JOINT PRE-ALIGNMENT AND ROBUST RIGID POINT SET REGISTRATION

Vladislav Golyanik, Bertram Taetz and Didier Stricker

University of Kaiserslautern, Germany
German Research Center for Artificial Intelligence (DFKI), Germany
{*Vladislav.Golyanik, Bertram.Taetz, Didier.Stricker*}*@dfki.de*

ABSTRACT

We present an elegant solution to joint pre-alignment and rigid point set registration, given prior matches. Instead of performing pre-alignment and the actual registration in the separate steps, prior matches explicitly influence the registration procedure in our approach. This results in several advantages. Firstly, our approach solves the pre-alignment task — an approximate resolving of rotation and translation — with an insufficient number of prior correspondences, when other methods fail. Secondly, it produces more accurate rigid registrations of noisy point sets than the state of the art Coherent Point Drift method. Combined with application specific methods for correspondence establishment, we demonstrate superiority of our approach in several synthetic and real-world scenarios.

Index Terms— Point set registration, image registration, prior matches, Coherent Point Drift

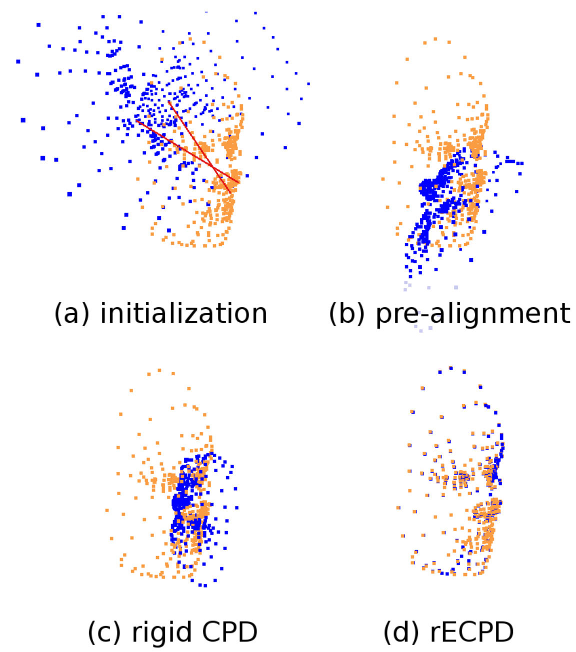

Fig. 1: Rigid registration of two point clouds (representing human faces) with prior matches. (a): input with prior matches shown as red connecting lines; (b): pre-alignment with Kabsch algorithm; (c): registration results of the pre-aligned point sets by CPD; (d): registration results by our method, the rECPD. Given only two priors, our algorithm is capable of resolving rotation and rigid transformation correctly, whereas Kabsch algorithm fails to pre-estimate rotation due to ambiguities in the solution space and CPD fails to register point sets correctly due to severe initial misalignment.

1. INTRODUCTION

The problem of point set registration, i.e. aligning point sets into a common coordinate frame, often arises in image processing and computer vision. Medical image registration, shape recognition and CAD modelling are notable application examples amongst many others. If an underlying transformation between a *reference* and a transformed *template* point set is entirely described by parameters of the rigid body motion, the problem is specified as *rigid point set registration*. A detailed overview of approaches for rigid point set registration can be found in [1, 2].

Generally, registration algorithms perform well on point sets representing noiseless objects. When this assumption is violated — (1) point sets represent partially overlapping parts of an object, (2) contain outliers or (3) differences in scene poses are significant — difficulties arise. All the violations often happen in practice. Although some approaches can handle different types of noise distributions [2, 3], it turns out that outliers often do not obey a particular probability distribution, but are rather *clustered*. Clustered outliers and regular points in a point set can not be distinguished from each other during the registration. For instance, the main advantage of mixture model based methods [4, 2] over Iterative Closest Point (ICP) methods [5, 6, 7] lies in the soft assignment of correspondences via probabilities, proving to be more robust in the presence of outliers. However, if outliers are clustered, the methods may fail (see §3). Partial shapes can also be comprehended as point sets with large areas of clustered outliers. Consequently, violation (1) is a special case of violation (2).

To compensate for the aforementioned violations, preprocessing steps are required. To cope with clustered outliers, weighting the correspondences with respect to overlapping parts is possible [7]. However, the performance of method

[7] strongly depends on the initialization. To cope with severe misalignment between point sets, registration is usually performed on a subset of reliable key points obtained by key point detectors [8, 9]. Through evaluation and comparison of the key points, correspondences are established [10]. Thus, the problem reduces to a transformation estimation problem, since *all* correspondences are known and the number of points in both sets is equal. An overview and comparison of methods for solving transformation estimation problems can be found in [11].

As already mentioned, often one or several correspondences between the template and the reference — prior matches — are known in advance. They can constraint the solution space in favourable way ensuring convergence to a disired range. At the same time, integration of prior matches is not straightforward: either prior matches are used in the pre-alignment step to assure an advantageous initial orientation, or a registration algorithm allows embedding of prior matches explicitly. In the first case prior matches are decoupled from the registration procedure. Considering probabilistic point set registration methods, especially the state of the art Coherent Point Drift algorithm [2], effective embedding of prior matches was recently shown in the literature for the non-rigid case [12]. To investigate embedding of prior matches into a rigid probabilistic point set registration algorithm is the main contribution of this paper. More specifically: (1) we derive a solution for embedding prior correspondences into the CPD algorithm. This results in a robust and convenient approach, the *rigid Extended Coherent Point Drift* (rECPD) combining rough pre-alignment and robust dense rigid point set registration; (2) we reveal the synergy arising from the interdependency between the generic rigid registration and predominance of particular correspondences. If number of prior matches $N_c < D$ (D is space dimensionality), instead of barely taking insufficient amount of correspondences to pre-estimate rotation, our method additionally exploits information contained in the point sets.

2. RIGID EXTENDED COHERENT POINT DRIFT

We extend CPD algorithm [2] by embedding prior correspondences, deliberately influencing the rigid point set registration. Given a set of prior matches $(\mathbf{y}_j, \mathbf{x}_k)$, $j, k \in N_c \subset \mathbb{N}^2$, we model them by the product of independent density functions

$$P_c(N_c) = \prod_{(j,k)\in N_c} p_c(\mathbf{x}_j, \mathbf{y}_k) \quad \text{with} \tag{1}$$

$$p_c(\mathbf{x}_j, \mathbf{y}_k) = \frac{1}{(2\pi\alpha)^{D/2}} \exp\left(-\frac{\|\mathbf{x}_j - T(\mathbf{y}_k, \theta)\|^2}{2\alpha^2}\right). \tag{2}$$

The parameter $0 < \alpha < 1$ reflects the prior's degree of reliability. We embed prior matches into the CPD by superposition of the prior probability $P_c(N_c)$ and the GMM. The modified probability density function for a particular point x reads

$$\tilde{p}(\mathbf{x}) = P_c(N_c) p(\mathbf{x}). \tag{3}$$

Taking the negative logarithm of the combined modified GMM we obtain

$$\begin{aligned}\tilde{E}(\theta, \sigma^2) &= -\log\left(P_c(N_c)\prod_{i=1}^{N} p(\mathbf{x}_n)\right) = \\ &= E(\theta, \sigma^2) - \sum_{(j,k)\in N_c} \log(p_c(\mathbf{x}_j, \mathbf{y}_k))\end{aligned} \tag{4}$$

— the modified energy function. Rewriting the last term of Eq. (4) and leaving out the constant terms w.r.t. θ and σ^2, the modified objective function reads

$$\begin{aligned}\tilde{Q} &= Q + \frac{1}{2\alpha^2} \sum_{(j,k)\in N_c} \|\mathbf{x}_j - T(\mathbf{y}_k, \theta)\|^2 = \\ &= Q + \frac{1}{2\alpha^2} \sum_{n=1}^{N}\sum_{m=1}^{M} \tilde{P}_{m,n} \|\mathbf{x}_n - T(\mathbf{y}_m, \theta)\|^2,\end{aligned} \tag{5}$$

where the sparse selection matrix $\tilde{\mathbf{P}}_{M\times N}$ contains 1 in the (j, k) entry, if $(j, k) \in N_c$. To specify the parameter set θ we impose the rules of rigid body dynamics on the GMM centroids. Thus, the rECPD objective function reads

$$\begin{aligned}\tilde{Q}(\mathbf{R}, \mathbf{t}, s, \sigma^2) &= \frac{1}{2\sigma^2} \sum_{n=1}^{N}\sum_{m=1}^{M} P^{old}(m|\mathbf{x}_n) \|\mathbf{x}_n - s\mathbf{R}y_m - \mathbf{t}\|^2 + \\ &+ \frac{N_p D}{2}\log\sigma^2 + \frac{1}{2\alpha^2}\sum_{n=1}^{N}\sum_{m=1}^{M} \tilde{P}_{m,n}\|\mathbf{x}_n - s\mathbf{R}y_m - \mathbf{t}\|^2,\end{aligned} \tag{6}$$

so that $\mathbf{R}^T\mathbf{R} = \mathbf{I}$, $\det(\mathbf{R}) = +1$ and $N_p = \sum_{n,m} P^{old}(m|\mathbf{x}_n)$. We find minimizer of the rECPD objective function in Eq. (6) taking advantage of lemma 1 from [13]. First, we reformulate it to match the form $\mathrm{tr}(\mathbf{A}^T\mathbf{R})$. To eliminate the translation term from $\tilde{Q}$, we compute its first derivative w.r.t. $\mathbf{t}$ and equate it to zero. This yields the modified term

$$\mathbf{t} = \frac{\left[\mathbf{X}^T\mathbf{P}^T\mathbf{1} + \frac{\sigma^2}{\alpha^2}\mathbf{X}^T\tilde{\mathbf{P}}^T\mathbf{1} - s\mathbf{R}(\mathbf{Y}^T\mathbf{P}\mathbf{1} + \frac{\sigma^2}{\alpha^2}\mathbf{Y}^T\tilde{\mathbf{P}}\mathbf{1})\right]}{N_P^c}, \tag{7}$$

where $\mathbf{1} = (1, 1, \ldots, 1)^T_{M\times 1}$ and $\mathbf{P}$ is a matrix with elements $p_{m,n} = P^{old}(m|x_n)$ (P^{old} is computed as in the original CPD). From Eq. (7) follows

$$N_P^c = \mathbf{1}^T\mathbf{P}\mathbf{1} + \frac{\sigma^2}{\alpha^2}\mathbf{1}^T\tilde{\mathbf{P}}\mathbf{1} = N_P + \frac{\sigma^2}{\alpha^2} N_{P_c} \tag{8}$$

with $N_{P_c} = \sum_{(j,k)\in N_c} p_c(\mathbf{x}_j, \mathbf{y}_k)$. Considering the mean vectors

$$\mu_x = \mathbf{E}(X) = \frac{\mathbf{X}^T\mathbf{P}^T\mathbf{1}}{N_P}, \quad \mu_y = \mathbf{E}(Y) = \frac{\mathbf{Y}^T\mathbf{P}\mathbf{1}}{N_P} \tag{9}$$

we define the modified mean vectors

$$\mu_x^c = \frac{1}{N_P^c}\left[N_P\mu_x + \frac{\sigma^2}{\alpha^2}\mathbf{X}^T\tilde{\mathbf{P}}^T\mathbf{1}\right] \text{ and} \tag{10}$$

$$\mu_y^c = \frac{1}{N_P^c}\left[N_P\mu_y + \frac{\sigma^2}{\alpha^2}\mathbf{Y}^T\tilde{\mathbf{P}}\mathbf{1}\right]. \tag{11}$$

Algorithm 1 Rigid Extended Coherent Point Drift

Input: Reference and template point sets $\mathbf{X}, \mathbf{Y}$, prior matches N_c

Initialization: $\mathbf{R} = \mathbf{I},\ \mathbf{t} = \mathbf{0},\ s = 1,\ 0 \leq w \leq 1, \sigma^2 = \frac{1}{DNM} \sum_{n=1}^{N} \sum_{m=1}^{M} \|x_n - y_m\|^2$, $\tilde{\mathbf{P}}$ precomputed as in Eq. (5).
EM-optimization
repeat
  E-step: compute $\mathbf{P}$ with $p_{m,n} = P^{old}(m|\mathbf{x}_n)$ as in the original CPD.
  M-step: solve for $\mathbf{R}, s, \mathbf{t}, \sigma^2$
    N_P^c as in Eq. (8) and μ_x^c, μ_y^c as in Eq. (10).
    $\hat{\mathbf{X}}_c = \mathbf{X} - \mathbf{1}(\mu_x^c)^T, \hat{\mathbf{Y}}_c = \mathbf{Y} - \mathbf{1}(\mu_y^c)^T$
    Compute $\mathbf{R}$ as in Eq. (13).
    $s = \frac{\text{tr}[\mathbf{A}^T\mathbf{R}]}{\text{tr}[\hat{\mathbf{Y}}_c^T d(\mathbf{P1})\hat{\mathbf{Y}}_c + \frac{\sigma^2}{\alpha^2}\hat{\mathbf{Y}}_c^T d(\tilde{\mathbf{P}}\mathbf{1})\hat{\mathbf{Y}}_c]}$ with $\mathbf{A}$ as in Eq. (13).
    Compute $\mathbf{t}$ as in Eq. (7).
    $\sigma^2 = \text{tr}[\hat{X}_c^T d(\mathbf{P1})\hat{X}_c] - s\,\text{tr}[\mathbf{A}^T\mathbf{R}] + s^2\frac{\sigma^2}{\alpha^2}\text{tr}[\hat{Y}_c^T d(\mathbf{P1})\hat{Y}_c]$
until step size of the EM algorithm $dL \leq \epsilon$ or number of iterations exceeds a predefined value.

Output: Aligned point set $T(\mathbf{Y}) = s\mathbf{Y}\mathbf{R}^T + \mathbf{1}\mathbf{t}^T$. Correspondence probabilities are given by $\mathbf{P}$.

By substituting $\mathbf{t}$ from Eq. (7) back into Eq. (6) and further defining centered point set matrices $\hat{\mathbf{X}}_c = \mathbf{X} - \mathbf{1}(\mu_x^c)^T$ and $\hat{\mathbf{Y}}_c = \mathbf{Y} - \mathbf{1}(\mu_y^c)^T$, we rewrite the obtective function as

$$\tilde{Q} = -\underbrace{\frac{s}{\sigma^2}}_{>0} \text{tr}\left[\underbrace{\left((\hat{\mathbf{X}}_c^T\mathbf{P}^T\hat{\mathbf{Y}}_c)^T + \frac{\sigma^2}{\alpha^2}(\hat{\mathbf{X}}_c^T\tilde{\mathbf{P}}^T\hat{\mathbf{Y}}_c)^T\right)}_{=:\mathbf{A}^T}\mathbf{R}\right] \tag{12}$$

(to achieve this form, we also utilize the orthogonality of the rotation matrix and the fact that the trace is linear and invariant under cyclic matrix permutations). Minimization of $\tilde{Q}$ is equal to maximization of tr$(\mathbf{A}^T\mathbf{R})$ as defined in Eq. (12), so that $\mathbf{R}^T\mathbf{R} = \mathbf{I}$ and $\det(\mathbf{R}) = +1$. We apply lemma 1 from [13] and obtain the rotation matrix

$$\mathbf{R} = \mathbf{UCV} \quad \text{with} \tag{13}$$

$$\mathbf{C} = d(1, 1, \dots, \det(\mathbf{U}\mathbf{V}^T)) \quad \text{and} \tag{14}$$

$$\mathbf{USV} = svd\left(\underbrace{\hat{\mathbf{X}}_c^T\mathbf{P}^T\hat{\mathbf{Y}}_c + \frac{\sigma^2}{\alpha^2}(\hat{\mathbf{X}}_c^T\tilde{\mathbf{P}}^T\hat{\mathbf{Y}}_c)}_{=\mathbf{A}}\right). \tag{15}$$

Analogously, in order to obtain the optimal s and σ^2, respective derivatives of $\tilde{Q}$ in Eq. (12) have to be computed and equated to zero. The whole method is summarized in Algorithm 1.

Now we are ready to investigate how rECPD contrasts from most closely related approaches. From CPD with a pre-alignment step our approach differs in the way that prior matches influence the registration procedure explicitly in every EM iteration. If $\alpha = 1$, the prior matches are not valid and our approach reduces to CPD. In opposite, if α positive and infinitely close to zero, rECPD operates similar to the Kabsch algorithm [14]. Indeed, the term $(\hat{\mathbf{X}}_c^T\tilde{\mathbf{P}}^T\hat{\mathbf{Y}}_c)^T$ in Eq. (12) can be rewritten as $(\hat{\mathbf{Y}}_c^T\tilde{\mathbf{P}}\hat{\mathbf{X}}_c)$. At the same time, our approach differs from the Kabsch algorithm in several ways. Firstly, all points of both point sets are involved in the optimization, whereas prior matches are usually several orders of magnitude stronger weighted. Secondly, different weighting and thus the uncertainty level adjustment for every distinct prior is possible. This allows to combine priors from various sources in a flexible manner. Depending on settings and number of prior matches, different effects are possible such as pre-alignment with insufficient number of prior matches or robust registration in presence of clustered outliers.

3. EVALUATION

In this section we evaluate performance of rECPD in synthetic and real-world scenarios. We use publicly available implementation of CPD [2] as well as Matlab's implementation of the Kabsch algorithm [15]. We implemented rECPD in C++ and ran experiments on the system with 32 GB RAM and Intel Xeon E3-1245 processor.

Experiments on Synthetic Data. In an experiment on synthetic data we take the sparse 3D face point set from [2], duplicate it and systematically change orientation of the copy jointly around the x, y and z axes with the angle-step size of 36°. This results in 1000 different initial orientations of the point sets. The scaling factor and the translation vector are chosen randomly. In the preprocessing step we establish prior matches between the point sets. Those are obtained through comparison of the Persistent Feature Histograms (PFHs) [10] at the 3D key points. We find the 3D key points with the Harris3D (H3D) [9] and the Intrinsic Shape Signatures (ISS) [8] 3D key point detectors. Only correspondences with the highest match scores are taken into account. Thus, we obtain two reliable matches in total. Note that in both cases the prior matches relate not exactly the same points. For every initial orientation we perform a preprocessing step with Kabsch algorithm and eventually register the point sets with CPD and rECPD. In the latter case, prior matches are input directly (no preprocessing step undertaken) as one of the algorithm's parameters. In Fig. 1 selected results are shown. CPD is able to restore correct rigid transformations in 242 cases out of 1000 (24% success rate), mostly when absolute values of the individual inclination angles do not exceed 72 degrees. CPD performs as if no rotation pre-estimate would be accomplished, since two prior matches do not suffice to unambiguously pre-estimate rotation in 3D space ($N_c \leq D$). In opposite, rECPD is able to recover correct rigid transformations in all cases, achieving 100% success rate in this experiment. This experiment shows that synergetic effect of joint point set registration and explicit incorporation of prior matches results in the correct transformation in this underconstrained case. In other words, taking into account all other point allows to compensate for missing information.

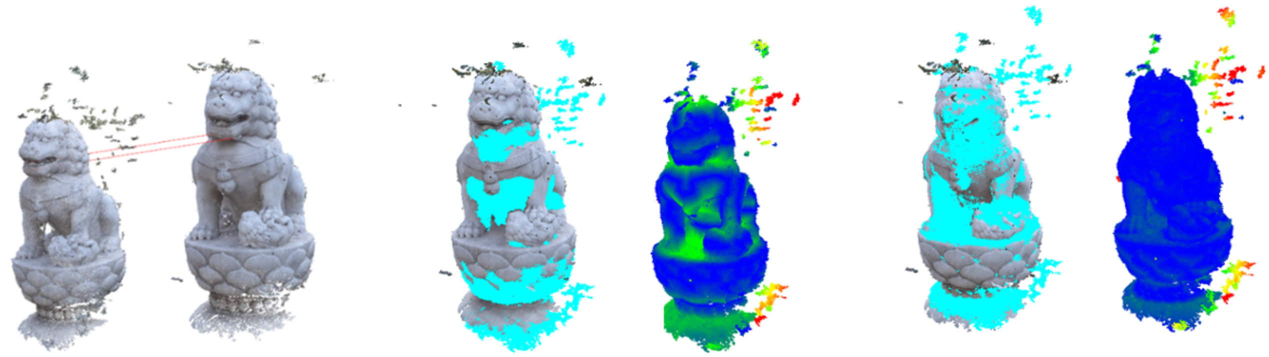

Fig. 2: Rigid registration of Lion data set. (a): inputs with prior matches shown on red; (b): registration results by CPD (left) and cloud-to-cloud distance in Blue<Green<Yellow<Red scale (right), saturation point = 7.0 distance units, (mean error; std. deviation) = (0.940; 0.902); (c): registration results by rECPD and cloud-to-cloud distance in Blue<Green<Yellow<Red scale (right), saturation point = 7.0 distance units, (mean error; std. deviation) = (0.483; 0.855).

Thereafter, the above-stated experiment is repeated in different configurations, swapping reference and template point set and using only one prior match at a time. In 25% cases CPD was able to resolve rotation correctly. Interestingly, rECPC achieves 98% success rate with only one prior match. For the sake of completeness we repeat the experiment also with three prior matches, obtaining 100% success rate for both methods, as expected. In this experiment rECPD — both with one and two matches — clearly outperforms the state-of-the art rigid registration algorithm CPD with the preprocessing step. The results also reveal that the performance of rECPD is equivalent to the performance of CPD with pre-alignment step if $N_c = D$. In this experiment point sets do not contain clustered outliers. Under real world conditions this prerequisite is not always fulfilled and further advantages of rECPD can be observed as shown below.

Experiments on Real Data. In the experiment on real data we evaluate out method on the scans form the Lion data set (Fig. 2, (a); reference is located on the right). Both point clouds represent the same object, namely the statue of a lion. The scans were reconstructed with the PMVS-2 algorithm [16] with different parameters. Therefore, they are noisy and contain areas with clustered outliers constituting 1.5% of points in the reference and 3% of points in the template point clouds. Totally the reference point cloud contains $5 \cdot 10^5$ points and the template point cloud $1.5 \cdot 10^5$ points.

We perform the preprocessing step to obtain prior matches that are used to pre-align the scans (in case of CPD) or as one of the input parameters (in case of rECPD). Correspondence establishment follows similar steps as in the experiment on synthetic data. The ISS key point detector is used and two correspondences with highest scores are taken. Eventually we register scans with CPD and rECPD. The running times amounted to 80 seconds and 72 seconds respectively. Results are shown in Fig. 2. Influence of clustered outliers is substantial and CPD fails to register the scans in the ROI correctly. In opposite, rECPD performs more accurately. Results are also reflected in the corresponding cloud-to-cloud comparisons (Fig. 2, (b), (c), right parts contain plots with root-mean-square errors). The measure reveals that rECPD, provided reliable prior matches, is more robust compared to CPD in presence of clustered outliers in this experiment.

4. CONCLUSIONS

To the best of our knowledge, the proposed method is the first allowing to embed prior matches into a rigid point set registration algorithm. On the one hand, our solution makes rigid point set registration convenient in cases with given prior matches, since no pre-alignment step is required. On the other hand, additional synergy from the explicit embedding emerges. Prior knowledge is sustained during the whole registration procedure and influences it in a favourable way. In scenarios with only one or two prior matches, rECPD clearly outperforms CPD with a pre-alignment step. In scenarios when scans contain clustered outliers, rECPD outperforms CPD with pre-alignment step in terms of cloud-to-cloud distance in ROI. Both statements are supported by the experiments on synthetic and real data. Future work aims at application of the proposed technique in medical scenarios.

5. ACKNOWLEDGEMENTS

Development of rECPD was partially funded by the EU 7th Framework Programme project AlterEgo (600610) and by the BMBF project DYNAMICS (01IW15003).

6. REFERENCES

[1] B. Jian and B. C. Vemuri, "Robust point set registration using gaussian mixture models," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 33, no. 8, pp. 1633–1645, 2011.

[2] A. Myronenko and X. Song, "Point-set registration: Coherent point drift," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2010.

[3] D. Gerogiannis, C. Nikou, and A. Likas, "The mixtures of student's t-distributions as a robust framework for rigid registration," *Image and Vision Computing*, vol. 27, no. 9, pp. 1285 – 1294, 2009.

[4] A. Myronenko, X. Song, and M. Á. Carreira-Perpiñán, "Non-rigid point set registration: Coherent point drift," in *Neural Information Processing Systems (NIPS)*, 2006, pp. 1009–1016.

[5] P. J. Besl and N. D. McKay, "A method for registration of 3-d shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 14, no. 2, pp. 239–256, 1992.

[6] S. Rusinkiewicz and M. Levoy, "Efficient variants of the ICP algorithm," in *3D Digital Imaging and Modeling (3DIM)*, 2001.

[7] A. Fitzgibbon, "Robust Registration of 2D and 3D Point Sets," in *British Machine Vision Conference (BMVC)*, 2001, pp. 43.1–43.10.

[8] Y. Zhong, "Intrinsic shape signatures: A shape descriptor for 3d object recognition.," in *International Conference on Computer Vision Workshops (ICCV Workshops)*, 2009, pp. 689–696.

[9] I. Sipiran and B. Bustos, "Harris 3d: A robust extension of the harris operator for interest point detection on 3d meshes," *The Visual Computer*, vol. 27, no. 11, pp. 963–976, 2011.

[10] R. B. Rusu, N. Blodow, Z. C. Marton, and M. Beetz, "Aligning Point Cloud Views using Persistent Feature Histograms," in *Intelligent Robots and Systems (IROS)*, 2008.

[11] D.W. Eggert, A. Lorusso, and R.B. Fisher, "Estimating 3-d rigid body transformations: a comparison of four major algorithms," *Machine Vision and Applications*, vol. 9, no. 5-6, pp. 272–290, 1997.

[12] V. Golyanik, B. Taetz, G. Reis, and D. Stricker, "Extended coherent point drift algorithm with correspondence priors and optimal subsampling," in *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2016.

[13] A. Myronenko and X. Song, "On the closed-form solution of the rotation matrix arising in computer vision problems," *Computing Research Repository (CoRR)*, vol. abs/0904.1613, 2009.

[14] W. Kabsch, "A solution for the best rotation to relate two sets of vectors," *Acta Crystallographica Section A*, vol. 32, no. 5, pp. 922–923, 1976.

[15] The MathWorks Inc., "*MATLAB version 9.0 (R2016a)*," `http://mathworks.com/products/matlab/`, 2016.

[16] Y. Furukawa and J. Ponce, "Accurate, dense, and robust multiview stereopsis," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 32, no. 8, pp. 1362–1376, 2010.