https://doi.org/10.2352/ISSN.2470-1173.2019.16.3DMP-008

Modified M-Estimation for Fast Global Registration of 3D Point Clouds

Faisal Azhar, Stephen Pollard and Guy Adams; HP Labs; Bristol, United Kingdom

Abstract

We present a modified M-estimation based method for fast global 3D point cloud registration which rapidly converges to an optimal solution while matching or exceeding the accuracy of existing global registration methods.The key idea of our work is to introduce weighted median based M-estimation for re-weighted least squares adeployed in a graduated fashion which takes into account the error distribution of the residuals to achieve rapid convergence to an optimal solution. The experimental results on synthetic and real data sets show the significantly improved convergence of our method with a comparable accuracy with respect to the state-of-the-art global registration methods.

Introduction

Three-dimensional (3D) point cloud registration is the problem of consistently and accurately aligning two or more point clouds, i.e., points representing the X, Y, and Z geometric coordinates of an underlying sampled surface. 3D registration is of growing importance in the field of computer vision and robotics. The registration can be used to align 3D model data in the form of CAD models or full 3D scans with partial scan data, depicting the current scene, to perform tasks such as 3D object retrieval. It can also be applied to perform 3D object interaction tasks such as robotic bin picking, 2D and 3D inspection and authentication. Furthermore, multiple partial scans can be aligned for scene/object reconstruction to generate unified/full 3D point clouds. These are very challenging problems because the 3D data captured from the scanning devices, e.g., structured light scanners, time-of-flight sensors, binocular stereo etc., tends to be noisy and contains only partial data.

The standard 3D registration approach is to first find an initial or coarse alignment and then use the Iterative Closest Point (ICP) [2] algorithm to refine and obtain a final alignment between a model and scene. The model is often a computer aided design (CAD) or a full 3D scan and the scene is a single scan or multiview scan of the object, alternatively the model and scene may be simply two overlapping views of the same object or scene. The initial alignment employs 3D feature descriptors in a Random Sample and Consensus (RANSAC) [7] model fitting scheme to find a small number of correct point correspondences between the model and the scene and from there obtain an initial rigid transformation which aligns the scene to the model. The sampling based initial alignment and iterative final alignment are computationally expensive because they require many alignments to be tested in order to robustly determine the optimal transformation. Each initial local estimate requires at least three nearest neighbour point to point correspondences to be tested at random. The majority of the computational burden is in testing candidate alignments which are later discarded as they prove to be sub-optimal.

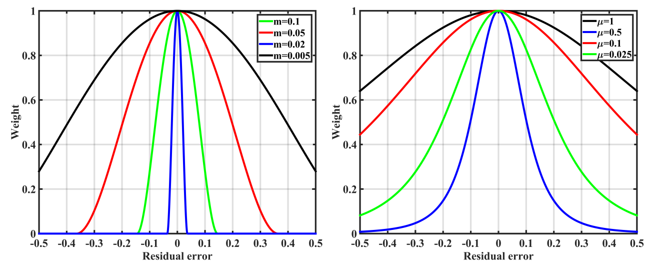

(a) (b)

Figure 1: Robust estimator weighting functions. (a) Tukey bisuquare and (b) Modified German McClure.

In order to address these issues, a fast global registration (FGR) method is proposed in [19] that does not involve iterative sampling, model fitting, or local refinement. It does not require initialization and can align noisy partially overlapping point clouds. It optimizes a modified German-McClure (GMC) objective function over a fixed set of point correspondences to directly produce a precise alignment similar to a well-initialized initial refinement algorithm. Therefore, it is an order of a magnitude faster than existing global registration methods. The FGR method uses a fixed regime of reducing the control parameter of the GMC objective function to increase the weight for point correspondences with smaller residuals. Due to this fixed regime, termed as graduated non convexity, which is based on the the dimensions of the point cloud the FGR method tends to converge rather slowly.

In this paper, we present a modified M-estimation based fast global registration (MFR) method which rapidly (i.e., significantly faster than FGR) converges to an optimal solution while matching or exceeding the accuracy of the FGR and other existing global registration methods. The contributions of our work are as follows. (1) weighted median based M-estimation for re-weighted least squares, and (2) graduated M-estimation which takes into account the error distribution of the residuals for rapid convergence to an optimal solution.

Related Work

The geometric registration of 3D point clouds or surfaces has been extensively researched [4, 19, 18, 6, 10, 16, 15, 17, 12, 11]. It generally consists of two stages: global alignment, which computes an initial estimate of the rigid motion between two point cloud, which is followed by local refinement, which refines this initial estimate to obtain a final registration .

Most global methods use hand crafted 3D feature detectors/descriptors to extract meaningful information from the 3D point clouds. Such descriptors are recovered from the dense point cloud to derive per point rotation invariant feature descriptors which are used to search for best candidate correspondences on a nearest neighbour basis. Some of the most popular 3D feature descriptors include point feature histogram (PFH) [16], fast point feature histogram (FPFH) [15], signature of histogram of orientations (SHOT) [17], etc. RANSAC is used to repeatedly estimate an alignment for a randomly chosen subset of correspondences which is validated on the entire point cloud. Noisy data and partially overlapping point clouds create a significant problem to these methods because they require many repetitions to find a good correspondence set which also results in a transformation that is sufficiently close to the required optimal solution.

The local refinement methods such as ICP [10] require an initial alignment to produce a final registration. ICP starts with an initial alignment and iterates a search for point correspondences via nearest Euclidian distance matches (point to point or point to plane) and then recomputes the alignment based on the new set of correspondences. ICP only obtains a good registration when its starting alignment is close to the optimal solution.The work presented in [8] is a non-linear least squares optimization with ICP employing robust estimation technique. More recent and related examples of ICP explicitly deploying variants of M-estimation are due to Ding et. al [5] and Bergstrom and Edlund [1]. However, all these methods require a good initial solution to be able to obtain a satisfactory final solution.

As discussed in the introduction, our work is motivated from the FGR [19] method in which a fixed subset of corresponding points, some correct and some not, are considered on mass. The aim of the optimisation in [19] is to iterate towards a solution that selects the correct matches while discarding the noise using a GMC objective function within a graduated non-convexity framework.

Our Contribution

The primary contribution of our MFR method is to employ a form of M-estimation so that the error statistics can directly determine the rate of convergence, resulting in much reduced computational requirements. Furthermore, we introduce two small modifications to the M-estimation approach to help preserve the graduated non-convexity and improve its effectiveness, namely using a weighted median to determine the error measure used in the M-estimation and reducing the control parameter of the robust bi-square weighting function through the iterative process to achieve better selection of the correctly corresponding points. These changes allow us to match or exceed the performance of the FGR method in terms of transformation accuracy and associated robust rejection of outlier correspondences, in a scheme that converges much more quickly to the optimal solution.

Mathematical Foundation

Let $\mathbf{P} = [\mathbf{p}_1, \mathbf{p}_2, ..., \mathbf{p}_n]$ and $\mathbf{Q} = [\mathbf{q}_1, \mathbf{q}_2, ..., \mathbf{q}_n]$ be two sets of n corresponding 3D points. We wish to find a rigid transformation $\mathbf{T} = [\mathbf{R} \quad \mathbf{t}]$ that optimally and robustly aligns the two sets of 3D points in the least squares sense, i.e., we seek a rotation $\mathbf{R}$ 3×3 matrix and a translation 3×1 vector $\mathbf{t}$ such that

$$(\mathbf{R}, \mathbf{t}) = \underset{\mathbf{R}, \mathbf{t}}{argmin} \sum_{i=1}^{n} w_i \|(\mathbf{R}\mathbf{p}_i + t) - \mathbf{q}_i\|^2 \tag{1}$$

where weights $w_i \geq 0$ provides a robust selection for the i^{th} point correspondence.

FGR Method

The FGR method uses the linearized form of $\mathbf{T}$, i.e., to approximate the global transformation through a process of successive composition in which each sub-transform is represented as a 6D vector containing three rotational ($\Theta \quad \Theta_y \quad \Theta_z$) and three translational components ($t_x \quad t_y \quad t_z$), to minimize Equation 1. The authors implementation available on the GitHub confirms that the FGR method uses Equation 1 and applies weights computed based on the modified GMC objective function,

$$w_i = \left(\frac{\mu}{\mu + \|\mathbf{T}\mathbf{p}_i - \mathbf{q}_i\|^2} \right)^2 \tag{2}$$

where weight is a function of the residual error $r_i = \|\mathbf{T}\mathbf{p}_i - \mathbf{q}_i\|^2$ measured at the previous iteration.

The FGR method uses Gauss Newton optimization to first solve a linear system $\mathbf{J}^{\mathrm{T}}\mathbf{J}\xi = -\mathbf{J}^{\mathrm{T}}\mathbf{r}$ which approximates Equation 1. $\mathbf{J}$ is the Jacobian of the residual error $\mathbf{r}$ based on the first order linear approximation

$$\mathbf{T}_L = \begin{bmatrix} 1 & -\Theta_z & \Theta_y & t_x \\ \Theta_z & 1 & -\Theta_x & t_y \\ -\Theta_y & \Theta_x & 1 & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Subsequently, at each iteration the 6D vector is mapped to a full transformation matrix $\mathbf{T}^k = \mathbf{T}_{\triangle}\mathbf{T}^{k-1}$where $\mathbf{T}_{\triangle}$ is an orthonormal transformation matrix based on the solution to the linear system $\mathbf{T}_L$ and superscript k denotes the transformation for k^{th} iteration (see [19] for detail). More importantly the FGR method employs a fixed regime of reducing the control parameter μ every four iterations by a fixed factor until a fixed threshold (0.025 units of the diameter of the surface) is reached. However, note that in practice as per the publicly available GitHub implementation the FGR method continues to iterate for a fixed 64 iterations. This reduction of control parameter μ is termed as graduated non-convexity [3] and begins with a very broad weighting function with μ set to the square of the diameter of the surface (representative examples of the modified GMC for a unit diameter surface are shown in Figure 1b). The drawback of this approach is that the weighting will only very slowly differentiate in favour of point correspondences with smaller residual error. Hence, the result is a slow convergence to the optimal solution.

Proposed MFR Method

We propose a modifed M-estimation based fast global registration (MFR) method which addresses the above-mentioned drawbacks of the FGR method. Our MFR method rapidly converges to an optimal solution while matching or exceeding the accuracy of the FGR and other existing global registration methods. The key idea of our work is to introduce weighted median based M-estimation for re-weighted least squares and deploy a graduated M-estimation which takes into account the error distribution of the residuals to achieve a rapid convergence to the optimal solution.

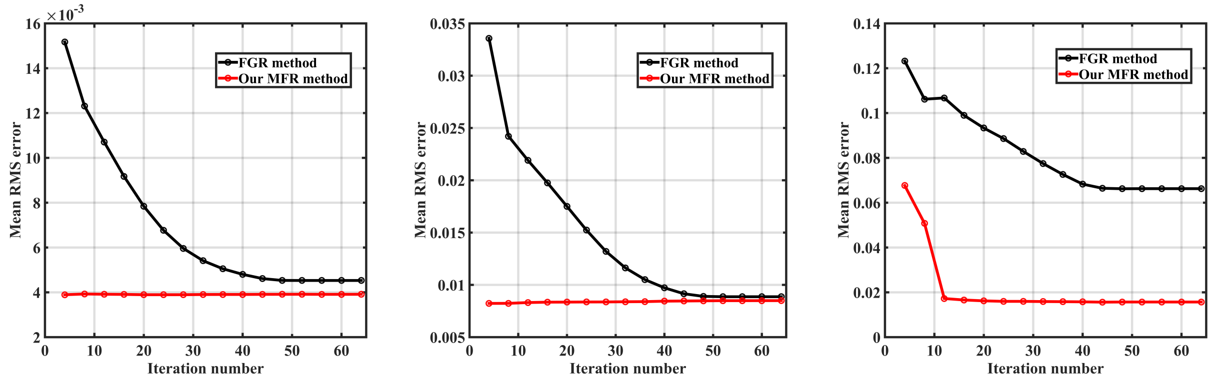

(a) (b) (c)

Figure 2: Comparison of convergence plots for our MFR versus the FGR method. Mean RMS error with respect to iteration number for each of the three Gaussian noise levels, (a) $\eta = 0$, (b) $\eta = 0.25\%$ and (c) $\eta = 0.5\%$.

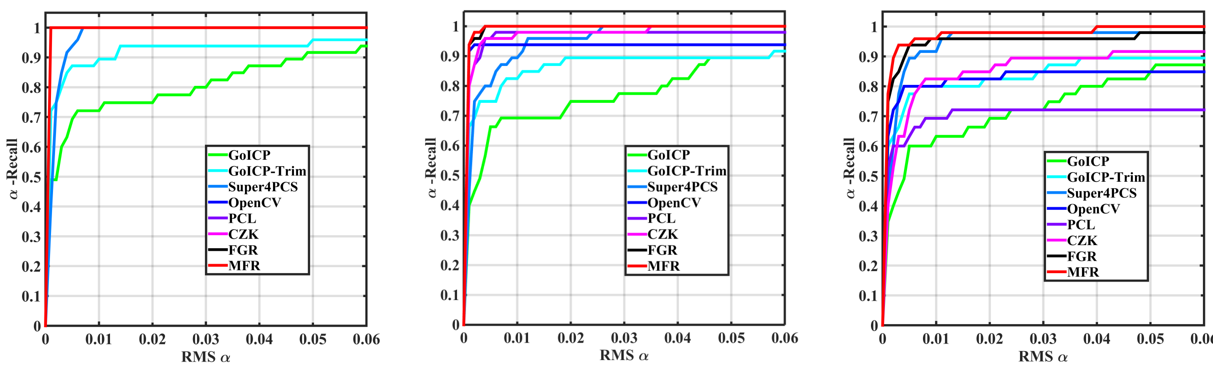

(a) (b) (c)

Figure 3: α-Recall is the fraction of test for which a method achieves RMS error less than a threshold set using α value (on the x-axis). RMS is defined in unit of diameter of surface. (a) $\eta = 0$, (b) $\eta = 0.25\%$ and (c) $\eta = 0.5\%$.

Tukey bisquare objective with weighted median:

The mathematical details of our MFR method are as follows. First, we estimate the residual error $r_i = \|\mathbf{T}\mathbf{p}_i - \mathbf{q}_i\|^2$ for the n point correspondences. Next, we use a well known Tukey bisquare objective function to compute a weighting function of the form

$$w_i = \begin{cases} \left[1 - \left(\frac{r}{k}\right)^2\right]^2 & if\,|r_i| \leq k \\ 0 & if\,|r_i| > k \end{cases} \tag{3}$$

where $k = \Psi\sigma$ is a scalar derived from the current set of residual errors r_i. The parameter σ is a measure of the deviation of the residual errors at the previous iteration and Ψ is a control parameter which is traditionally set to $\Psi = 4.8651$. We compute a suitable σ for penalizing outliers using

$$\sigma = 1.4826\left(1 + \frac{5}{n-3}\right)m \tag{4}$$

For n point correspondences $[\mathbf{p}_1, \mathbf{p}_2, ..., \mathbf{p}_n]$ ordered according to their residual errors $[r_1, r_2, \ldots, r_n]$ with normalised positive weights $[w_1, w_2, ..., w_n]$ such that $\sum w_i = 1$, the weighted median m of the residuals is defined as the residual r_i of the element $\mathbf{p}_i$ which satisfies

$$m = \begin{cases} r_i & if \sum_{k=1}^{i-1}{}_i \leq 0.5 \\ & and \sum_{k=i+1}^{n}{}_i \leq 0.5 \end{cases} \tag{5}$$

The weighted median is both robust against noise and outliers, and allows for non-uniform statistical weights related to the residual errors as opposed to the median which assigns uniform weights. Representative examples of the weighting function for a unit diameter surface are shown in Figure 1a for various measures of the weighted mean m for a fixed value of the control parameter $\Psi = 4.8651$.

Graduated non-convexity M-estimation:

In contrast to the two step linear optimization of the FGR method, we solve for a full transformation matrix $\mathbf{T} = [\mathbf{R} \quad \mathbf{t}]$ in a single step using the traditional singular value decomposition (SVD) [9] method to determine the re-weighted least squares solution for Equation 1. The SVD based procedure is as follows. Given two 3D point sets $\mathbf{P} = [\mathbf{p}_1, \mathbf{p}_2, ..., \mathbf{p}_n]$ and $\mathbf{Q} = [\mathbf{q}_1, \mathbf{q}_2, ..., \mathbf{q}_n]$, first compute the weighted centroids using

Algorithm 1 Our modified fast global registration (MFR) method

Input:
Two sets of n corresponding 3D points, i.e., $\mathbf{P} = [\mathbf{p}_1, \mathbf{p}_2, ..., \mathbf{p}_n]$ and $\mathbf{Q} = [\mathbf{q}_1, \mathbf{q}_2, ..., \mathbf{q}_n]$.
Output: Transformation matrix $\mathbf{T}_{Mest}$ which aligns $\mathbf{P}$ to $\mathbf{Q}$.

Set $\Psi_s = 4.6851$, $\Psi_f = 3$, $step = {}^{(\Psi_s - \Psi_f)}/_{16}$ (subscript s and f is for starting and final value), weights $w_{1,...,n} = 1$ initialized to one, and Transformation matrix$\mathbf{T}_{Mest}$ set to identity.
MFR estimation:
FOR $Iteration < MaxItr$

1. Compute a rigid transformation using SVD on the two sets $\mathbf{P}$ to $\mathbf{Q}$ using Equation 6 - 10 to get full transformation matrix $\mathbf{T} = [\mathbf{R} \quad \mathbf{t}]$.
2. Estimate the residual error $r_i = \|\mathbf{T}\mathbf{p}_i - \mathbf{q}_i\|^2$ for $i = 1, ..., n$ point correspondences.
3. Apply the Tukey bisquare objective function using Equation 3 - 5 to get corresponding updated weights $W_i = w_i^2$.
4. Graduated reduction of the control parameter such that $\Psi_s = max(\Psi_s - step, \Psi_f)$.

END

5. Final transformation matrix $\mathbf{T}_{Mest}$ =$\mathbf{T}$.

$$\bar{\mathbf{p}} = \frac{\sum_{i=1}^{n} w_i \mathbf{p}_i}{\sum_{i=1}^{n} {}_i}, \mathbf{q} = \frac{\sum_{i=1}^{n} w_i \mathbf{q}_i}{\sum_{i=1}^{n} {}_i} \tag{6}$$

then compute a set $\mathbf{A}$ and $\mathbf{B}$ containing centered vectors such as

$$\mathbf{A} = [\mathbf{p}_1 - \bar{\mathbf{p}}, ..., \mathbf{p}_n - \bar{\mathbf{p}}], \quad B = [\mathbf{q}_1 - \bar{\mathbf{q}}, ..., \mathbf{q}_n - \bar{\mathbf{q}}] \tag{7}$$

Next, compute the 3×3 covariance matrix $\mathbf{C}$ using

$$\mathbf{C} = \mathbf{A}\mathbf{W}\mathbf{B}^T \tag{8}$$

where $\mathbf{A}$ and $\mathbf{B}$ are $3 \times n$ matrices which have centered vectors as their columns, $\mathbf{W} = diag(w_1, w_2, ..., w_n)$ and superscript T is the transpose.

Finally, we apply SVD to the covariance matrix $\mathbf{C}$ to obtain $\mathbf{U}\mathbf{S}\mathbf{V}^T = \mathbf{C}$, where $\mathbf{U}$ and $\mathbf{V}$ are orthogonal matrices, and $\mathbf{S}$ is a diagonal matrix of non-negative singular values[9]. The optimal rotation matrix $\mathbf{R}$ is given by

$$\mathbf{R} = \mathbf{V}\mathbf{U}^T \tag{9}$$

and the translation vector $\mathbf{t}$ by

$$\mathbf{t} = \bar{\mathbf{q}} - \mathbf{R}\,\bar{\mathbf{p}} \tag{10}$$

We apply our weighted median based M-estimator in a graduated non-convexity scheme by reducing the control parameter Ψ which controls the shape of the Tukey bisquare objective function. Our weighted median based graduated M-estimation obtains rapid convergence and high accuracy of registration (see Experiments Section).

Implementation details:

We are motivated from the work of Rusu et. al [15, 10] on 3D descriptor and Qian et. al [19] on generating fixed initial sets of corresponding 3D points. For each point $\mathbf{p}$, all of its neighbours enclosed in the sphere with a given radius (10 units/mm) are selected. Next, for every pair of points in the given radius of $\mathbf{p}$ and their estimated normals, a Darboux frame [15] is defined to compute 33 dimensional FPFH descriptor for the point $\mathbf{p}$. Subsequently, we use the FPFH descriptor to find nearest neighbour correspondences similar to the FGR [19] method. First, a set of correspondences is built which contains for each point $\mathbf{p}$ nearest neighbour in set $\mathbf{Q}$ and for each point $\mathbf{q}$ nearest neighbour $\mathbf{P}$. Next, reciprocity is used to select correspondence pair $(\mathbf{p}, \mathbf{q})$ if and only if $\mathbf{p}$ has the nearest neighbor in set $\mathbf{Q}$ and $\mathbf{q}$ has the nearest neighbor in $\mathbf{P}$. Finally, tuple test is used to select correspondence pair which meet the inlier ratio test. In tuple test, three correspondence pairs are randomly picked to test $0.9 < {}^{\|\mathbf{p}_i - \mathbf{p}_j\|}/_{\|\mathbf{q}_i - \mathbf{q}_j\|} < 1.11$. Our modified fast global registration method which used the initial correspondence sets from the above-described process is detailed in the Algorithm 1.

Experiments

We have conducted a series of experiments to compare the performance of our MFR method against seven state-of-the-art global registration methods on two data sets, i.e., a synthetic range dataset [19] and the UWA benchmark dataset [13]. We are primarily motivated by the FGR [19] method and are grateful for its publically available evaluation functions. The synthetic range and UWA dataset are used to compare our MFR method with the FGR and 6 other global registration methods (termed as GoICP, GoICP-Trim, Super4PCS, OpenCV PCL and CZK). ICP is known to be susceptible to local minima, GoICP [18] integrates local ICP into the branch-and-bound [14] scheme to guarantee a globally optimal solution. GoICP-Trim is its 10% trimming variant with 1000 data points that supports partial overlap. Super4PCS [12] is an optimal linear time output-sensitive algorithm which uses an efficient data structure to obtain a global alignment. OpenCV [6] is the surface registration algorithm which uses point pair features with hash table lookup and voting with pose clustering to obtain a global registration. PCL [10, 15] is the sample consensus initial alignment algorithm which uses histogram of point pair features, i.e., FPFH, to obtain a global registration. CZK [4] is a method which combines geometric registration of scene fragments with robust global optimization based on line processes for 3D scene reconstruction. We use the evaluation function of FGR method and compute root mean square (RMS) point-to-point error for all the global registration methods.

Synthetic range data

This dataset consists of 5 models (Bimba, Dancing Children, Chinese Dragon, Angel and Bunny) with 5 pairs of partially overlapping range data with 3 Gaussian noise levels ($\eta = 0$, $\eta = 0.25\%$ and $\eta = 0.5\%$), i.e., a total of 75 pairs for testing registration (25 pairs of tests for each noise level). The noise levels are scaled with respect to the normalized diameter (1.0) of the models. The number of points vary between $8,868$ and $19,749$ points with overlap ratio between 47% and 90%.

In Figure 2, we show the convergence plots of mean RMS error with respect to every four iterations with the FGR μ and

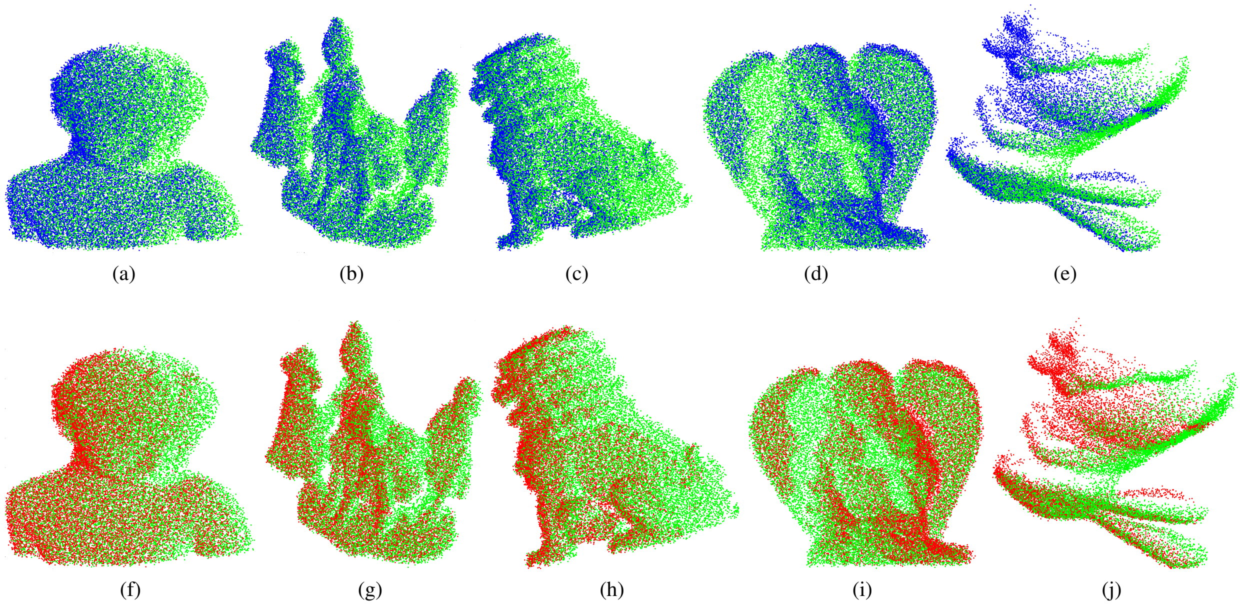

Figure 4: Visualization example of two overlapping points clouds after scene to model (*Green points*) registration is obtained using: (a)-(e) FGR (*Blue points*) method and (f)-(j) MFR (*Red points*) method. Bimba, Dancing Children, Chinese Dragon, Angel and Bunny model left to right column.

the MFR Ψ control parameter reduced by a fixed factor until the maximum 64 iterations is reached. In practice μ reduces from unity by a factor of 1.4 every 4 iterations stopping at 0.025 after 44 iterations, while Ψ reduces linearly from 4.8651 to 3.0 over the first 16 iterations. Figure 2a-2c compares the convergence of the FGR and the MFR method for Gaussian noise levels $\eta = 0.0\%$, $\eta = 0.25\%$ and $\eta = 0.5\%$ respectively. It can be seen from Figure 2 that our MFR method obtains a rapid convergence and high accuracy of registration. Our MFR method uses only 4 iterations as opposed to approximately 50 iterations required by the FGR method to reach an optimal solution for $\eta = 0.0\%$ and $\eta = 0.25\%$ as shown in Figure 2a and 2b. Also, our MFR method uses only 12 iterations to reach an optimal solution for high noise level $\eta = 0.5\%$ as shown in Figure 2c. Note that in all variations of Gaussian noise levels our MFR method achieves a lower mean RMS error than the FGR method. In particular, the MFR method outperforms the FGR method in terms of low mean RMS error for high noise $\eta = 0.5\%$. Hence, the results demonstrate that the MFR uses approximately $1/10^{th}$ the number of iterations of the FGR method to rapidly converge to an optimal solution which is better than the state-of-the-art FGR method.

Figure 3 compares the accuracy of our MFR method versus GoICP, GoICP-Trim, Super4PCS, OpenCV, PCL, CZK and FGR methods. Figure 3a, 3b and 3c show the α-Recall as a fraction of 25 pair wise registrations with RMS $< \alpha$ with respect to $\eta = 0.0\%$, $\eta = 0.25\%$ and $\eta = 0.5\%$ respectively. It can be seen from Figure 3 that our MFR method matches or surpasses all the global registration methods. In particular, for $\eta = 0.5\%$ our MFR method is prominently better than the FGR and Super4PCS methods, and significantly better than CZK, GoICP-Trim, OpenCV, PCL and GoICP methods.

Figure 4 shows visual examples of registration obtained using the FGR (*Blue points*) and the MFR (*Red points*) method on the synthetic range data set models (*Green points*). It shows that our MFR method with smaller number of iterations matches the alignment accuracy of the FGR.

The Figure 5 below shows the effectiveness of graduated non-convexity (GNC) and weighted median (WM) in the context of M-Estimation. The recall and convergence for our weighted median graduated non-convexity (WM-GNC) versus graduated non-convexity (GNC), weighted median (WM) and standard M-Estimation (M-Est) for high noise $\eta = 0.5\%$ is shown in Figure (a) and (b) respectively. Note that the performance differences were not significant for the lower noise examples. In conclusion it can be seen that the modifications provide a small but significant improvement in performance that go beyond the state of the art provided by the FGR.

UWA benchmark dataset

This dataset consists of 4 models (Cheff, Chicken, Parasaurolophus, T-rex) with 50 scenes of multiple objects, i.e., all 4 objects are present in most of the scenes. We exclude the Rhino model of UWA (as was the case for the FGR paper [19]) because its ground truth is not provided. Also, some scenes do not contain any 3D data for an object due to occlusion or viewpoint of the scanner, hence yielding a total of 188 pairs of model and scene pairs for testing registration. It is a challenging data set due to clutter, occlusion and low overlap. In Figure 6a, we plot mean RMS registration error measured between each model and the scene in the units in which each are represented (presumably mm) over the whole set of model and scene pairs against the number of iterations for each of the FGR and MFR methods. Notice again that the MFR has much improved convergence. In Figure 6b, we show an example scene point cloud (in Red) and examples of each model matched against it (in Green) for the MFR method.

Computational Time

In Figure 7, we present graphs of computational time in seconds of our MFR method versus the FGR global registration method. The computational time was measured in C++ on a HP Z book with an Intel i7 2.6 GHz processor with 16 GB RAM. In Figure 7a, we show results for each of the 75 point cloud pairs from the synthetic data set while in Figure 7b we show similar graphs for the 188 model and scene pairs for the UWA data sets. For each case we plot the time for all 64 iterations for each method, and also the time for the reduced number of iterations generally required by our MFR method to achieve a stable optimal result. In summary, the average time across the synthetic data set for the FGR method is 0.199 secs and the MFR method is 0.0089 secs. Hence, the MFR method is 22 times faster than the FGR method. In addition, for the UWA dataset the average time for the FGR method is 0.4351 secs and for the MFR method is 0.0167 secs. Hence, the MFR method is 26 times faster than the FGR method.

Conclusions

We have presented MFR, a fast global registration method that uses a form of M-estimation to provide improved rate of convergence when compared to the state of the art methods whilst maintaining overall performance. The reason behind this is that the bi-square weighting function used in M-estimation is derived in part from the measured residuals at the current iteration and so can adapt to the situation (amount of noise and/or number of outliers present in the data). In comparison the state of the art FGR method uses a re-weighting scheme that operates open loop and is not dependent upon the measured residuals. Instead it uses the dimension of the measured surface/object to determine a re-weighting strategy that is fixed for all situations. For the data at hand, which is the same data that was presented for the FGR method in [19], we have shown that it was not necessary to use such a conservative graduated non-convexity (GNC) form of re-weighting strategy in order to achieve the required performance. We have no doubt that it would be possible to change the design of the GNC component of the FGR algorithm to achieve better convergence speed for the data sets covered in this paper. However, it seems better to use a standard robust statistical technique, as provided by the bi-square weighting function, to provide a re-weighting scheme that is tied closely to the data itself. In practice, we made some minor modifications to the standard M-estimation technique to ensure it always outperformed the FGR method. Most significantly, we introduced a small amount of GNC into the technique where we iteratively reduced the control parameter of weighting function to essentially tighten the selection criteria and more significantly reject outliers. This improvement was derived directly from the FGR method and provided a small but significant advantage to the MFR method. The second improvement was less significant and just involved using a weighted median, instead of the standard median, in the re-weighting formula. This adds extra computational cost and provides a small improvement in performance. Even with this extra cost, we found that the MFR approach, which is based on the standard SVD method, is computational less expensive per iteration than the linearised Jacobian method of the FGR. Furthermore as the SVD approach computes a full transformation matrix (rather than a first order approximation), and uses a stronger re-weighting if the data permits, each of these iterations is much more significant and result in considerably improved convergence behaviour.

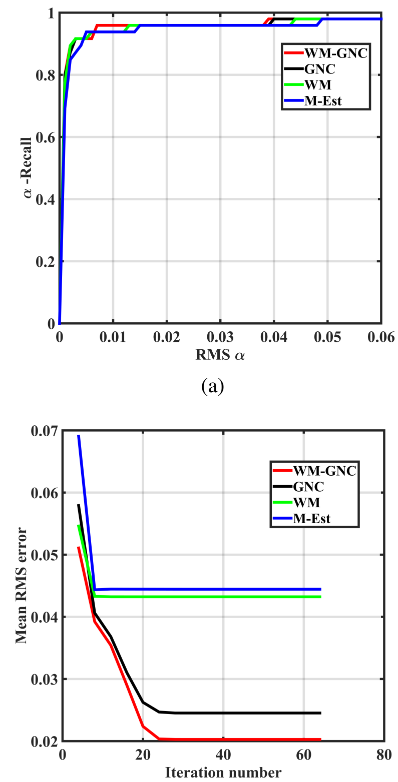

(a)

(b)

Figure 5: Effectiveness of weighted median (WM) plus graduated non-convexity (GNC) versus GNC alone, WM alone and standard M-Estimation (M-Est) for high noise $\eta = 0.5\%$. (a) α-Recall for tests in which a method achieves RMS$<\alpha$, and (b) Convergence plot of mean RMS error with respect to iteration number.

References

[1] Per Bergström and Ove Edlund. Robust registration of surfaces using a refined iterative closest point algorithm with a trust region approach. *Numerical Algorithms*, 74(3):755–779, Mar 2017.

[2] P. J. Besl and N. D. McKay. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.

[3] 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, Jul 1996.

[4] Sungjoon Choi, Qian-Yi Zhou, and Vladlen Koltun. Robust reconstruction of indoor scenes. In *CVPR*, pages 5556–5565, 2015.

[5] J. Ding, Q. Liu, P. Sun, and J. Wang. Robust registration of 3d point sets for free-form surface inspection. In *IEEE International Conference on Mechatronics and Automation*, pages 1012–1017, Aug 2017.

[6] B. Drost, M. Ulrich, N. Navab, and S. Ilic. Model globally, match locally: Efficient and robust 3d object recognition. In

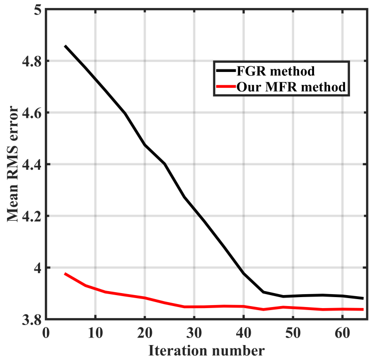

(a)

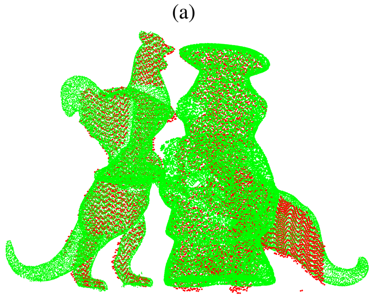

(b)

Figure 6: Comparison of convergence plots for our MFR versus the FGR method. (a) Mean RMS error with respect to iteration number for all 188 pair wise registrations, and (b) Visualization example of registration using our MFR method.

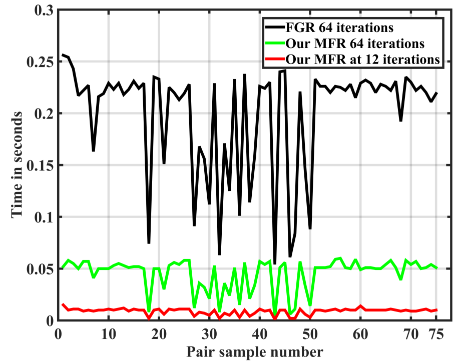

(a)

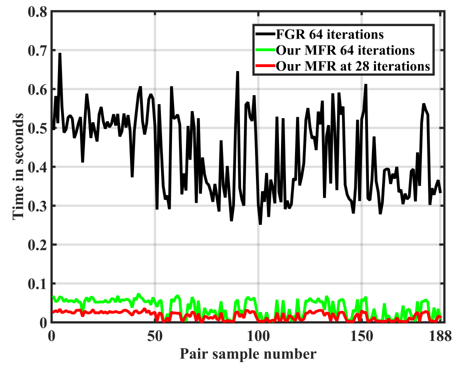

(b)

Figure 7: Comparison of computational time in seconds for our MFR versus the FGR method. (a) Synthetic dataset and (b) UWA dataset.

IEEE Conference on Computer Vision and Pattern Recognition, pages 998–1005, June 2010.

[7] M. A. Fischler and R.C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *ACM Commun.*, 24(6):381–395, June 1981.

[8] A. W. Fitzgibbon. Robust registration of 2D and 3D point sets. In *British Machine Vision Conference*, pages 662–670, 2001.

[9] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2003.

[10] D. Holz, A. E. Ichim, F. Tombari, R. B. Rusu, and S. Behnke. Registration with the point cloud library: A modular framework for aligning in 3-d. *IEEE Robotics Automation Magazine*, 22(4):110–124, Dec 2015.

[11] Q. Y. Zhou J. Park and V. Koltun. Colored point cloud registration revisited. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 143–152, Oct 2017.

[12] Nicolas Mellado, Dror Aiger, and Niloy J. Mitra. Super 4pcs fast global pointcloud registration via smart indexing. *Comput. Graph. Forum*, 33:205–215, Aug 2014.

[13] Ajmal S. Mian, Mohammed Bennamoun, and Robyn Owens. Three-dimensional model-based object recognition and segmentation in cluttered scenes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28:1584–1601, Oct 2006.

[14] C. Olsson, F. Kahl, and M. Oskarsson. Branch-and-bound methods for euclidean registration problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(5):783–794, May 2009.

[15] Radu Bogdan Rusu, Nico Blodow, and Michael Beetz. Fast point feature histograms (fpfh) for 3d registration. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 1848–1853, 2009.

[16] Radu Bogdan Rusu, Nico Blodow, Zoltan Csaba Marton, and Michael Beetz. Aligning point cloud views using persistent feature histograms. In *Proceedings of the 21st IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2008.

[17] Federico Tombari, Samuele Salti, and Luigi Di Stefano. Unique signatures of histograms for local surface description. In *Proceedings of the 11th European Conference on Computer Vision Conference on Computer Vision: Part III*, ECCV'10, pages 356–369, 2010.

[18] Jiaolong Yang, Hongdong Li, Dylan Campbell, and Yunde Jia. Go-icp: A globally optimal solution to 3d icp point-set registration. *IEEE Trans. Pattern Anal. Mach. Intell.*, 38:2241–2254, Nov 2016.

[19] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Fast global registration. In *Computer Vision - ECCV 2016*, pages 766–782. Springer International Publishing, 2016.

Author Biography

Faisal Azhar is a Senior Imaging Researcher at HP Labs, Bristol. He joined HP in 2014 after completing a PhD on video based human activity recognition at the University of Warwick. His research interests include 2D/3D computer vision techniques and deep learning approaches for tasks such as recognition, registration and robot-object interaction.