Hurricane, Venus and Neptune) GRP fails to reach the optimal global registration: in fact, such bigger cases are more liable to incur into problematic phenomena such as error propagation and loop closure. On the other hand, the cascade of GRP followed by IOM (that is, the HYB approach) obtained similar alignment error than the one reached by IOM approach, suggesting once more that the proposed approach is capable of recovering from conditions of potential error propagation misalignments.

4.3.3. Visual quality assessment

Global registration performance obtained for each of the approaches on a small set of the datasets can also be visually inferred in Fig. 10 where the greater perceived color interpenetration denotes a better alignment result. For each dataset its constituting range images are shown, each associated with a different color: in principle, the lower the residual alignment error is, the better the color interpenetration appears. In the first column of Fig. 10 the Crankcase dataset (93 views) is displayed. For the GRP configuration it is easy to identify three distinct colors near the two bigger holes of the model (purple, sky and orange), while for the IOM and HYB approaches a greater

interpenetration is obtained. The second column shows the alignment results for the Neptune dataset (169 views). For the GRP configuration, one can distinguish separate colors on the statue stomach (purple and sky) and pedestal (yellow) more easily than for the alignments obtained with the other techniques. In the third column, for the Capital dataset (23 views), one can see separate colors in the upper part of the object for the GRP, while again a better interpenetration is obtained for the other techniques. At last, for the Cupid dataset (45 views), three main colors (green, blue and purple) can be easily identified for the GRP, while this is not the case for the other two techniques.

4.3.4. Computational speed

Table 2 shows the computational times required for each technique. With respect to OOM algorithm, IOM presents a significantly reduced computational burden: up to 10 times less on smaller datasets, while higher gains seem to be possible for bigger ones, as in the case of Hurricane dataset. Nevertheless, GRP is faster than IOM in reaching its error minimum, especially for bigger datasets: in fact, while GRP is composed by a sequence of ICP runs, which are very fast to compute, IOM needs to consider the alignment globally, causing the computational burden of such approach to increase exponentially with the number of scans to be processed. This is the main reason why we proposed the HYB approach: we try to exploit the GRP speed to bring the alignment closer to optimum, so that the IOM can then converge faster. However, the results we obtained do not allow us to draw any bold conclusion at this stage. In fact, if HYB registrations are as accurate as the IOM ones, the related computational behavior is somehow fluctuating; in some cases (as for Hurricane, Decoration and Crankcase) HYB takes moderately longer than the pure IOM approach, while in others (such as Horse, Shell and Neptune) the HYB grants a significant improvement. Our best guess to explain such behavior is that it may be possible that, for a given initial alignment condition, the GRP results in a local minimum which is more distant (on the $SO(3)^N$ manifold) than the initial alignment with respect to the optimal solution, therefore the IOM has to compensate for the augmented distance. In short, we can conclude that HYB computational performance are dataset-dependent, however, the HYB cascade seem to be an effective way to reduce the computational burden of bigger datasets. Finally, we believe that computational performance can be sensibly improved

by surpassing some limitations of our current implementation. First of all, we verified that a significant portion of time for IOM is spent in computing matrices A , B and \mathcal{M} , and since some of these computations involve sparse matrices, this time could be reduced by switching from Newmat toward a matrix library that better supports sparse calculus. Moreover, correspondence selection at each iteration is parallelizable and can thus be computed on GPU hardware.

[Table 1 about here.]

[Table 2 about here.]

[Figure 9 about here.]

5. Conclusion

In this work, a new approach capable of guaranteeing accurate and robust global registration of collections of range images is presented. Based on the valuable Optimization-on-a-Manifold engine, originally proposed in [6] (OOM), a series of major algorithmic improvements and technical solutions are introduced and described in detail. In particular, the updating of correspondence points at each iteration is made feasible by a fast procedure and produces a considerable qualitative and quantitative change in the convergence properties of the algorithm. The resulting Improved Optimization-on-a-Manifold (IOM) demonstrates a significant boost in terms of both accuracy and computational efficiency with respect to the original OOM approach. When IOM is compared to the Pulli’s multiview registration technique (GRP), the benefit of using an unconstrained minimization approach on the manifold $SO(3)$ compared to heuristics-based methods for global registration is evident in terms of achievable alignment accuracy (GRP presented in all cases worse and less consistent accuracy results compared to IOM). Moreover, a robustness test was performed on both GRP and IOM evidencing the sensitivity of the former and the resilience of the latter to various misalignment scenarios. On the other hand, GRP tends to scale better in terms of resource allocation (and computation time) when the number of views increase (order of a hundred and more range images were used in our tests). Therefore, a hybrid scheme where IOM is preceded by GRP has been also explored with the idea to facilitate convergence. Even if not systematically, this scheme led to computational improvements for the most demanding

analyzed cases, where the discrepancy between GRP and IOM computation time is more evident. Nevertheless, keeping in mind that the main and fully reached objective of the present work was the improvement of the original OOM approach, future works on parallelized computing (multi-thread or GPU based) and sparse matrix exploitation (e.g. by using dedicated software libraries) are expected to yield a further boost in computational performance, while fully maintaining the achieved outperforming accuracy. What can be eventually said is that the presented IOM demonstrates to be a viable solution for the global registration of large collections (dozens of views) of dense (hundreds to million points per single view) range images in modern high-quality 3D object modeling pipelines.

[Figure 10 about here.]

References

- [1] F. Bernardini, H. Rushmeier, The 3D model acquisition pipeline, *Computer Graphics Forum* 21 (2) (2002) 149 – 172.
- [2] K. Pulli, Multiview registration for large data sets, in: 2nd International Conference on 3-D Digital Imaging and Modeling, 1999, pp. 160 – 168.
- [3] C. J. Taylor, D. J. Kriegman, Minimization on the lie group $SO(3)$ and related manifolds, Tech. Rep. 9405, Yale University (1994).
- [4] P. Y. Lee, Geometric optimization for computer vision, Ph.D. thesis, Australian National University (2005).
- [5] S. Krishnan, P. Y. Lee, J. B. Moore, S. Venkatasubramanian, Global registration of multiple 3D point sets via optimization-on-a-manifold, in: 3rd Eurographics Symposium on Geometry Processing, 2005.
- [6] S. Krishnan, P. Y. Lee, J. B. Moore, S. Venkatasubramanian, Optimisation-on-a-manifold for global registration of multiple 3D point sets, *International Journal of Intelligent Systems Technologies and Applications* 3 (2007) 319 – 340.
- [7] V. Raja, K. J. Fernandes, *Reverse Engineering: An Industrial Perspective*, Springer, 2008.

- [8] D. Breitenreicher, C. Schnörr, Model-based multiple rigid object detection and registration in unstructured range data, *International Journal of Computer Vision* 92 (2011) 32 – 52.
- [9] A. Vrabel, O. Bellon, L. Silva, A 3D reconstruction pipeline for digital preservation, in: *22nd IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 2687 – 2694.
- [10] D. Borrmann, J. Elseberg, K. Lingemann, A. Nchter, J. Hertzberg, Globally consistent 3D mapping with scan matching, *Robotics and Autonomous Systems* 56 (2) (2008) 130 – 142.
- [11] D. Rueckert, J. A. Schnabel, Medical image registration, in: *Biomedical Image Processing*, 2011, pp. 131 – 154.
- [12] F. Bonarrigo, A. Signoroni, R. Leonardi, Multi-view alignment with database of features for an improved usage of high-end 3D scanners, *EURASIP Journal on Advances in Signal Processing* 2012 (1) (2012) 148.
- [13] Y. Chen, G. Medioni, Object modelling by registration of multiple range images, *Image and Vision Computing* 10 (1992) 145 – 155.
- [14] P. Besl, H. McKay, A method for registration of 3-d shapes, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14 (2) (1992) 239 – 256.
- [15] A. J. Stoddart, S. Lemke, A. Hilton, T. Renn, Estimating pose uncertainty for surface registration, *Image and Vision Computing* 16 (2) (1998) 111 – 120.
- [16] S. Rusinkiewicz, M. Levoy, Efficient variants of the icp algorithm, in: *3rd International Conference on 3-D Digital Imaging and Modeling*, 2001, pp. 145 – 152.
- [17] G. C. Sharp, S. W. Lee, D. K. Wehe, ICP registration using invariant features, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24 (1) (2002) 90 – 102.
- [18] S.-Y. Park, M. Subbarao, An accurate and fast point-to-plane registration technique, *Pattern Recognition Letters* 24 (16) (2003) 2967 – 2976.

- [19] H. Pottmann, S. Leopoldseder, M. Hofer, Registration without icp, *Computer Vision and Image Understanding* 95 (1) (2004) 54 – 71.
- [20] N. J. Mitra, N. Gelfand, H. Pottmann, L. Guibas, Registration of point cloud data from a geometric optimization perspective, in: *2nd Eurographics Symposium on Geometry Processing, 2004*, pp. 23 – 31.
- [21] L. Ikemoto, N. Gelfand, M. Levoy, A hierarchical method for aligning warped meshes, in: *Proc. 3-D Digital Imaging and Modeling, 3DIM 2003, 2003*, pp. 434 – 441.
- [22] M. Wand, P. Jenke, Q. Huang, M. Bokeloh, L. Guibas, A. Schilling, Reconstruction of deforming geometry from time-varying point clouds, in: *Proc. Symp. on Geometric Processing, SGP '07, 2007*, pp. 49–58.
- [23] M. Kazhdan, M. Bolitho, H. Hoppe, Poisson surface reconstruction, in: *4th Eurographics Symposium on Geometry Processing, 2006*, pp. 61 – 70.
- [24] T. Masuda, N. Yokoya, A robust method for registration and segmentation of multiple range images, *Computer Vision and Image Understanding* 61 (1995) 295 – 307.
- [25] G. Blais, M. Levine, Registering multiview range data to create 3D computer objects, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 17 (8) (1995) 820 – 824.
- [26] L. Silva, O. R. P. Bellon, K. Boyer, Enhanced, robust genetic algorithms for multiview range image registration, in: *3-D Digital Imaging and Modeling, 2003. 3DIM 2003. Proceedings. Fourth International Conference on, 2003*, pp. 268–275.
- [27] R. Bergevin, M. Soucy, H. Gagnon, D. Laurendeau, Towards a general multi-view registration technique, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 18 (5) (1996) 540 – 547.
- [28] R. Benjemaa, F. Schmitt, Fast global registration of 3D sampled surfaces using a multi-z-buffer technique, *Image and Vision Computing* 17 (2) (1999) 113 – 123.

- [29] J. Williams, M. Bennamoun, Simultaneous registration of multiple corresponding point sets, *Computer Vision and Image Understanding* 81 (1) (2001) 117 – 142.
- [30] X. Pennec, Multiple registration and mean rigid shapes - application to the 3D case, in: 16th Leeds Annual Statistical Workshop, 1996, pp. 178 – 185.
- [31] A. Stoddart, A. Hilton, Registration of multiple point sets, in: 13th International Conference on Pattern Recognition, Vol. 2, 1996, pp. 40 – 44.
- [32] D. W. Eggert, A. W. Fitzgibbon, R. B. Fisher, Simultaneous registration of multiple range views for use in reverse engineering of cad models, *Computer Vision and Image Understanding* 69 (3) (1998) 253 – 272.
- [33] D. F. Huber, M. Hebert, Fully automatic registration of multiple 3D data sets, *Image and Vision Computing* 21 (7) (2003) 637 – 650.
- [34] P. J. Neugebauer, Reconstruction of real-world objects via simultaneous registration and robust combination of multiple range images, *International Journal of Shape Modeling* 3 (1997) 71 – 90.
- [35] B. K. Horn, Closed-form solution of absolute orientation using unit quaternions, *Journal of the Optical Society of America* 4 (4) (1987) 629 – 642.
- [36] R. Benjemaa, F. Schmitt, A solution for the registration of multiple 3D point sets using unit quaternions, in: 5th European Conference on Computer Vision, 1998, pp. 34 – 50.
- [37] G. Sharp, S. Lee, D. Wehe, Multiview registration of 3D scenes by minimizing error between coordinate frames, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26 (8) (2004) 1037 – 1050.
- [38] A. Torsello, E. Rodola, A. Albarelli, Multiview registration via graph diffusion of dual quaternions, in: *Computer Vision and Pattern Recognition (CVPR)*, 2011 IEEE Conference on, 2011, pp. 2441–2448.
- [39] H. Pottmann, S. Leopoldseder, M. Hofer, Simultaneous registration of multiple views of a 3D object, in: *International Society for Photogrammetry and Remote Sensing*, 2001.

- [40] H. Pottmann, Q.-X. Huang, Y.-L. Yang, S.-M. Hu, Geometry and convergence analysis of algorithms for registration of 3D shapes, *International Journal of Computer Vision* 67 (3) (2006) 277 – 296.
- [41] D. Breitenreicher, C. Schnörr, Intrinsic second-order geometric optimization for robust point set registration without correspondence, in: D. Cremers, Y. Boykov, A. Blake, F. Schmidt (Eds.), *Energy Minimization Methods in Computer Vision and Pattern Recognition*, Vol. 5681 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, 2009, pp. 274–287.
- [42] S. Cunnington, A. Stoddart, N-view point set registration: a comparison, in: *10th British Machine Vision Conference*, 1999, pp. 234 – 244.
- [43] Y. Liu, Automatic range image registration in the markov chain, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32 (1) (2010) 12 – 29.
- [44] F. Bonarrigo, A. Signoroni, An enhanced 'optimization-on-a-manifold' framework for global registration of 3D range data, in: *1st International Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission*, 2011, pp. 350 – 357.
- [45] R. Davies, NEWMAT library package, http://www.robertnz.net/nm_intro.htm.
- [46] N. Pezzotti, F. Bonarrigo, A. Signoroni, Boosting the computational performance of feature-based multiple 3d scan alignment by iat-k-means clustering, in: *3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT)*, 2012 Second International Conference on, 2012, pp. 89 –96.

List of Tables

1	Global registration results	36
2	Execution time	37

Table 1

Dataset	Avg # points/RI	#RIs	Initial dist. [μm]	Avg distance [μm]			Std deviation [μm]			
				GRP	OOM	IOM	HYB	GRP	OOM	IOM
Denture	400 k	8	21	17	17	17	5	3	4	4
Dog	120 k	14	30	19	27	17	6	7	3	3
Cherub	650 k	20	39	30	25	20	5	3	2	2
Dolphin	400 k	20	56	38	48	33	4	11	3	3
Capital	750 k	23	34	26	29	23	4	3	3	2
Rose	150 k	24	220	33	143	13	6	25	1	1
Hurricane	700 k	32	155	86	88	36	140	41	6	6
Cupid	300 k	45	330	23	—	17	3	—	2	3
Decoration	500 k	48	194	172	—	24	167	—	6	6
Venus	800 k	61	65	58	—	17	52	—	4	4
Horse	800 k	81	165	69	—	27	62	—	5	5
Crankcase	550 k	93	145	105	—	35	62	—	7	7
Shell	750 k	98	225	63	—	50	12	—	9	7
Neptune	750 k	169	395	72	—	41	58	—	9	7

Table 2

Dataset	Avg # points/RI	#RIs	Tot # points	Execution time [hh:mm:ss]			
				GRP	OOM	IOM	HYB
Denture	400 k	8	3.0 M	00:00:03	00:00:12	00:00:02	00:00:05
Dog	120 k	14	2.0 M	00:00:02	00:00:30	00:00:07	00:00:06
Cherub	650 k	20	13.0 M	00:00:07	00:00:38	00:00:31	00:00:39
Dolphin	400 k	20	8.0 M	00:00:08	00:00:49	00:00:10	00:00:18
Capital	750 k	23	17.0 M	00:00:06	00:01:37	00:00:14	00:00:19
Rose	150 k	24	3.5 M	00:00:11	00:03:22	00:00:45	00:00:49
Hurricane	700 k	32	22.5 M	00:02:21	00:14:55	00:00:29	00:02:48
Cupid	300 k	45	13.5 M	00:00:44	—	00:01:40	00:01:36
Decoration	500 k	48	24.0 M	00:00:49	—	00:01:40	00:02:35
Venus	800 k	61	49.0 M	00:01:02	—	00:01:53	00:02:15
Horse	800 k	81	65.0 M	00:00:58	—	00:05:43	00:04:49
Crankcase	550 k	93	51.0 M	00:01:47	—	00:10:50	00:16:01
Shell	750 k	98	73.5 M	00:03:49	—	00:47:04	00:13:58
Neptune	750 k	169	127.0 M	00:11:10	—	03:50:30	02:17:46

List of Figures

1	The datasets employed for global registration assessment. From upper left: Venus, Crankcase, Dog, Neptune, Dolphin, Hurricane, Cupid, Capital, Cherub, Horse, Denture, Rose, Shell, Decoration.	39
2	Impact of correspondences update	40
3	Impact of correspondences matching strategy	41
4	Different correspondences matching strategies.	42
5	Standard (left) and optimized (right) correspondences update processes	43
6	Convergence results for the two datasets at varying misalignment scenarios	44
7	Different misalignment scenarios for the Cupid dataset	45
8	Convergence results for a given misalignment scenario of Hurricane dataset	46
9	Convergence rates for eight of the considered datasets.	47
10	Comparison between datasets aligned through GRP, IOM and HYB methods: the better the colors are interpenetrated, the lower the dataset error associated.	48




Figure 1

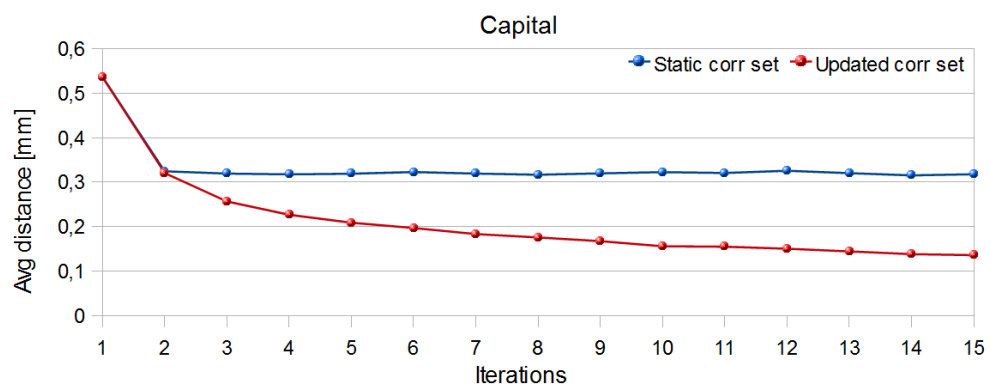


Figure 2

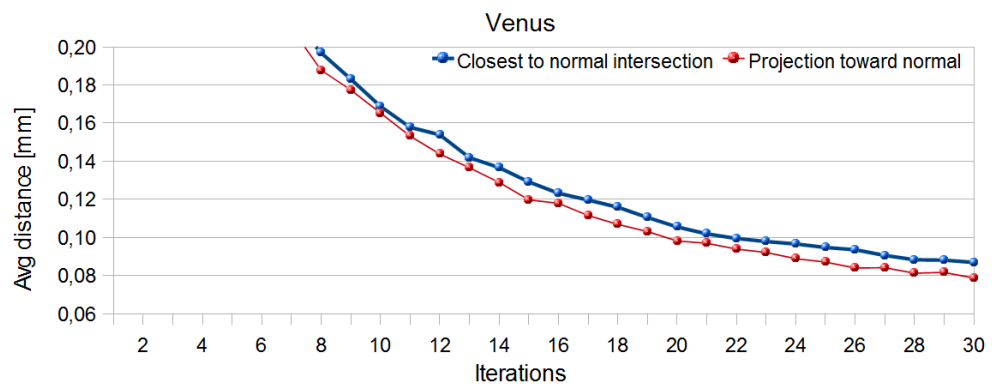


Figure 3




Figure 4

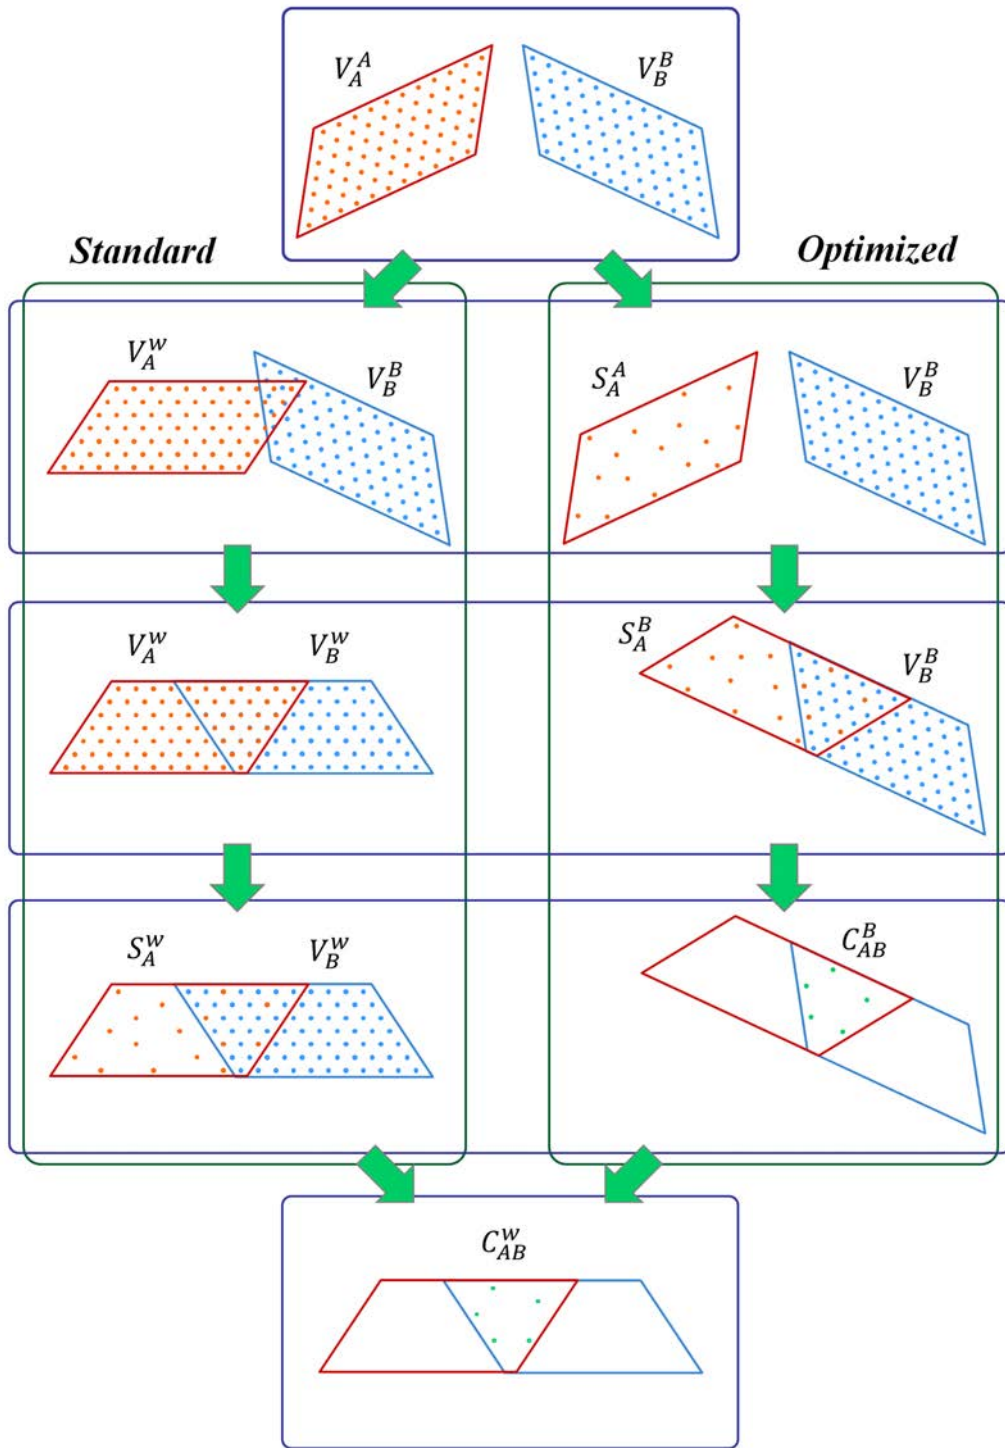


Figure 5

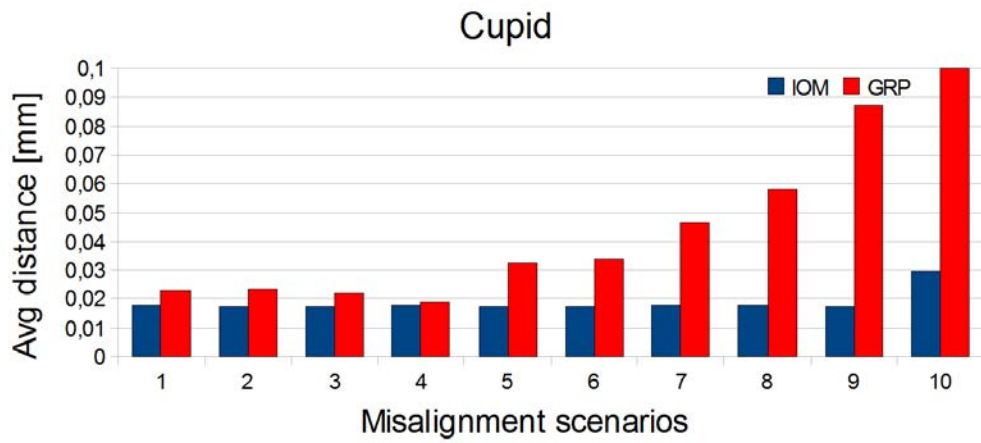
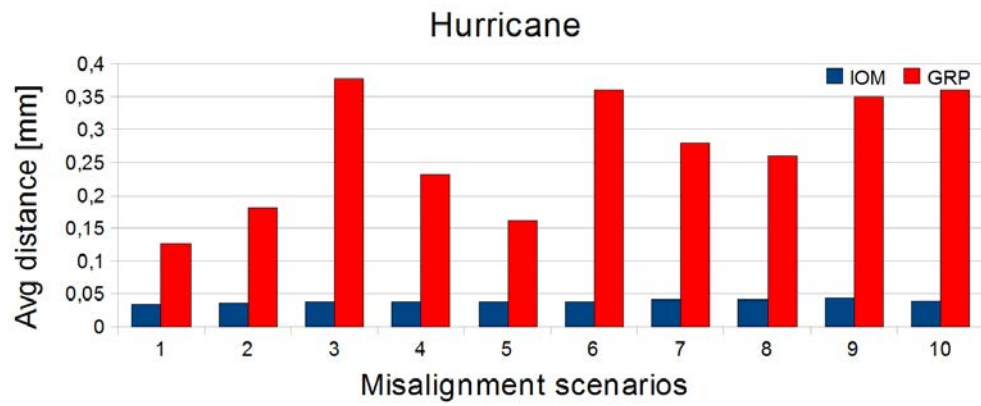



Figure 6



(a)

(b)

(c)

Figure 7

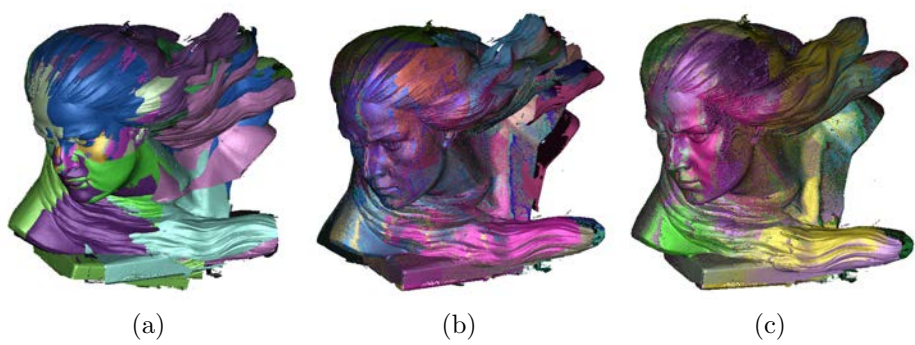


Figure 8

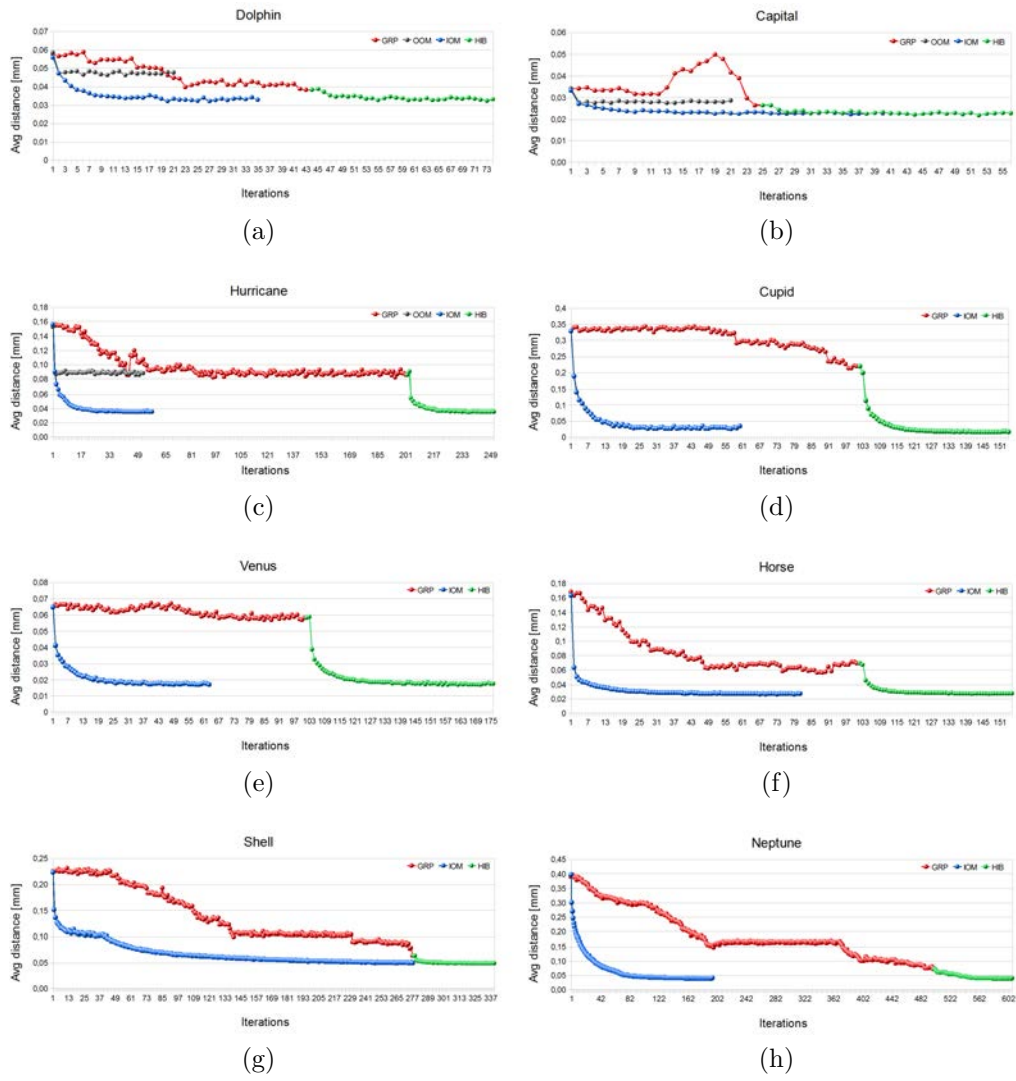


Figure 9




Figure 10