Joint Global ICP for Improved Automatic Alignment of Full Turn Object Scans

Torben Fetzer[1], Gerd Reis[2] and Didier Stricker[1,2]

[1] University of Kaiserslautern
[2] German Research Center for Artificial Intelligence (DFKI)
{torben.fetzer, gerd.reis, didier.stricker}@dfki.de

Abstract. Point cloud registration is an important task in computer vision, computer graphics, robotics, odometry and many other disciplines. The problem has been studied for a long time and many different approaches have been established. In the case of existing rough initializations, the ICP approach is still widely used as the state of the art. Often only the pairwise problem is treated. In case of many applications, especially in 3D reconstruction, closed rotations of sequences of partial reconstructions have to be registered. We show that there are considerable advantages if ICP iterations are performed jointly instead of the usual pairwise approach (Pulli's approach). Without the need for increased computational effort, lower alignment errors are achieved, drift is avoided and calibration errors are uniformly distributed over all scans. The joint approach is further extended into a global version, which not only considers one-sided adjacent scans, but updates symmetrically in both directions. The result is an approach that leads to a much smoother and more stable convergence, which moreover enables a stable stopping criterion to be applied. This makes the procedure fully automatic and therefore superior to most other methods, that often tremble close to the optimum and have to be terminated manually. We present a complete procedure, which in addition addresses the issue of automatic outlier detection in order to solve the investigated problem data independently without any user interaction.

Keywords: Point Cloud Registration · Automatic Alignment · Drift Reduction

1 Introduction

The task of point cloud registration is to align two point sets so that they resemble each other as closely as possible in as many regions as possible. To make this problem well-defined, it must be assumed that the point clouds represent the same scene or at least that sufficiently large parts of the point clouds represent overlapping parts of the scene. Otherwise, no matching areas can be identified and the problem cannot be solved.

A distinction is made between rigid and non-rigid registration. For the rigid case, two point clouds are aligned only by rotation and translation (in some cases also scaling). The appearance and proportions are fully preserved. In contrast, in the non-rigid registration, deformable objects are aligned by non-linear transformations.

In classical computer vision and robotics, rigid alignment is by far the most common case and has been extensively studied. For this purpose, methods have been established,

which simultaneously detect point correspondences and align them iteratively. In particular the method *Iterative Closest Point* (ICP) [3], [18] has to be named, which has been successfully applied for decades and that will also form the basis of our procedure.

A special case, which occurs in many practical applications, is given by a sequence of point clouds, which partially overlap pairwise and whose last point cloud closes up with the first one. In this case, it is no longer a matter of several pairwise registration problems but a global overdetermined registration problem. This is because each point cloud has two neighbors (last and next one) with whom it must be aligned. In the case of real data, such as the partial reconstructions of a 3D scanner, pairwise sequential alignment would usually lead to a drift, i.e. a large gap or too much overlap between the last and first position. This drift occurs when the partial alignment errors and possible calibration errors in the partial reconstructions add up to a large error. To avoid such drift, it is common to apply *Pulli's procedure* [11], which involves aligning and merging opposite pairs of adjacent point clouds. The resulting merged larger point clouds are then further treated together. In this way, the error is not concentrated between two scans and the drift is distributed over a larger number of scans.

In this work we will show that there are nevertheless better ways with much better properties to solve the global alignment problem in a stable way and to actually distribute the drift evenly without higher computational effort. We will show that a joint iteration of the pairwise registrations distributes the drift uniformly and achieves lower alignment errors than the state of the art. Furthermore, we present how the standard procedure can be extended into a global procedure by symmetrically registering each scan with the next and last scan in the sequence. This results in a global approach that leads to a much smoother convergence, which allows the reliable use of automatic stopping criteria. In contrast, standard procedures usually begin to tremble near the minimum, which often requires a manual termination of the iterations. Finally, a practical approach to the automatic detection of outliers is presented. This is to provide a complete stable solution to the problem without any user interaction. To allow maximum reproducibility, the entire procedure is attached as pseudo-code at the end of the paper.

2 Related Work

The problem of point cloud registration has been well studied for several decades. Explicit methods for rigidly aligning given point correspondences from two data sets have already been developed in the last century [2], [15]. They are based on the singular value decomposition and are still the basis of the modern state of the art. In [1] these methods are robustified by additional weights, based on the certainty of the correspondences and in [5] they are extensively evaluated and compared with other approaches.

For many applications, there is more than two views to be aligned. In order to treat multiple point clouds jointly, an extension of the orthogonal procrustes problem has been introduced in [13]. In [14] the orthogonal constraints are relaxed and the jointly obtained solutions are projected to the space of the orthogonal matrices afterwards. [10] transferred the problem into a simple semidefinite programming in order to ease solving. These methods are no longer explicit and require higher computational effort. In the context of point cloud alignment performed in the upcoming task, the given correspon-

dences are erroneous approximations and change from iteration to iteration. Therefore, a higher accuracy to the costs of additional internal iterations is not reasonable.

Usually, no exact point correspondences are available. A famous principle proven in practice is *Iterative Closest Point* [3], [11], [18], which iteratively selects the closest points of the data sets as correspondences and calculates infinitesimal updates accordingly. There are also variants that take the normals of the point clouds into account [8], [6] and thus improve the alignment for badly sampled and very smooth objects.

In order to accelerate convergence of the methods, the point clouds are either adeptly sampled [12], [6] and outliers efficiently detected and rejected [17], [4], [12] or the iteration updates extrapolated like in [16]. There are also methods to accelerate by a multi-resolution approach [7] or recently by Anderson acceleration [9].

3 Background: Rigid Point Cloud Alignment

The most common algorithm for rigidly aligning point clouds is *Iterative Closest Point*. Thereby, the closest points of two point sets are chosen as correspondences and optimally aligned with each other. Afterwards, new correspondences are chosen, based on the improved alignment. Iteratively, the alignment of the point clouds is improved. For given point correspondences there is a closed form of the optimal rotation matrix and the translation vector for the pairwise case (*Procrustes Analysis*). Since this is also the basis of the method presented in the following and in order to make the paper independently, we will briefly present the procedure for the pairwise case and then show how the method is applied to a full turn according to the current state of the art.

3.1 Orthogonal Procrustes Problem

Assume two sets of point clouds $P = \{\mathbf{p}_0, ..., \mathbf{p}_{N-1}\}$ and $P' = \{\mathbf{p}'_0, ..., \mathbf{p}'_{N-1}\}$ consisting of matching point pairs $\mathbf{p}_n \leftrightarrow \mathbf{p}'_n$, for $n = 0, ..., N-1$ are given. The task is to find an optimal rotation matrix $\mathbf{R}$ and translation vector $\mathbf{t}$ in order to align points $\mathbf{p}_n$ from P by $\mathbf{R}\mathbf{p}_n + \mathbf{t}$ to points $\mathbf{p}'_n$ from P'. Therefore, the sum of Euclidean distances between all point pairs is minimized:

$$\underset{\mathbf{R},\,\mathbf{t}}{\operatorname{argmin}} \sum_{n=0}^{N-1} \|\mathbf{p}'_n - \mathbf{R}\mathbf{p}_n - \mathbf{t}\|_2^2 \tag{1}$$

Setting the derivative of (1) with respect to translation vector $\mathbf{t}$ equal to zero leads to the minimizing $\mathbf{t}$ of the energy:

$$\mathbf{t} = \frac{1}{N}\sum_{n=0}^{N-1} \mathbf{p}'_n - \mathbf{R}\frac{1}{N}\sum_{n=0}^{N-1} \mathbf{p}_n = \mu_{P'} - \mathbf{R}\mu_P \tag{2}$$

Thereby μ_P and $\mu_{P'}$ denote the centroids of the point clouds computed by the mean of the point sets. Inserting Eq. (2) into the problem formulation (1) decouples the problem. It is equivalent to aligning point clouds with zero centroids by optimal rotation only:

$$\underset{\mathbf{R}}{\operatorname{argmin}} \sum_{n=0}^{N-1} \|\mathbf{q}'_n - \mathbf{R}\mathbf{q}_n\|_2^2, \quad \begin{matrix}\text{with} & \mathbf{q}_n = \mathbf{p}_n - \mu_P \\ \text{and} & \mathbf{q}'_n = \mathbf{p}'_n - \mu_{P'}\end{matrix} \tag{3}$$

Calculating the norm explicitly and replacing the remaining scalar product by the trace formulation leads to the following formulation of the problem, that can be solved in terms of the singular value decomposition of matrix $\mathbf{H}$. The validity of this optimizer can be shown by application of *Cauchy-Schwartz Inequality*.

$$\underset{\mathbf{R}}{\operatorname{argmax}} \operatorname{Tr}\big(\mathbf{R} \sum_{n=0}^{N-1} \mathbf{q}_n \mathbf{q'}_n^{\mathsf{T}}\big) = \underset{\mathbf{R}}{\operatorname{argmax}} \operatorname{Tr}(\mathbf{RH}) \quad \rightarrow \mathbf{R} = \mathbf{V}\mathbf{U}^{\mathsf{T}}, \text{ with } \mathbf{H} = \mathbf{U}\boldsymbol{\Lambda}\mathbf{V}^{\mathsf{T}} \quad (4)$$

Weighted Case When working with real data it is usual to apply certainty weights $w_n \geq 0$ with $\sum_{n=0}^{N-1} w_n = 1$ with respect to the point pairs to the alignment error (1) in order to robustify the approach. The problem is solved similarly to the unweighted case using weighted versions of the centroids and matrix $\mathbf{H}$ (see Algorithm 1).

3.2 Iterative Closest Point (ICP)

Usually, no point correspondences are available between two point clouds. In the procedure of *ICP*, these are approximately chosen in each iteration as the nearest points between the data sets and infinitesimal updates are calculated by *Orthogonal Procrustes Analysis*. Since it is a least-squares formulation, it is important to assess the quality of the correspondence. For this purpose, outliers, i.e. points that obviously have no matches in the other data set, are rejected. All other points are weighted according to their quality, which is often done by the point-to-point distance. Sampling rates and methods can also have a strong influence on performance and should not be disregarded.

Initialization For this procedure to work, an initial alignment is urgently required. This prevents the procedure from getting stuck in a local minimum. Based on feature points in the object, be it from texture or geometry, usually a few matches can be found which allow a rough alignment of the point clouds as initialization. Based on this initial registration, the ICP algorithm has proven over a long period of time to be a good choice for refining the alignment.

3.3 Full Turn Registration: Pulli's Approach

In a variety of practical applications, full turns of overlapping partial reconstructions are captured as depicted in Figure 1. Usually, the last scan overlaps with the first one and therefore completes the reconstruction process. In sequential pairwise registrations of the scans, a drift error between the last and the first position often occurs. To avoid or at least reduce this drift error, Pulli's approach [11] has always been the undisputed state of the art. One after the other, scans are registered and merged with their neighbors. These merged point clouds are then registered again until the whole object is composed. In fact, the error is distributed more evenly than in the naive approach and is not added up to a single gap, but it is far from uniform. While the first registration procedures only contain the local alignment errors, the last step combines the alignment errors of several sub-alignments and possible calibration errors.

Fig. 1: Partial scans of a full turn and the aligned point cloud (middle).

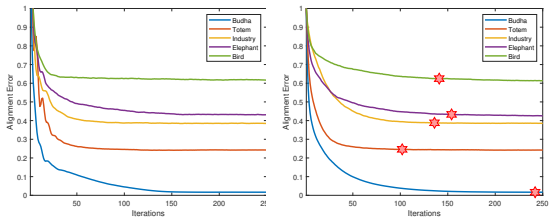

Fig. 2: Alignment errors of the proposed joint sequential (left) and global (right) ICP variants. Red stars mark automatic stopping points of the global method.

4 Joint Rigid Point Cloud Alignment

In the following, the alignment problem of a full rotation of scans is formulated as common optimization problem. We assume that two successive scans have at least some overlap and that the last scan closes up to the first one, thus well defining the problem.

Joint Minimization Problem Let be given a full turn consisting of S scans $\{\mathcal{S}_0, ..., \mathcal{S}_{S-1}\}$, where the last position $\mathcal{S}_{S-1}$ is assumed to be overlapping with the first one $\mathcal{S}_0$. Between two subsequent scans, say scan s and scan $s+1$, we assume to have N point matches each, given by $\mathbf{p}_n^{(s,s+1)} \leftrightarrow \mathbf{p}_n^{(s+1,s)}$, for $n = 0, ..., N-1$. The objective error function, that has to be minimized is then given by

$$\underset{\mathbf{R}^{(s)},\, \mathbf{t}^{(s)}}{\operatorname{argmin}} \sum_{s=0}^{S-1} \sum_{n=0}^{N-1} \|\mathbf{R}^{(s+1)}\mathbf{p}_n^{(s+1,s)} + \mathbf{t}^{(s+1)} - \mathbf{R}^{(s)}\mathbf{p}_n^{(s,s+1)} - \mathbf{t}^{(s)}\|_2^2. \tag{5}$$

Note that we assume a periodic arrangement, so that the scans' indices are treated modulo S, which means $S \equiv 0$. Setting the partial derivative with respect to any translation vector $\mathbf{t}^{(s)}$ equal to zero, we get:

$$2\mathbf{t}^{(s)} - \mathbf{t}^{(s-1)} - \mathbf{t}^{(s+1)} = \mathbf{R}^{(s-1)}\mu_{s-1,s} + \mathbf{R}^{(s+1)}\mu_{s+1,s} - \mathbf{R}^{(s)}(\mu_{s,s-1} + \mu_{s,s+1}) \tag{6}$$

which is sufficiently fulfilled for

$$\mathbf{t}^{(s+1)} - \mathbf{t}^{(s)} = \mathbf{R}^{(s)}\mu_{s,s+1} - \mathbf{R}^{(s+1)}\mu_{s+1,s}. \tag{7}$$

Therefore, the objective function (5) can be decoupled to:

$$\underset{\mathbf{R}^{(s)}}{\operatorname{argmin}} \sum_{s=0}^{S-1} \sum_{n=0}^{N-1} \|\mathbf{R}^{(s+1)}(\mathbf{p}_n^{(s+1,s)} - \mu_{s+1,s}) - \mathbf{R}^{(s)}(\mathbf{p}_n^{(s,s+1)} - \mu_{s,s+1})\|_2^2 \tag{8}$$

$$= \underset{\mathbf{R}^{(s)}}{\operatorname{argmin}} \sum_{s=0}^{S-1} \sum_{n=0}^{N-1} \|\mathbf{R}^{(s+1)}\mathbf{q}_n^{(s+1,s)} - \mathbf{R}^{(s)}\mathbf{q}_n^{(s,s+1)}\|_2^2 \tag{9}$$

This work was partially funded by the projects MARMORBILD (03VP00293) and VIDETE (01W18002) of the German Federal Ministry of Education and Research (BMBF).

Joint Sequential ICP Solving the terms of the joint minimization problem (9) sequentially for one s after the other by simply applying the standard strategy (4) leads to the pairwise approach with joint iterations. Iteratively, closest points between each neighboring pair are chosen and alignment updates by *Orthogonal Procrustes Problem* (4) are applied to each pair. This procedure already avoids drift and the errors are uniformly distributed without additional computational effort.

Joint Global ICP In order to derive a global formulation, that does not only take pairwise point clouds into account, but also treats the global arrangement information, we are minimizing functional (9) with respect to each $\mathbf{R}^{(s)}$ while fixing the others. This is equivalent to solving the following optimization problem:

$$\begin{aligned} &\underset{\mathbf{R}^{(s)}}{\operatorname{argmax}} \sum_{s=0}^{S-1} \operatorname{Tr}(\mathbf{R}^{(s)}\mathbf{H}_{s,s+1}\mathbf{R}^{(s+1)^\mathsf{T}}) \\ &= \underset{\mathbf{R}^{(s)}}{\operatorname{argmax}} \operatorname{Tr}(\mathbf{R}^{(s)}(\underbrace{\mathbf{H}_{s,s+1}\mathbf{R}^{(s+1)^\mathsf{T}} + \mathbf{H}_{s,s-1}^\mathsf{T}\mathbf{R}^{(s-1)^\mathsf{T}}}_{\mathbf{H}_s})) \end{aligned} \tag{10}$$

This is a form of a symmetric alignment update of scan $\mathcal{S}_S$ towards last and next adjacent scans $\mathcal{S}_{S-1}$ and $\mathcal{S}_{S+1}$. The problem can be solved similar to (4) using singular value decomposition and does not require special treatment.

Efficient Point Matching To efficiently find the nearest points between two point clouds, the use of k-d-trees has been established for a long time. Building them means a not inconsiderable effort, but once they are created, the nearest points can be found in logarithmic time. A special feature is, that for each point cloud of a scan only one tree has to be set up, which can be further used after transformation by applying the inverse transformation to the input points, as shown in Algorithm 1. Especially in the iterative application to large point sets, this means an enormous time saving.

5 Outlier Rejection

In order to achieve an automatic procedure that can be applied to a possibly large number of configurations, outliers must be reliably detected in every set of correspondences. Standard procedures, such as rejecting the 10% of correspondence with largest point-to-point distances in each iteration, are widely used, but rely on a well-chosen value. In order to be independent of a fixed value, we have carried out investigations on a large number of data sets.

The task is to separate a set of N point correspondences into two subsets. The separation should divide the outliers as well as possible from the eligible correspondences. Let $D = \{d_0, ..., d_{N-1}\}$ be the set of point-to-point distances of respective correspondences, sorted in a descending order ($d_0 \geq d_1 \geq ... \geq d_{N-1}$). Tests on approximately 25000 different point sets and configurations have shown that a good partition

$$D = D_{\text{outliers}} \cup D_{\text{inliers}} = \{d_0, ..., d_t\} \cup \{d_{t+1}, ..., d_{N-1}\} \tag{11}$$

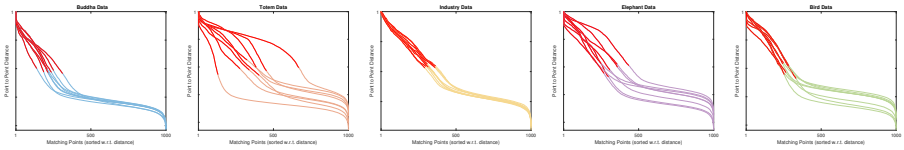

Fig. 3: Outlier rejection strategy (12) applied to the point clouds shown in Figure 4 bottom right. The lines result from point-to-point distances of matches sorted in descending order. The red segments visualize the detected subset of outliers.

is achieved at a split point $t \in \{0, ..., N-1\}$ if the *coefficients of variation* of both subsets is equal or as close as possible. Therefore, t can be successively increased until the following equation holds approximately true:

$$\frac{\frac{1}{t+1}\sum_{n=0}^{t} d_n^2}{\left(\frac{1}{t+1}\sum_{n=0}^{t} d_n\right)^2} = \frac{\frac{1}{N-t}\sum_{n=t+1}^{N-1-1} d_n^2}{\left(\frac{1}{N-t-1}\sum_{n=t+1}^{N-1} d_n\right)^2} \tag{12}$$

Figure 3 shows the behavior of the rejection strategy for the point sets that are evaluated in the upcoming section (see Fig. 4, bottom right). Each reconstruction consists of eight partial scans as it usually comes up from 3D scanners. Between each adjacent pair of scans, matches are computed and outliers are detected by the proposed strategy. Each of the subplots in Fig. 3 shows the 8 curves which result from sorting 1000 matches between each of the 8 point pairs. The red segments visualize the detected set of outliers. For better visualization the plots are given in a logarithmic scale and normalized.

6 Evaluation

For perfect artificial data or uniformly added noise, all alignment strategies work satisfactorily. The situation is different for the real use case of recorded data. In the following we evaluate the considered ICP methods for registration of full turns on a number of sample data sets as they appear from typical 3D scanners. For five independent objects (*Buddha*, *Totem*, *Industry*, *Elephant*, *Bird*), full rotations of eight partial reconstructions each were created. In order to fully align them, the registration methods must be able to deal with both, local alignment errors of the partial point clouds and calibration errors that can have an impact on the overall fit. The standard procedures were compared to the presented joint ICP variants. Figure 4 bottom right shows the resulting aligned point clouds of the *Joint Global ICP* approach to represent the objects under investigation.

The plots in Figure 4 show the convergence behaviour of the alignment strategies for increasing numbers of iterations. Both methods that were proposed converge to significantly lower errors than the trivial sequential pairwise alignment (black line) and Pulli's drift preventing procedure (red line) which reflects the current state of the art. Although both, the *Joint Sequential ICP* and the *Joint Global ICP* converge to the same optimum, the alignment error of the sequential variant occasionally alternates depending on the data (see *Totem*). In contrast, the global approach converges completely smoothly and evenly, which leads to a more stable convergence in general.

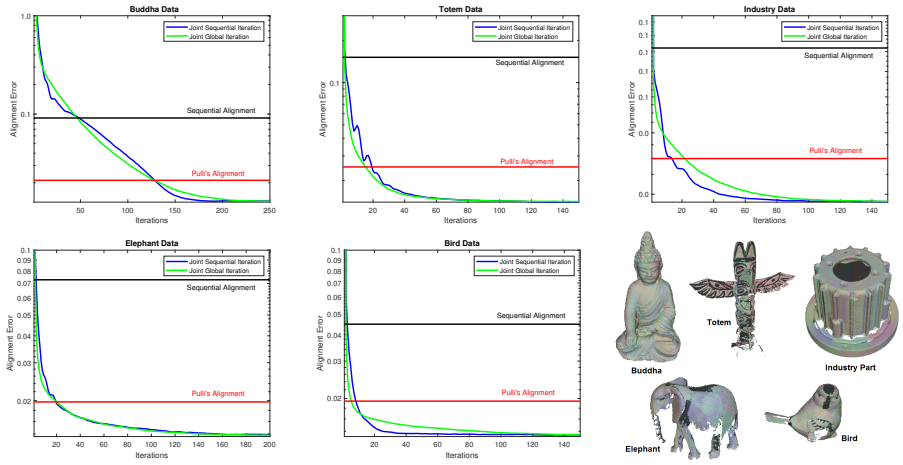

Fig. 4: Convergence behavior for five independent data sets. Alignment error of the naive pairwise approach is given by the black line (contains drift error). State of the art is given by Pulli's approach (red line). Jointly iterating approaches converge to much lower errors. While the sequential procedure (blue) may alternate depending on the data, the global approach (green) converges smoothly.

6.1 Stopping Criterion

The smooth convergence behaviour of the presented global ICP variant provides a considerable advantage over all previous ICP methods. Most of them alternate during the procedure, due to iteratively updated point correspondences, which increases the chance of getting stuck in local minima.

Moreover, the error often starts to alternate around the minimum. Most papers write "we iterate until the alignment error does not reasonable improve any more" without further information. Standard stopping criteria for convergence do not hold in most situations, since the differences between two subsequent iterations may not fall under a given threshold. This is the reason why in many practical implementations the alignment does continue and start to tremble until it is manually stopped.

Figure 2 shows the behavior of the weighted alignment errors for the investigated data sets. Left plot shows the behavior for the *Joint Sequential ICP* and right plot for *Joint Global ICP*. A simple smooth stopping strategy like checking for improvements of the alignment within the last few iterations enables the stable automatic termination of the procedure after a reasonable number of iterations. The stopping points are visualized by the stars in Figure 2.

7 Conclusion

In this paper we presented a procedure that aligns complete closed turns of partial point clouds jointly in a global manner. Not only pairwise adjacent point clouds are considered but also corrected symmetrically to all neighbors. The usual, widely spread

ICP procedure can be applied in a slightly adapted way. By iterating the subproblems jointly, alignment and calibration errors are evenly distributed over all scans and drift

Algorithm 1: Joint Global ICP for Full Turn Alignment

Input: Initially aligned partial point clouds $P_0, ..., P_{S-1}$ of scans $\mathcal{S}_r, ..., \mathcal{S}_{S-1}$.

Initialize parameters $\mathbf{R}^{(s)} = \mathbf{I}$, $\mathbf{t}^{(s)} = \mathbf{0}$ for all partial scans $\mathcal{S}_s$, $s = 0, ..., S-1$.

Setup a k-d-tree $\mathcal{T}_s$ for each point cloud P_s.

Sample point clouds P_s to a size of N elements.

for $i = 0, 1, 2, ...$ **do**

  for $s = 0, ..., S-1$ **do**

    Search in k-d-trees of adjacent scans for correspondences:

$$\mathcal{T}_{s-1}\big(\mathbf{R}^{(s-1)^\mathsf{T}}(\mathbf{R}^{(s)}P_s + \mathbf{t}^{(s)} - \mathbf{t}^{(s-1)})\big) \rightarrow P^{(s,s-1)}$$

$$\mathcal{T}_{s+1}\big(\mathbf{R}^{(s+1)^\mathsf{T}}(\mathbf{R}^{(s)}P_s + \mathbf{t}^{(s)} - \mathbf{t}^{(s+1)})\big) \rightarrow P^{(s,s+1)}$$

    Reject outliers as introduced in Sec. 5.

    Weight correspondences with respect to point distances.

    Subtract centroids from point sets:

$$\mu_{s,s-1} = \sum_{n=0}^{N-1} w_n \mathbf{p}_n^{s,s-1} \quad \rightarrow Q^{(s,s-1)} = P^{(s,s-1)} - \mu_{s,s-1}$$

$$\mu_{s,s+1} = \sum_{n=0}^{N-1} w_n \mathbf{p}_n^{s,s+1} \quad \rightarrow Q^{(s,s+1)} = P^{(s,s+1)} - \mu_{s,s+1}$$

  end

  for $s = 0, ..., S-1$ **do**

    Set up symmetric system matrices:

$$\mathbf{H}_{s,s-1} = \sum_{n=0}^{N-1} w_n \mathbf{q}_n^{(s,s-1)} \mathbf{q}_n^{(s-1,s)^\mathsf{T}}, \quad \mathbf{H}_{s,s+1} = \sum_{n=0}^{N-1} w_n \mathbf{q}_n^{(s,s+1)} \mathbf{q}_n^{(s+1,s)^\mathsf{T}}$$

$$\rightarrow \mathbf{H}_s = \mathbf{H}_{s,s+1}\mathbf{R}^{(s+1)^\mathsf{T}} + \mathbf{H}_{s,s-1}^\mathsf{T}\mathbf{R}^{(s-1)^\mathsf{T}}$$

    Compute SVD $\mathbf{H}_s = \mathbf{U}_s \boldsymbol{\Lambda}_s \mathbf{V}_s^\mathsf{T}$ and compose updated rotation $\mathbf{R}^{(s)} = \mathbf{V}_s \mathbf{U}_s^\mathsf{T}$.

  end

  for $s = 0, ..., S-1$ **do**

    Update translation vectors $\mathbf{t}^{(s)}$ by Eq. 6.

  end

  Compute weighted alignment error and check for improvement withing the last say 10 iterations. If no improvement break.

end

Output: Optimal transformations $\mathbf{R}^{(s)}$, $\mathbf{t}^{(s)}$.

is prevented. The global approach leads to a smooth convergence behaviour, which enables the credible application of automatic stopping criteria. Together with an introduced outlier rejection strategy, this results in an extremely stable automatic procedure that generates better results than the previous state of the art. This is moreover achieved without any user interaction or additional computational effort. To ease reproduction, the procedure is finally attached as pseudo-code.

References

1. Akca, D.: Generalized procrustes analysis and its applications in photogrammetry. Tech. rep., ETH Zurich (2003)
2. Arun, K.S., Huang, T.S., Blostein, S.D.: Least-squares fitting of two 3-d point sets. IEEE Transactions on pattern analysis and machine intelligence (5), 698–700 (1987)
3. Besl, P.J., McKay, N.D.: Method for registration of 3-d shapes. In: Sensor fusion IV: control paradigms and data structures. vol. 1611, pp. 586–606. International Society for Optics and Photonics (1992)
4. Dorai, C., Wang, G., Jain, A.K., Mercer, C.: Registration and integration of multiple object views for 3d model construction. IEEE Transactions on pattern analysis and machine intelligence **20**(1), 83–89 (1998)
5. Eggert, D.W., Lorusso, A., Fisher, R.B.: Estimating 3-d rigid body transformations: a comparison of four major algorithms. Machine vision and applications (1997)
6. Gelfand, N., Ikemoto, L., Rusinkiewicz, S., Levoy, M.: Geometrically stable sampling for the icp algorithm. In: Fourth International Conference on 3-D Digital Imaging and Modeling, 2003. 3DIM 2003. pp. 260–267. IEEE (2003)
7. Jost, T., Hügli, H.: Fast icp algorithms for shape registration. In: Joint Pattern Recognition Symposium. pp. 91–99. Springer (2002)
8. Masuda, T., Sakaue, K., Yokoya, N.: Registration and integration of multiple range images for 3-d model construction. In: Proceedings of 13th international conference on pattern recognition. vol. 1, pp. 879–883. IEEE (1996)
9. Pavlov, A.L., Ovchinnikov, G.W., Derbyshev, D.Y., Tsetserukou, D., Oseledets, I.V.: Aa-icp: Iterative closest point with anderson acceleration. In: IEEE International Conference on Robotics and Automation (ICRA). pp. 1–6 (2018)
10. Pizarro, D., Bartoli, A.: Global optimization for optimal generalized procrustes analysis. In: CVPR. IEEE (2011)
11. Pulli, K.: Multiview registration for large data sets. In: Second International Conference on 3-D Digital Imaging and Modeling (Cat. No. PR00062). pp. 160–168. IEEE (1999)
12. Rusinkiewicz, S., Levoy, M.: Efficient variants of the icp algorithm. In: Proceedings third international conference on 3-D digital imaging and modeling. IEEE (2001)
13. Ten Berge, J.M.: Orthogonal procrustes rotation for two or more matrices. Psychometrika **42**(2), 267–276 (1977)
14. Trendafilov, N.T., Lippert, R.A.: The multimode procrustes problem. Linear algebra and its applications (2002)
15. Umeyama, S.: Least-squares estimation of transformation parameters between two point patterns. IEEE Transactions on Pattern Analysis & Machine Intelligence (1991)
16. Yaniv, Z.: Rigid registration: The iterative closest point algorithm. School of Engineering and Computer Science, The Hebrew University, Israel (2001)
17. Zhang, Z.: Iterative point matching for registration of free-form curves and surfaces. International journal of computer vision **13**(2), 119–152 (1994)
18. Zinßer, T., Schmidt, J., Niemann, H.: Point set registration with integrated scale estimation. In: International conference on pattern recognition and image processing (2005)