Efficient and Robust Registration on the 3D Special Euclidean Group

Uttaran Bhattacharya*

Department of Computer Science
University of Maryland
College Park, MD 20740, USA

uttaranb@cs.umd.edu

Venu Madhav Govindu[†]

Department of Electrical Engineering
Indian Institute of Science
Bengaluru 560012, INDIA

venug@iisc.ac.in

Abstract

We present a robust, fast and accurate method for registration of 3D scans. Using correspondences, our method optimizes a robust cost function on the intrinsic representation of rigid motions, i.e., the Special Euclidean group $\mathbb{SE}(3)$. We exploit the geometric properties of Lie groups as well as the robustness afforded by an iteratively reweighted least squares optimization. We also generalize our approach to a joint multiview method that simultaneously solves for the registration of a set of scans. Our approach significantly outperforms the state-of-the-art robust 3D registration method based on a line process in terms of both speed and accuracy. We show that this line process method is a special case of our principled geometric solution. Finally, we also present scenarios where global registration based on feature correspondences fails but multiview ICP based on our robust motion estimation is successful.

1. Introduction

The availability of consumer depth cameras has allowed the easy and reliable scanning of a 3D scene or object [32, 42, 51]. To build a geometrically consistent 3D reconstruction, we need to solve the key problem of aligning or registering all scans in a global frame of reference [40, 44, 20, 30]. There is a wide variety of solutions to the 3D registration problem in the literature [17, 28, 24, 11, 52]. All of these methods involve a trade-off between speed and accuracy. Recently, [52] has presented a method for fast global registration (henceforth denoted FGR) of 3D scans based on the robustness of a line process. This method has been shown to outperform existing methods in terms of both speed and accuracy.

*This work was done while the author was a student at the Indian Institute of Science. Supported in part by a research grant from Science and Engineering Research Board, Department of Science and Technology, Government of India.

[†]Corresponding author.

As in [52], given a set of 3D feature correspondences, we pose the registration problem as one of solving for the rigid Euclidean motion that minimizes a robust cost function. However, unlike their approach, our solution systematically utilises the rich geometric structure of the underlying Lie group representation for 3D motion, i.e., the Special Euclidean group $\mathbb{SE}(3)$. In this context, we achieve robustness via the iteratively reweighted least squares (IRLS) method. The key observation of our paper is that our specific combination of geometry of rigid Euclidean motion and the robustness of IRLS allows our method to significantly outperform the fast global registration method of [52] without having to take recourse to tuning of parameters. Additionally, we show that the solution proposed in [52] is a special case of our more general solution. In the process, we demonstrate that we can gain both theoretical insight as well as improved performance by utilizing the rich structure of the underlying geometric representation of rigid Euclidean motions.

Furthermore, our work also addresses two important considerations. Firstly, in order to achieve robustness, some loss functions used for optimization have parameters that need to be either provided *a priori* or estimated *in situ*. This is the case with [52] which uses the Geman-McClure loss function. We argue that for the problem of 3D registration the statistical efficiency (i.e. accuracy) trade-off inherent to all robust estimators can be addressed using loss functions that are parameter free. This obviates the need to estimate any parameter during the course of optimization. Secondly, we argue that 3D registration using 3D point feature correspondences has certain limitations. While such feature correspondences can be reliably obtained when the camera motion is small (equivalently there is significant overlap between scans), there are certain scenarios where the feature correspondences break down. For such cases, we demonstrate that accurate joint registration of multiple 3D scans can be achieved by incorporating our robust motion estimation method into local methods such as ICP.

2. Literature Survey

Although the literature of 3D registration is extensive, in the following we only focus on aspects that are directly relevant to our method. A large number of methods for registering 3D scans using point correspondences have two key aspects, (a) a method for establishing point correspondences between two 3D scans and (b) a method for estimating motion given an adequate number of such correspondences. We may further classify methods according to whether they use a fixed set of point correspondences [11, 52, 45, 21] or update them [30, 37, 2, 12, 7]. In the former instance, we use invariant feature representations to find point correspondences across a pair of scans which are then used for motion estimation. In the latter approach, we alternately update point correspondences using nearest neighbors and motion estimation till convergence, such as in the classical approach of ICP and its variants (see [37, 7, 9, 49] and references therein). Independent of the method for establishing correspondences, we need a method for robust motion estimation given a set of correspondences. This problem of robust motion estimation is the main focus of this paper.

The least squares solution to the motion estimation problem is the classical method of Umeyama [47]. However, the least squares solution breaks down in the presence of outliers. The requisite robustness is often achieved using variants of RANSAC [30, 24, 2, 35, 38, 33, 36] or motion clustering [30, 12, 39]. Other classes of approaches are based on the branch-and-bound framework [17, 22, 27, 14, 48]. However, all of these approaches often require expensive iterations and are slow to converge. Other solutions based on expectation-maximization using Gaussian mixture model representations of the scans [25, 15, 13] are similarly slow. Another approach that is relatively efficient is to perform robust averaging of pairwise motions between scans [19, 46, 3], or use variants of pose-graph optimization [11], to produce tightly registered scans. Yet another approach is to use IRLS [23] to optimize robust cost functions [1]. The recent fast global registration method of [52] made use of the duality between robust estimators and line processes [6] to develop a fast approach for global registration. This method has been shown to produce the best results till date in terms of speed as well as accuracy.

3. Lie Group Structure of Euclidean Motions

Our method utilizes the geometric structure of Lie groups, specifically that of the Special Euclidean group $\mathbb{SE}(3)$. In this section, we provide a very brief description of the group structure and relevant Lie group properties. The Euclidean motion between two 3D scans consists of a rotation and a translation. While rotations can be represented in a variety of ways, we represent them as 3×3 orthonormal matrices $\mathbf{R}$, *i.e.*, $\mathbf{R}\mathbf{R}^\top = \mathbf{I}_3$ and $|\mathbf{R}| = +1$ (here and for

the rest of this paper we use $\mathbf{I}_n$ to denote the $n \times n$ identity matrix). A rigid Euclidean motion $(\mathbf{R}, \mathbf{t})$ can then be compactly represented by a 4×4 matrix

$$\mathbf{M} = \begin{bmatrix} \mathbf{R} & | & \mathbf{t} \\ \mathbf{0} & | & 1 \end{bmatrix} \quad (1)$$

where $\mathbf{t} \in \mathbb{R}^3$. The matrices $\mathbf{R}$ and $\mathbf{M}$ satisfy the properties of a matrix group and also form smooth, differentiable manifolds, *i.e.*, they are Lie groups. Thus, $\mathbf{R} \in \mathbb{SO}(3)$ and $\mathbf{M} \in \mathbb{SE}(3)$ where $\mathbb{SO}(3)$ and $\mathbb{SE}(3)$ are the Special Orthogonal and the Special Euclidean groups respectively. We note that $\mathbf{R}$ and $\mathbf{M}$ have 3 and 6 degrees of freedom respectively.

Lie Groups: The Lie group structure of $\mathbb{SE}(3)$ plays an important role in our formulation. Detailed expositions on the geometric properties of this representation are presented in [10, 41]. For our purposes, it will suffice to note that for finite-dimensional Lie groups (matrix groups) the product and inverse operations are differentiable mappings. Every point in a Lie group has a local neighborhood (tangent space) called its Lie algebra which has the properties of a vector space. For $\mathbf{R} \in \mathbb{SO}(3)$ and $\mathbf{M} \in \mathbb{SE}(3)$, the corresponding Lie algebra are denoted as $[\omega]_\times \in \mathfrak{so}(3)$ and $\mathbf{m} \in \mathfrak{se}(3)$ respectively. Here $[\omega]_\times \in \mathfrak{so}(3)$ is the skew-symmetric form of the axis-angle rotation representation ω . In this representation, $\mathbf{R}$ represents a rotation by angle $\|\omega\|$ about the axis $\frac{\omega}{\|\omega\|}$. Further, we can move from a Lie group to its Lie algebra and vice-versa using the logarithm and the exponential mappings respectively. Thus, $\mathbf{R} = \exp([\omega]_\times)$ and $[\omega]_\times = \log(\mathbf{R})$. Similarly, we have $\mathbf{M} = \exp(\mathbf{m})$ and $\mathbf{m} = \log(\mathbf{M})$ with the forms

$$\begin{aligned} \mathbf{m} &= \log(\mathbf{M}) = \begin{bmatrix} [\omega]_\times & | & \mathbf{u} \\ \mathbf{0} & | & 0 \end{bmatrix}; \\ \mathbf{M} &= \exp(\mathbf{m}) = \sum_{k=0}^{\infty} \frac{\mathbf{m}^k}{k!} = \begin{bmatrix} \mathbf{R} & | & \mathbf{t} \\ \mathbf{0} & | & 1 \end{bmatrix}. \end{aligned} \quad (2)$$

Further, we note that the $\exp(\cdot)$ and $\log(\cdot)$ mappings for $\mathbb{SO}(3)$ and $\mathbb{SE}(3)$ have closed form expressions that can be efficiently implemented, as shown in Eqn. 7 in [18].

4. Robust Motion Estimation

Given $\mathcal{S} \geq 3$ pairs of point correspondences $\{(\mathbf{p}^s, \mathbf{q}^s) | 1 \leq s \leq \mathcal{S}\}$ between two scans, we can solve for the motion $\mathbf{M}$ required for aligning the scan pair. In the case of global approaches, where the motion can potentially be large, such point correspondences are obtained by matching geometric features such as FPFH [38] whereas in iterative schemes like ICP the correspondence of points is obtained by finding the nearest neighbor match on the second scan

for a point on the first one. Due to a variety of sources of error, the correspondence pairs are not in perfect alignment or could be grossly incorrect, *i.e.*, $e^s = \|\mathbf{p}^s - \mathbf{M}\mathbf{q}^s\| \neq 0$, where e^s denotes the norm of the discrepancy in registration for the s -th correspondence pair $(\mathbf{p}^s, \mathbf{q}^s)$. When we assume that the individual errors have a zero-mean, iid Gaussian distribution, the optimal estimate for $\mathbf{M}$ is obtained by a least squares minimization and has a closed form [47]. However, since correspondences could be highly incorrect in practice, instead of a least squares formulation, we pose motion estimation as the optimization of a robust cost function

$$\min_{\mathbf{M} \in \mathbb{SE}(3)} \mathcal{C}(\mathbf{M}) = \sum_{s=1}^S \rho(\|\mathbf{p}^s - \mathbf{M}\mathbf{q}^s\|) = \sum_{s=1}^S \rho(e^s(\mathbf{M})) \quad (3)$$

where $\rho(\cdot)$ is a robust loss function. We also note that the individual error terms $e^s(\mathbf{M})$ are a function of the motion $\mathbf{M}$. The use of robust estimators is well studied in statistics as an M-estimation problem and has been successfully used in a variety of vision problems [29]. However, in addition to robustness, in Eqn. 3 we require our solution to satisfy the geometric constraints for $\mathbf{M} \in \mathbb{SE}(3)$. These requirements of robust estimation under geometric constraints are satisfied by our solution.

4.1. Proposed Solution for Pairwise Registration

We propose to minimize the cost function $\mathcal{C}(\mathbf{M})$ in Eqn. 3 in an iterative fashion. Let the estimate for $\mathbf{M}$ at the $(k-1)$ -th iteration be denoted $\mathbf{M}(k-1)$. In the k -th iteration, let us update the motion matrix by $\Delta\mathbf{M}(k)$, *i.e.*, $\mathbf{M}(k) = \Delta\mathbf{M}(k)\mathbf{M}(k-1)$. Noting that here the matrix multiplication is not commutative, our formulation uses a left-invariant metric on $\mathbb{SE}(3)$ [34, 50]. Using a first-order approximation for the motion update matrix $\Delta\mathbf{M}(k)$, we have

$$\Delta\mathbf{M}(k) \approx \mathbf{I}_4 + \Delta\mathbf{m}(k) \quad (4)$$

$$\begin{aligned} \Rightarrow \mathcal{C}(\mathbf{M}(k)) & \left(= \sum_{s=1}^S \rho(\|\mathbf{p}^s - \mathbf{M}(k)\mathbf{q}^s\|) \right) \\ & = \sum_{s=1}^S \rho(\|\mathbf{p}^s - (\mathbf{I}_4 + \Delta\mathbf{m}(k))\mathbf{M}(k-1)\mathbf{q}^s\|). \end{aligned} \quad (5)$$

The Lie algebra matrix $\Delta\mathbf{m}(k)$ encodes the 6 parameters that we need to estimate for the update $\Delta\mathbf{M}(k)$. We can obtain the vector representation for these 6 parameters using the ‘vee’ operator, *i.e.*, $\mathbf{v} = \Delta\mathbf{m}^\vee = [\boldsymbol{\omega} \ \mathbf{u}]^\top$. Since we use a first-order approximation in Eqn. 4, the cost $\mathcal{C}(\mathbf{M}(k))$ is linear in $\Delta\mathbf{m}(k)$. We equivalently note that it is also linear in $\mathbf{v}(k)$. Thus, we rewrite the individual error terms $e^s(\mathbf{M}(k))$ as

$$\begin{aligned} e^s(\mathbf{M}(k)) & = \|\mathbf{p}^s - (\mathbf{I}_4 + \Delta\mathbf{m}(k))\mathbf{M}(k-1)\mathbf{q}^s\| \\ & = \|\mathbf{A}^s \mathbf{v} - \mathbf{b}^s\| \end{aligned} \quad (6)$$

where $\mathbf{A}^s$ and $\mathbf{b}^s$ are the appropriate matrices. The derivation of the explicit forms of $\mathbf{A}^s$ and $\mathbf{b}^s$ are given in Section A in the Appendix. To obtain the update in the k -th iteration, we now optimize the cost $\mathcal{C}(\mathbf{M}(k)) = \sum_{s=1}^S \rho(e^s(\mathbf{v}))$ w.r.t. $\mathbf{v}$, and get,

$$\frac{\rho'(e^s)}{e^s} (\mathbf{A}^s)^\top \mathbf{A}^s \mathbf{v} = \frac{\rho'(e^s)}{e^s} (\mathbf{A}^s)^\top \mathbf{b}^s \quad (7)$$

for each summand indexed by s , where $\rho'(\cdot) = \frac{\partial \rho}{\partial e}$ is the influence function of the robust loss $\rho(\cdot)$. We may further denote $w^s = \frac{\rho'(e^s)}{e^s}$, which is the relative weight accorded to the s -th equation in Eqn. 7. Collecting all such relationships obtained for each pair of correspondences $(\mathbf{p}^s, \mathbf{q}^s)$ into a single system of equations we have

$$\mathbf{A}^\top \mathbf{W} \mathbf{A} \mathbf{v} = \mathbf{A}^\top \mathbf{W} \mathbf{b} \quad (8)$$

where $\mathbf{A} = \begin{bmatrix} \mathbf{A}^1 \\ \vdots \\ \mathbf{A}^S \end{bmatrix}$, $\mathbf{b} = \begin{bmatrix} \mathbf{b}^1 \\ \vdots \\ \mathbf{b}^S \end{bmatrix}$, and $\mathbf{W} = \text{diag}(w^1 \mathbf{I}_3, \dots, w^S \mathbf{I}_3)$.

Eqn. 8 is a weighted linear system of equations with the solution $\mathbf{v} = (\mathbf{A}^\top \mathbf{W} \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{W} \mathbf{b}$. However, it should be noted that each individual weight w^s is a function of the error e^s which, in turn, is dependent on $\mathbf{v}$ since $e^s = \|\mathbf{A}^s \mathbf{v} - \mathbf{b}^s\|$. Thus, in Eqn. 8, the equivalent relationship is $\mathbf{A}^\top \mathbf{W}(\mathbf{v}) \mathbf{A} \mathbf{v} = \mathbf{A}^\top \mathbf{W}(\mathbf{v}) \mathbf{b}$. The solution for this system of equations is the well-known iteratively reweighted least squares (IRLS) method [23, 43]. In the IRLS method, in each iteration the weights w^s are estimated based on the current estimate of $\mathbf{v}$. Given these weights, $\mathbf{v}$ is re-estimated. This process is repeated till convergence.

Given a solution for $\mathbf{v}$, we can estimate $\Delta\mathbf{M}(k) = \exp(\hat{\mathbf{v}}(k))$ where the ‘hat’ operator $\hat{\mathbf{v}}$ converts the estimated Lie algebra parameters $\mathbf{v}$ into its equivalent matrix representation $\Delta\mathbf{m}(k)$. We emphasize here that although in Eqn. 4, we assumed a first-order approximation, we map the estimated $\mathbf{v}(k)$ into an intrinsic estimate of the motion update, *i.e.*, $\Delta\mathbf{M}(k) = \exp(\hat{\mathbf{v}}(k)) \in \mathbb{SE}(3)$. In other words, a first-order approximation in an intermediate step does not mean that the actual cost function is approximated. The mapping $\Delta\mathbf{m}(k) = \exp(\hat{\mathbf{v}}(k))$ ensures that the estimated motion $\mathbf{M}(k)$ is always a valid member of the $\mathbb{SE}(3)$ group. We now state our solution for robust motion estimation in Algorithm 1.

Algorithm 1 IRLS estimation of pairwise 3D registration

Input: $\{(\mathbf{p}^1, \mathbf{q}^1) \cdots (\mathbf{p}^S, \mathbf{q}^S)\}$ (S correspondences across a pairs of scans)

Output: $\mathbf{M} \in \mathbb{SE}(3)$ (Robust motion estimate)

Initialization: $\mathbf{M} = \mathbf{I}_4$

```

while  $\|\mathbf{v}\| > \epsilon$  do
  1. Compute  $\{(\mathbf{A}^s, \mathbf{b}^s) | \forall s \in [1 \cdots S]\}$  using Eqn. 6
  2. Compute weights  $w^s = \frac{\rho'(e^s)}{e^s}$  as defined by Eqn. 7
  3. Estimate  $\mathbf{v}$  as IRLS solution for Eqn. 8
  4. Update  $\mathbf{M} \leftarrow \exp(\hat{\mathbf{v}})\mathbf{M}$ 
end while

```

Algorithm 1 is an iterative algorithm with a nested iteration. The outer loop is defined by the `while` statement and we denote its number of iterations as K_{outer} . The inner loop consists of the IRLS step of line 3 in Algorithm 1 since IRLS is itself an iterative method. We denote the number of iterations of the IRLS step as K_{IRLS} .

4.2. Extension to Joint Multiview Registration

In Algorithm 1 we presented the registration solution for two scans. This approach can be extended to the simultaneous or joint multiview registration of a set of scans. Towards this end, we define a viewgraph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ where $v_i \in \mathcal{V}$ represents the Euclidean motion of the i -th scan (equivalently camera) and an edge $(i, j) \in \mathcal{E}$ signifies that the relative Euclidean motion between the i -th and the j -th scans can be determined from available matches. Further, we denote the number of scans as $N = |\mathcal{V}|$. We may now define the cost function to be optimized for joint multiview registration as follows

$$\begin{aligned}
 \mathcal{C}(\mathbb{M}) &= \sum_{(i,j) \in \mathcal{E}} \sum_{s=1}^{S_{ij}} \rho(\|\mathbf{M}_i \mathbf{p}_i^s - \mathbf{M}_j \mathbf{p}_j^s\|) \\
 &= \sum_{(i,j) \in \mathcal{E}} \sum_{s=1}^{S_{ij}} \rho(e_{ij}^s(\mathbb{M}))
 \end{aligned} \tag{9}$$

where $\mathbb{M} = \{\mathbf{M}_1 \cdots \mathbf{M}_N\}$ denotes the set of absolute motions of each of the N scans w.r.t. to a global frame of reference and S_{ij} is the number of correspondences between the i -th and the j -th scans. We again use an iterative approach to minimize the cost function $\mathcal{C}(\mathbb{M})$ in Eqn. 9 w.r.t. $\mathbb{M}$. In the k -th iteration, we update each motion matrix $\mathbf{M}_i$ by $\Delta \mathbf{M}_i(k)$, i.e., $\mathbf{M}_i(k) = \Delta \mathbf{M}_i(k) \mathbf{M}_i(k-1) \forall i \in [1 \cdots N]$. Using a first-order approximation for each update matrix $\Delta \mathbf{M}_i(k)$, we have

$$e_{ij}^s(\mathbb{M}(k)) = \|(\mathbf{I}_4 + \Delta \mathbf{m}_i(k)) \mathbf{M}_i(k-1) \mathbf{p}_i^s - (\mathbf{I}_4 + \Delta \mathbf{m}_j(k)) \mathbf{M}_j(k-1) \mathbf{p}_j^s\| \tag{10}$$

$$= \|\mathbf{A}_{ij}^s \mathbf{v} - \mathbf{b}_{ij}^s\| \tag{11}$$

where $\mathbf{v} = [\mathbf{v}_1 \cdots \mathbf{v}_N]^\top$ collates the corresponding vector representations of each of the Lie algebra matrices $\Delta \mathbf{m}_i(k)$, and $\mathbf{A}_{ij}^s$ and $\mathbf{b}_{ij}^s$ are constructed analogous to Eqn. 6. The subsequent update in the k -th iteration is then analogously obtained from the relation

$$\mathbf{A}^\top \mathbb{W}(\mathbf{v}) \mathbf{A} \mathbf{v} = \mathbf{A}^\top \mathbb{W}(\mathbf{v}) \mathbf{b} \tag{12}$$

where $\mathbf{A}$, $\mathbf{b}$ and $\mathbb{W}(\mathbf{v})$ correspondingly collate all $\mathbf{A}_{ij}^s$, $\mathbf{b}_{ij}^s$ and $\mathbf{w}_{ij}^s = \frac{\rho'(e_{ij}^s)}{e_{ij}^s}$. As earlier, the solution for the system of equations in Eqn. 12 is the IRLS method, where we estimate the weights w_{ij}^s and the collated vector representation $\mathbf{v}$ in alternating iterations till convergence.

Given a solution for $\mathbf{v}$, we can estimate, for each $i \in [1 \cdots N]$, $\Delta \mathbf{M}_i(k) = \exp(\hat{\mathbf{v}}_i(k))$ and thereby update each member of $\mathbb{M}(k)$ as $\mathbf{M}_i(k) \leftarrow \Delta \mathbf{M}_i(k) \mathbf{M}_i(k-1)$. It should also be noted that in multiview registration, the choice of the global frame of reference for the set $\mathbb{M}$ is arbitrary. For our implementation, we fix it to the first scan, i.e., we set $\mathbf{M}_1 = \mathbf{I}_4$ and do not update $\mathbf{M}_1$ throughout course of optimization. The detailed algorithm is presented as Algorithm B.1 in the Appendix.

5. Results

We test our algorithm on 3 different choices of the loss function $\rho(\cdot)$, namely $L_{\frac{1}{2}}$: $\rho(x) = \sqrt{|x|}$, L_1 : $\rho(x) = |x|$, and scaled Geman-McClure: $\rho(x) = \frac{\mu x^2}{\mu + x^2}$, where μ is the scale factor. The FGR approach of [52] uses only the scaled Geman-McClure loss function and anneals μ in fixed steps per iteration. To enable comparison, we anneal μ in an identical manner in our tests with the scaled Geman-McClure loss function. We present results on both pairwise and multiview registration tests. For the pairwise and multiview tests, we terminate using $\epsilon = 10^{-5}$ and $\epsilon = 10^{-7}$ respectively. All our reported running times are measured on a single thread on an Intel Xeon(R) E5-2650 v3 processor clocked at 2.30 GHz.

5.1. Pairwise Registration

We present the performance of our proposed pairwise registration on the synthetic range data provided by [52] and on the 4 indoor sequences in the Augmented ICL-NUIM dataset provided by Choi et al. [11]. In all our pairwise tests, we use $K_{IRLS} = 2$. We compare the registration errors given by the 3 versions of our method with the following prior methods: Super4PCS [30], GoICP [48], Choi et al. [11], FGR [52] and DARE [26] (using the hyperparameters suggested by the authors). Registration errors comprise of statistical measures on the distances between the ground-truth point correspondences between pairs of scans post alignment. We also report the trajectory errors, which include statistical measures on both the rotation errors and

Table 1: Median rotation angle error (Md.RAE) (in degrees), median translation norm error (Md.TNE), mean RMSE (Mn.RMSE) and maximal RMSE (Mx.RMSE) (all in units of the surface diameters) achieved by each method for each noise level σ on the synthetic range datasets.

Method	$\sigma = 0.0025$			
	Md.RAE	Md.TNE	Mn.RMSE	Mx.RMSE
Super4PCS [30]	0.864	0.008	0.014	0.029
GoICP [48]	1.207	0.011	0.032	0.133
Choi et al. [11]	0.778	0.006	0.008	0.022
FGR [52]	0.749	0.005	0.004	0.011
DARE [26]	1.851	0.013	0.035	0.176
Our $L_{\frac{1}{2}}$	0.545	0.004	0.004	0.011
Our L_1	0.566	0.004	0.004	0.011
Our GM	0.725	0.005	0.004	0.011

Method	$\sigma = 0.005$			
	Md.RAE	Md.TNE	Mn.RMSE	Mx.RMSE
Super4PCS [30]	1.468	0.012	0.017	0.095
GoICP [48]	1.736	0.019	0.037	0.127
Choi et al. [11]	1.533	0.015	0.035	0.274
FGR [52]	1.146	0.008	0.006	0.017
DARE [26]	2.005	0.025	0.054	0.312
Our $L_{\frac{1}{2}}$	0.959	0.008	0.006	0.017
Our L_1	1.516	0.011	0.007	0.017
Our GM	1.146	0.008	0.006	0.017

the translation norm errors between the pairs of cameras (corresponding to the given scans) w.r.t. the ground-truth camera pair, for all the methods.

Synthetic range dataset: We perform the set of controlled experiments described in [52] on each of their 5 synthetic datasets, at the given Gaussian noise levels $\sigma = 0.0025$ and $\sigma = 0.005$ (for each model, σ is in units of the surface diameter). Adding Gaussian noise of this scale to a pair of scans in the synthetic range dataset is sufficient to introduce outliers in the point correspondences that are computed between that pair. Since depth cameras produce noisy scans in practice, this is a realistic way of increasing the outlier percentage in the point correspondences for synthetic scans. Table 1 lists the mean and maximal RMSE on the aligned ground truth correspondences for each method and for each noise level. For both noise levels, our method with the $L_{\frac{1}{2}}$ loss function attains the lowest registration error together with FGR. Table 2 reports the mean running time of the motion step of each method for each of the 5 models in the dataset. The $L_{\frac{1}{2}}$ of our method is more than $3\times$ faster than all prior methods on the average, and the L_1 version is more than $5\times$ faster. We present more pairwise results on synthetic datasets of [31] in Section C in the Appendix.

Augmented ICL-NUIM dataset: Each of the 4 sequences in the Augmented ICL-NUIM dataset provided by Choi et al. [11] consist of 2350 to 2870 scans of an indoor scene. Moreover, the given scans are provided in a smooth temporal sequence, *i.e.*, pairs of scans with proximity in time-tamp also have sufficient view overlap with each other. This, in turn, leads to reliable FPFH feature matches be-

Table 2: Running time (in milliseconds) of the motion step of each method for each model in the synthetic range dataset ($K = \times 1,000$).

Dataset	Mean # points	Super4PCS [30]	GoICP [48]	Choi et al. [11]	FGR [52]	DARE [26]	Our $L_{\frac{1}{2}}$	Our L_1	Our GM
Bimba	9,416	16,230	1,550	650	11.9	920	3.9	2.5	5.5
Child'n	11,148	18,410	1,620	890	16.8	960	4.8	3.3	8.0
Dragon	11,232	20,520	1,840	970	17.6	1,090	5.0	3.4	8.2
Angel	12,072	29,640	3,000	1,090	19.1	1,770	5.3	3.8	8.9
Bunny	13,357	38,470	5,530	1,170	21.6	3,310	7.4	5.1	9.9
Mean	11,445	24,650	2,710	960	17.4	1,610	5.3	3.6	8.1

tween such pairs. We therefore tested the performance of all the methods for all pairs of scans (i, j) in each sequence such that $|i - j| \leq 10$. Table 3 lists the results on each dataset for the various methods under consideration. For each dataset and corresponding method, we list the median rotation angle error and the median translation norm error of the recovered pairwise camera motions, as well as the mean computation time on the pairs. The $L_{\frac{1}{2}}$ version of our method performs the best in terms of the trajectory error statistics. It can also be seen to be significantly faster than the FGR method of [52].

5.2. Joint Multiview Registration

We present the performance of our joint multiview registration algorithm on the 4 sequences in the Augmented ICL-NUIM dataset, specifically, on the 47 to 57 scene fragments that were provided for each sequence by Choi et al. [11]. We use $K_{IRLS} = 3$ for multiview registration, as the joint optimization variable and its corresponding search space are both large ($\mathbf{v} \in \mathbb{R}^{6(N-1)}$ in Eqn. 12 for N cameras). First, we compute pairwise motion estimates between fragments followed by a robust motion averaging step on $\mathbb{SE}(3)$ that is similar to the one used for rotation averaging in [8]. The output of this two-stage approach is used to initialize the joint multiview optimization in Eqn. 9. The main drawback of only using the two-step approach for global registration is that it is not a true global method. It only considers local point correspondences, and then averages out the errors made by the pairwise motion estimates. Conversely, the joint multiview approach deals with point correspondences in the global setting and solves for the global cost function. The relative improvement in reconstruction error gained by using the joint multiview approach on top of the two-stage approach is shown in Table C.2 in the Appendix.

We compare our results with those of Choi et al. [11], FGR [52] and Super4PCS [30] as the best-performing RANSAC variant. While we are aware of other approaches to multiview registration including closed form approaches, we omit them from our comparisons because they failed to provide a solution in the large-scale dataset for multiview registration we have used. For example, Bartoli et al. [4] and Bergström et al. [5], among others, compute transformations using a closed form SVD solution that do not scale to large-scale data. Approaches such as that of Fitzgibbon

Table 3: Median rotation angle error (RAE) (in degrees), median translation norm error (TNE) (in meters) and mean running time (in milliseconds) of the motion step of each method for each sequence in the Augmented ICL-NUIM dataset.

Method	livingroom 1			livingroom 2			office 1			office 2		
	Md.RAE	Md.TNE	Mn.Time	Md.RAE	Md.TNE	Md.Time	Md.RAE	Md.TNE	Mn.Time	Md.RAE	Md.TNE	Mn.Time
Super4PCS [30]	1.104	0.039	368,030	0.616	0.033	344,720	0.932	0.038	367,980	0.844	0.027	345,460
GoICP [48]	1.336	0.071	35,110	0.992	0.058	33,420	1.365	0.066	34,450	1.104	0.047	32,530
Choi et al. [11]	0.941	0.041	14,740	0.551	0.031	13,850	0.811	0.036	14,720	0.765	0.029	13,990
FGR [52]	0.793	0.029	272	0.482	0.021	181	0.707	0.020	272	0.669	0.016	177
DARE [26]	1.305	0.044	21,500	1.172	0.059	20,320	1.716	0.037	21,110	1.286	0.068	20,920
Our $L_{\frac{1}{2}}$	0.595	0.023	61	0.380	0.017	50	0.474	0.014	59	0.437	0.011	45
Our L_1	0.964	0.025	33	0.419	0.019	27	0.569	0.017	33	0.524	0.013	25
Our GM	0.793	0.029	118	0.482	0.021	87	0.707	0.020	118	0.669	0.016	89

Table 4: Mean registration error (MRE) (in meters) and running time (in seconds) for each method for full reconstruction from the fragments of each sequence in the Augmented ICL-NUIM dataset.

Method	livingroom 1		livingroom 2	
	MRE	Time	MRE	Time
Super4PCS [30]	0.13	221,160	0.14	82,640
Choi et al. [11]	0.04	8,940	0.07	3,360
FGR [52]	0.05	131	0.06	81
Our $L_{\frac{1}{2}}$	0.04	71	0.05	49
Our L_1	0.07	62	0.09	40
Our GM	0.05	88	0.06	70

Method	office 1		office 2	
	MRE	Time	MRE	Time
Super4PCS [30]	0.09	112,430	0.11	100,720
Choi et al. [11]	0.03	4,500	0.04	4,080
FGR [52]	0.03	69	0.05	48
Our $L_{\frac{1}{2}}$	0.03	42	0.04	32
Our L_1	0.04	36	0.06	28
Our GM	0.03	55	0.05	41

et al. [16] use an LM-based approach, which are slow and do not exploit the geometry of the problem.

Table 4 lists the mean registration error from the ground truth surface achieved by each compared method on each sequence as well as the time taken to complete execution. For estimating the mean registration error, we use CloudCompare from <http://www.cloudcompare.org>.

Once again, the $L_{\frac{1}{2}}$ version of our method performs the best overall in terms of registration error and is significantly faster than the FGR method of [52]. Also our L_1 method is the fastest amongst all methods with a slight drop in accuracy compared to our $L_{\frac{1}{2}}$ method. Finally, Figure 1 shows a complete reconstruction of the sequence livingroom 2 from the Augmented ICL-NUIM dataset produced by the $L_{\frac{1}{2}}$ version of our multiview method. A full-blown version of this figure, as well as reconstructions of other sequences, are shown in Figures C.1, C.2, C.3 and C.4 in the Appendix.

6. Discussion

As we demonstrated in Section 5, our motion estimation method in Algorithm 1 is both fast and accurate. More

Figure 1: Reconstruction of the livingroom 2 sequence from the Augmented ICL-NUIM dataset, as given by our method with the $L_{\frac{1}{2}}$ loss function

specifically, our method outperforms the state-of-the-art FGR method of [52] in terms of both speed and accuracy. Given their strong similarities, in Section 6.1 we examine the relationship of the method of [52] and our approach in Algorithm 1. Subsequently we discuss the limitations of using FPFH for feature matching and our approach to overcome those in Section 6.2.

6.1. Comparison with FGR [52]

In [52], the cost function to be minimized is the same as that of Eqn. 3. However, for minimizing this cost function, they use a line process optimization. Originally developed in [6] for modelling discontinuities, a line process optimization can be shown to be equivalent to optimizing a robust estimator. Recalling that $e^s = \|\mathbf{p}^s - \mathbf{M}\mathbf{q}^s\|$, we define a cost function

$$E(\mathbf{M}, \mathbf{L}) = \sum_{s=1}^S (e^s)^2 l^s + \psi(l^s) \quad (13)$$

where $\mathbf{L} = [l^1 \dots l^S]$ is the collection of line processes l^s for each correspondence pair $(\mathbf{p}^s, \mathbf{q}^s)$ and $\psi(\cdot)$ is a penalty term for each line process. Here $\psi(l)$ is a monotonically decreasing function designed such that when $l = 0$, $\psi(l)$ is

a fixed non-negative constant and when $l = 1$, $\psi(l) = 0$. Thus, varying the line process l in the interval $[0, 1]$ allows us to move between a least-squares and a robust regime. The authors of [6] have shown that for every choice of loss function $\rho(\cdot)$, there is an equivalent $\psi(\cdot)$ such that minimization in Eqn. 13 yields the same solution as that of Eqn. 3. The FGR method of [52] utilizes the robustness of the line process to estimate the desired $\mathbf{M} \in \mathbb{SE}(3)$. Using a first-order approximation of the motion update, [52] arrives at a Gauss-Newton solution (Eqn. 8 of [52]). It can be easily shown that this solution is identical to solving the system of equations in Eqn. 8 in our notation. In other words, while solving for the update step $\mathbf{v}(k)$, the FGR method of [52] implicitly carries out a **single** iteration of our IRLS step in line 3 of Algorithm 1. By contrast, we carry out $K_{IRLS} > 1$ iterations of the IRLS step to achieve better convergence.

Although the difference between our method and that of [52] is only in the number of iterations of the inner IRLS step, its implication is significant. If we were solving a single IRLS problem in a vector space setting, this difference would have been immaterial. However, we note that in both of our approaches, the linear solution for the updates $\mathbf{v}(k)$ are interleaved with the non-linear motion updates $\mathbf{M} \leftarrow \exp(\hat{\mathbf{v}})\mathbf{M}$, resulting in significantly different trajectories of the solution $\mathbf{M}(k)$ on the $\mathbb{SE}(3)$ group. Specifically, in our case, by iterating the IRLS step to convergence we obtain the best possible estimate of $\mathbf{v}(k)$ in each intermediate step which, in turn, results in the best improvement of the estimate $\mathbf{M}(k)$ (equivalently the most reduction in the cost function $\mathcal{C}(\cdot)$ in the k -th step).

Another noteworthy difference between the two methods is the choice of the parametrization of the motion representation. We use the geometrically correct form of $\Delta\mathbf{m}(k)$ in Eqn. 4, *i.e.*, $\mathbf{v} = [\boldsymbol{\omega} \quad \mathbf{u}]^\top$ for our update step. However, for their update step, the authors of [52] use an extrinsic form of motion parametrization, *i.e.*, $[\boldsymbol{\omega} \quad \mathbf{t}]^\top$. While our parametrization is geometrically consistent with Lie group theory, we can recognize from Eqn. 7 in [18] that the choice in [52] is approximately close to our representation for small motion, *i.e.*, $\mathbf{u} \rightarrow \mathbf{t}$ if and only if $\theta \rightarrow 0$ in that equation. Conversely, for sufficiently large θ , the approximate representation $[\boldsymbol{\omega} \quad \mathbf{t}]^\top$ is notably different from the exact representation $[\boldsymbol{\omega} \quad \mathbf{u}]^\top$ of the $\mathfrak{se}(3)$ form. Therefore the improvement per iteration in the cost function for the method of [52] is lower compared to our single iteration IRLS form. The result is that the method of [52] has slower convergence.

We highlight both these differences for a pair of scans from the Augmented ICL-NUIM dataset in Figure 2. For the purpose of illustration, we consider the $L_{\frac{1}{2}}$ loss function for the line process optimization routine proposed in [52] as well as for our optimization routine. We do not consider

Figure 2: Comparison of the line process solution with our method for different K_{IRLS} . For ease of visualization, we show performance only for the iterations between 2 and 10.

Table 5: Number of iterations K_{outer} taken by each method to reach each convergence criterion ϵ

ϵ	LP [52]	Ours, $K_{IRLS} =$		
		1	2	3
10^{-1}	2	2	2	2
10^{-2}	5	5	4	4
10^{-3}	13	13	8	6
10^{-4}	21	18	11	8
10^{-5}	27	24	13	10
10^{-6}	33	32	19	14
10^{-7}	63	46	25	18

the Geman-McClure loss function since it has a scale factor which, in practice, is initialized at a large value and has to be progressively annealed during optimization. In other words, we use $L_{\frac{1}{2}}$ because it does not alter the fundamental properties of the two optimization routines, and at the same time lends itself to a clean illustration of our argument. We also note that in Figure 2, we represent the cost $\mathcal{C}(\mathbf{M})$ as a function of the number of iterations K_{outer} , and that it is depicted on a $\log_{10}$ scale. For ease of visualization, we only show the plot for the iteration range $[2, 10]$. Alongside the figure, Table 5 reports the number of iterations K_{outer} taken by each method to reach different convergence criteria specified by ϵ .

We observe from Table 5 that, firstly, even though the line process is conceptually equivalent to our procedure with a single IRLS step, the proposed optimization routine of [52] takes more iterations than our actual procedure with a single IRLS step to converge to the same convergence criterion. This is because of the extrinsic (approximate) parametrization $[\boldsymbol{\omega} \quad \mathbf{t}]^\top$ used in [52] as opposed to the correct Lie algebraic representation $[\boldsymbol{\omega} \quad \mathbf{u}]^\top$. Secondly, it is clear from both Figure 2 and Table 5 that the cost function converges in progressively fewer iterations as we increase the number of IRLS steps. However, it should be noted that increasing the number of IRLS steps makes each iteration of our optimization routine slower as well. Therefore, a

balance has to be struck between the speed of each iteration and the number of iterations required for convergence. In practice, we have found that using 2 IRLS steps per iteration for pairwise registration and 3 IRLS steps per iteration for joint multiview registration yield our desired results. In any event, the key observation is that the single iteration of the FGR is insufficient and yields poorer convergence compared to our formulation.

Choice of Loss Function: The choice of the loss function $\rho(\cdot)$ to be used is a critical factor in our estimation procedure. All loss functions achieve robustness to outliers by a trade-off with statistical efficiency (*i.e.* accuracy). In practice, the accuracy achievable by a chosen loss function depends on the empirical nature of the error distribution. As discussed earlier, in [52] the authors use the Geman-McClure loss function $\rho(x) = \frac{\mu x^2}{\mu + x^2}$. Here the performance critically depends on choosing a good value for μ that reflects the outlier distribution inherent to the data used. In [52], the authors start with a large μ and progressively reduce it in fixed amounts with every iteration. However, if the scale of the scans in the dataset varies significantly, such a fixed annealing procedure may not produce the best possible results. To overcome this limitation of the Geman-McClure method, we have also tested for L_1 and $L_{\frac{1}{2}}$ as shown in Section 5, and demonstrated that the latter provides the best performance. Apart from improved performance, an important desirable property of using $L_{\frac{1}{2}}$ is that it is entirely parameter free, hence we do not need to follow any additional annealing procedure. We could conceivably further improve performance by determining the optimal exponent p in $L_p(\cdot) = \|\cdot\|_2^p$ for $0 < p < 1$. However, for a broad range of possible error distributions, we find that $L_{\frac{1}{2}}$ is adequate.

6.2. Limitation of Feature-Based Registration

In Section 5, we have demonstrated the potency of the FPFH feature-matching based registration paradigm. However, we note that these experiments, derived from those presented in [52], have special data characteristics. Specifically, in these examples, either the datasets are small or the motions between adjacent scans are small in magnitude and exhibit temporal smoothness. However, we note that registration using correspondences cannot work in all scenarios. For example, when scans in the input dataset have sufficiently large depth differences or large motions between them in both rotation and translation, FPFH feature-based matches become few and unreliable. Consequently, the reconstructions given by both FGR [52] and our proposed algorithms become incorrect. In such a scenario, we need to take recourse to using a robust ICP-based multiview registration method (albeit with a greater computational cost) which converges to the correct solution.

In this alternative approach, following [19], we again

Figure 3: 3D reconstruction of a statue of Mahatma Gandhi. The close-up on the top left shows that joint multiview registration using FPFH features fails whereas the close-up on the bottom left shows successful registration using our robust pairwise motion estimation within a multiview ICP routine. The full reconstruction is shown on the right.

consider the camera viewgraph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$. For each edge available in the set $\mathcal{E}$, we estimate the pairwise motion using a robust version of ICP. Specifically, for the motion estimation step in our robust ICP, we use the motion obtained using the $L_{\frac{1}{2}}$ loss optimization for pairwise registration method as described in Algorithm 1. We estimate the absolute motion of all cameras by motion averaging of all pairwise motions on edges $\mathcal{E}$. Here, our solution for robust motion averaging for rigid Euclidean motions is similar to the robust rotation averaging method of [8]. We find that typically, we achieve the desired convergence in 3 full iterations of this procedure. To illustrate our argument, we show the qualities of reconstructions achieved for a life-size statue of Mahatma Gandhi located at Sabarmati Ashram in Ahmedabad, India, where the input set of scans is small and mostly have low overlap. Moreover, overlapping scans have significant depth differences between them, leading to significantly different FPFH features and consequently, a high percentage of incorrect matches. A visual representation of this scenario is shown in Figure D.2 in the Appendix for better understanding. As we can see in Figure 3, the joint multiview registration using FPFH features-based matches fails to correctly align some of the scans, whereas the robust multiview-based ICP routine successfully produces a correct reconstruction. A larger version of Figure 3 is shown in Figure D.1 in the Appendix for better viewing.

7. Conclusion

We have presented a robust and efficient 3D point registration method that uses the geometry of the $\mathbb{SE}(3)$ group and outperforms the state-of-the-art. We have also shown that when feature correspondences are unreliable, our method can make a multiview ICP method robust and effective.

References

- [1] Khurrum Aftab and Richard Hartley. Convergence of Iteratively Re-weighted Least Squares to Robust M-estimators. In Applications of Computer Vision (WACV), 2015 IEEE Winter Conference on, pages 480–487. IEEE, 2015.
- [2] Dror Aiger, Niloy J. Mitra, and Daniel Cohen-Or. 4-Points Congruent Sets for Robust Pairwise Surface Registration. In ACM Transactions on Graphics (ToG), volume 27 (3), page 85. ACM, 2008.
- [3] Federica Arrigoni, Beatrice Rossi, and Andrea Fusiello. Global Registration of 3D Point Sets via LRS Decomposition. In European Conference on Computer Vision, pages 489–504, 2016.
- [4] Adrien Bartoli, Daniel Pizarro, and Marco Loog. Stratified Generalized Procrustes Analysis. International journal of computer vision, 101(2):227–253, 2013.
- [5] Per Bergström and Ove Edlund. Robust Registration of Point Sets Using Iteratively Reweighted Least Squares. Computational Optimization and Applications, 58(3):543–561, 2014.
- [6] Michael J. Black and Anand Rangarajan. On the Unification of Line Processes, Outlier Rejection, and Robust Statistics with Applications in Early Vision. International Journal of Computer Vision, 19(1):57–91, 1996.
- [7] Sofien Bouaziz, Andrea Tagliasacchi, and Mark Pauly. Sparse Iterative Closest Point. In Proceedings of the Eleventh Eurographics/ACM SIGGRAPH Symposium on Geometry Processing, pages 113–123. Eurographics Association, 2013.
- [8] Avishek Chatterjee and Venu Madhav Govindu. Robust Relative Rotation Averaging. IEEE Transactions on Pattern Analysis and Machine Intelligence, 40(4):958–972, April 2018.
- [9] Dmitry Chetverikov, Dmitry Svirkov, Dmitry Stepanov, and Pavel Krsek. The Trimmed Iterative Closest Point Algorithm. In Proceedings of the 16th International Conference on Pattern Recognition, volume 3, pages 545–548. IEEE, 2002.
- [10] Gregory S. Chirikjian. Stochastic Models, Information Theory, and Lie Groups, Volume 2: Analytic Methods and Modern Applications. Birkhauser Boston, 2011.
- [11] Sungjoon Choi, Qian-Yi Zhou, and Vladlen Koltun. Robust Reconstruction of Indoor Scenes. In IEEE Conference on Computer Vision and Pattern Recognition, pages 5556–5565. IEEE, 2015.
- [12] Bertram Drost, Markus Ulrich, Nassir Navab, and Slobodan Ilic. Model Globally, Match Locally: Efficient and Robust 3D Object Recognition. In IEEE Conference on Computer Vision and Pattern Recognition, pages 998–1005. IEEE, 2010.
- [13] Ben Eckart, Kihwan Kim, and Jan Kautz. HGMR: Hierarchical Gaussian Mixtures for Adaptive 3D Registration. In The European Conference on Computer Vision (ECCV), September 2018.
- [14] Olof Enqvist, Klas Josephson, and Fredrik Kahl. Optimal Correspondences from Pairwise Constraints. In IEEE 12th International Conference on Computer Vision, pages 1295–1302. IEEE, 2009.
- [15] Georgios Evangelidis and Radu Horaud. Joint Alignment of Multiple Point Sets with Batch and Incremental Expectation-Maximization. IEEE Transactions on Pattern Analysis and Machine Intelligence, 40(6):1397–1410, June 2018.
- [16] Andrew W. Fitzgibbon. Robust Registration of 2D and 3D Point Sets. Image and Vision Computing, 21(13-14):1145–1153, 2003.
- [17] Natasha Gelfand, Niloy J. Mitra, Leonidas J. Guibas, and Helmut Pottmann. Robust Global Registration. In Symposium on Geometry Processing, volume 2 (3), page 5, 2005.
- [18] Venu Madhav Govindu. Lie-Algebraic Averaging for Globally Consistent Motion Estimation. In Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, volume 1, pages I–I. IEEE, 2004.
- [19] Venu Madhav Govindu and A. Pooja. On Averaging Multi-view Relations for 3D Scan Registration. IEEE Transactions on Image Processing, 23(3):1289–1302, March 2014.
- [20] Yulan Guo, Mohammed Bennamoun, Ferdous Sohel, Min Lu, and Jianwei Wan. 3D Object Recognition in Cluttered Scenes with Local Surface Features: A Survey. IEEE Transactions on Pattern Analysis and Machine Intelligence, 36(11):2270–2287, 2014.
- [21] Yulan Guo, Mohammed Bennamoun, Ferdous Sohel, Min Lu, Jianwei Wan, and Ngai Ming Kwok. A Comprehensive Performance Evaluation of 3D Local Feature Descriptors. International Journal of Computer Vision, 116(1):66–89, 2016.
- [22] Richard I. Hartley and Fredrik Kahl. Global Optimization Through Searching Rotation Space and Optimal Estimation of the Essential Matrix. In IEEE 11th International Conference on Computer Vision, pages 1–8. IEEE, 2007.
- [23] Paul W. Holland and Roy E. Welsch. Robust Regression Using Iteratively Reweighted Least-Squares. Communications in Statistics - Theory and Methods, 6(9):813–827, 1977.
- [24] Dirk Holz, Alexandru E. Ichim, Federico Tombari, Radu B. Rusu, and Sven Behnke. Registration with the Point Cloud Library: A Modular Framework for Aligning in 3D. IEEE Robotics & Automation Magazine, 22(4):110–124, 2015.
- [25] Bing Jian and Baba C. Vemuri. Robust Point Set Registration Using Gaussian Mixture Models. IEEE Transactions on Pattern Analysis and Machine Intelligence, 33(8):1633–1645, Aug 2011.
- [26] Felix Järemo Lawin, Martin Danelljan, Fahad Khan, Per-Erik Forssén, and Michael Felsberg. Density Adaptive Point Set Registration. In IEEE Conference on Computer Vision and Pattern Recognition. Computer Vision Foundation, June 2018.
- [27] Hongdong Li and Richard Hartley. The 3D-3D Registration Problem Revisited. In IEEE 11th International Conference on Computer Vision, pages 1–8. IEEE, 2007.
- [28] Ameesh Makadia, Alexander Patterson, and Kostas Daniilidis. Fully Automatic Registration of 3D point Clouds. In IEEE Computer Society Conference on Computer Vision and Pattern Recognition, volume 1, pages 1297–1304. IEEE, 2006.

- [29] Peter Meer, Doron Mintz, Azriel Rosenfeld, and Dong Yoon Kim. Robust Regression Methods for Computer Vision: A Review. International Journal of Computer Vision, 6(1):59–70, 1991.
- [30] Nicolas Mellado, Dror Aiger, and Niloy J. Mitra. Super 4PCS Fast Global Pointcloud Registration via Smart Indexing. In Computer Graphics Forum, volume 33 (5), pages 205–215. Wiley Online Library, 2014.
- [31] Ajmal S. Mian, Mohammed Bennamoun, and Robyn Owens. Three-Dimensional Model-Based Object Recognition and Segmentation in Cluttered Scenes. IEEE Transactions on Pattern Analysis and Machine Intelligence, 28(10):1584–1601, 2006.
- [32] Richard A. Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J. Davison, Pushmeet Kohli, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. KinectFusion: Real-Time Dense Surface Mapping and Tracking. In Proceedings of the 10th IEEE International Symposium on Mixed and Augmented Reality, 2011.
- [33] Chavdar Papazov, Sami Haddadin, Sven Parusel, Kai Krieger, and Darius Burschka. Rigid 3D Geometry Matching for Grasping of Known Objects in Cluttered Scenes. The International Journal of Robotics Research, 31(4):538–553, 2012.
- [34] Frank C. Park. Distance Metrics on the Rigid-Body Motions with Applications to Mechanism Design. Journal of Mechanical Design, 117(1):48–54, 1995.
- [35] Rahul Raguram, Jan-Michael Frahm, and Marc Pollefeys. A Comparative Analysis of RANSAC Techniques Leading to Adaptive Real-Time Random Sample Consensus. In European Conference on Computer Vision, pages 500–513. Springer, 2008.
- [36] Carolina Raposo and Joao P. Barreto. Using 2 Point + Normal Sets for Fast Registration of Point Clouds with Small Overlap. In 2017 IEEE International Conference on Robotics and Automation (ICRA), pages 5652–5658. IEEE, 2017.
- [37] Szymon Rusinkiewicz and Marc Levoy. Efficient Variants of the ICP Algorithm. In Proceedings of the Third International Conference on 3D Digital Imaging and Modeling, pages 145–152. IEEE, 2001.
- [38] Radu Bogdan Rusu, Nico Blodow, and Michael Beetz. Fast Point Feature Histograms (FPFH) for 3D Registration. In IEEE International Conference on Robotics and Automation, pages 3212–3217. IEEE, 2009.
- [39] Renato F. Salas-Moreno, Richard A. Newcombe, Hauke Strasdat, Paul H.J. Kelly, and Andrew J. Davison. SLAM++: Simultaneous Localisation and Mapping at the Level of Objects. In IEEE Conference on Computer Vision and Pattern Recognition, pages 1352–1359. IEEE, 2013.
- [40] Joaquim Salvi, Carles Matabosch, David Fofi, and Josep Forest. A Review of Recent Range Image Registration Methods with Accuracy Evaluation. Image and Vision Computing, 25(5):578–596, 2007.
- [41] Jon M. Selig. Geometrical Methods in Robotics. Springer New York, 2013.
- [42] Jürgen Sturm, Nikolas Engelhard, Felix Endres, Wolfram Burgard, and Daniel Cremers. A Benchmark for the Evaluation of RGB-D SLAM Systems. In International Conference on Intelligent Robots and Systems (IROS), IEEE/RSJ, pages 573–580. IEEE, 2012.
- [43] Richard Szeliski. Computer Vision: Algorithms and Applications. Springer-Verlag New York, Inc., 2010.
- [44] Gary K. L. Tam, Zhi-Quan Cheng, Yu-Kun Lai, Frank C. Langbein, Yonghuai Liu, David Marshall, Ralph R. Martin, Xian-Fang Sun, and Paul L. Rosin. Registration of 3D Point Clouds and Meshes: A Survey from Rigid to Nonrigid. IEEE Transactions on Visualization and Computer Graphics, 19(7):1199–1217, 2013.
- [45] Federico Tombari, Samuele Salti, and Luigi Di Stefano. Performance Evaluation of 3D Keypoint Detectors. International Journal of Computer Vision, 102(1-3):198–220, 2013.
- [46] Andrea Torsello and Andrea Albarelli. Multi-View Registration via Graph Diffusion of Dual Quaternions. In IEEE Conference on Computer Vision and Pattern Recognition, pages 2441–2448, 2011.
- [47] Shinji Umeyama. Least-Squares Estimation of Transformation Parameters Between Two Point Patterns. IEEE Transactions on Pattern Analysis and Machine Intelligence, 13(4):376–380, 1991.
- [48] Jiaolong Yang, Hongdong Li, Dylan Campbell, and Yunde Jia. Go-ICP: a Globally Optimal Solution to 3D ICP Point-Set Registration. IEEE Transactions on Pattern Analysis and Machine Intelligence, 38(11):2241–2254, 2016.
- [49] Jiaolong Yang, Hongdong Li, and Yunde Jia. Go-ICP: Solving 3D Registration Efficiently and Globally Optimally. In IEEE International Conference on Computer Vision, pages 1457–1464. IEEE, 2013.
- [50] Ernesto Zacur, Matias Bossa, and Salvador Olmos. Left-Invariant Riemannian Geodesics on Spatial Transformation Groups. SIAM Journal on Imaging Sciences, 7(3):1503–1557, 2014.
- [51] Zhengyou Zhang. Microsoft Kinect Sensor and its Effect. IEEE Multimedia, 19(2):4–10, 2012.
- [52] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Fast Global Registration. In European Conference on Computer Vision, pages 766–782. Springer, 2016.