

LSG-CPD: Coherent Point Drift with Local Surface Geometry for Point Cloud Registration

Weixiao Liu^{1,2} Hongtao Wu^{1,2} Gregory S. Chirikjian^{1,2*}

¹National University of Singapore, Singapore ²Johns Hopkins University, USA

{mpewx1, mpehtw, mpegre}@nus.edu.sg

Abstract

Probabilistic point cloud registration methods are becoming more popular because of their robustness. However, unlike point-to-plane variants of iterative closest point (ICP) which incorporate local surface geometric information such as surface normals, most probabilistic methods (e.g., coherent point drift (CPD)) ignore such information and build Gaussian mixture models (GMMs) with isotropic Gaussian covariances. This results in sphere-like GMM components which only penalize the point-to-point distance between the two point clouds. In this paper, we propose a novel method called CPD with Local Surface Geometry (LSG-CPD) for rigid point cloud registration. Our method adaptively adds different levels of point-to-plane penalization on top of the point-to-point penalization based on the flatness of the local surface. This results in GMM components with anisotropic covariances. We formulate point cloud registration as a maximum likelihood estimation (MLE) problem and solve it with the Expectation-Maximization (EM) algorithm. In the *E* step, we demonstrate that the computation can be recast into simple matrix manipulations and efficiently computed on a GPU. In the *M* step, we perform an unconstrained optimization on a matrix Lie group to efficiently update the rigid transformation of the registration. The proposed method outperforms state-of-the-art algorithms in terms of accuracy and robustness on various datasets captured with range scanners, RGBD cameras, and LiDARs. Also, it is significantly faster than modern implementations of CPD. The source code is available at <https://github.com/ChirikjianLab/LSG-CPD.git>.

1. Introduction

Point cloud registration plays an important role in various computer vision tasks [7, 30, 34, 39]. It is a task of aligning two point clouds and estimating the relative transformation between them. When the correspondences between the points are known, methods have been proposed to solve it efficiently [5, 41]. Our paper focuses on the correspondence-free registration. The most popular method, ICP [2], iteratively computes the alignment from point-to-point corre-

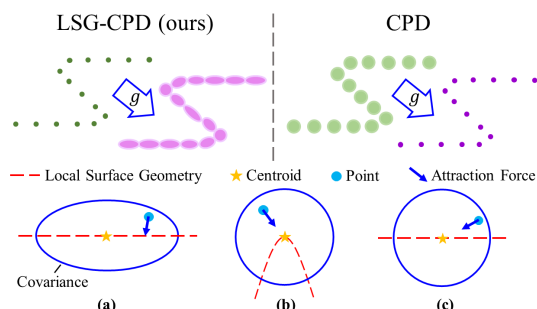


Figure 1. CPD with Local Surface Geometry (LSG-CPD). Green denotes the source set; purple denotes the target set. Unlike CPD [29] which constructs a sphere-like GMM component (c) for each point in the source set, the proposed method constructs GMM components on the target set. In the proposed method, plane-like GMM components (a) are created for points with small surface variation [31], *i.e.*, flat local surface; sphere-like GMM components (b) are created for points with large surface variation.

spondence obtained via the closest point. Various point-to-plane ICP variants [7, 8, 35, 36] make use of surface normals and take advantage of the tendency that most point cloud data are locally planar. However, ICP and its variants require good initialization and are susceptible to outliers, noise, and missing points.

Probabilistic methods formulate point cloud registration as a probability density estimation problem [14, 22, 29, 37]. In particular, methods such as CPD [29] and FilterReg [14] describe the geometry of the first point cloud via a probability distribution over Euclidean space with a GMM. The GMM is constructed by placing a Gaussian component with isotropic covariance at every point in the set. The second set is fitted to it via MLE. The introduction of the outlier component in the GMM and the sphere-like isotropic covariance makes these methods robust. However, unlike point-to-plane ICP variants [7, 8, 35, 36] which take advantage of the local planar geometry in point cloud data, these methods ignore such information. While these methods are robust, the isotropic covariance does not faithfully capture the local planar geometry of the point clouds, leading to inaccuracy in registration.

In this paper, we propose a novel method, CPD with Local Surface Geometry (LSG-CPD), which takes into account local surface geometry: surface normal and surface variation. Surface variation [31] is closely related to curvature and surface flatness. Intuitively, the flatter the local surface,

*Corresponding author

the smaller the curvature, the smaller the surface variation. We construct a GMM on the target set. The source set is fitted to it via MLE (Fig. 1). However, unlike previous methods [14, 29] which restrict covariances to be isotropic, we adaptively impose different levels of point-to-plane penalization on top of the original point-to-point penalization, resulting in anisotropic covariances (Sec. 3.1). If the local surface is flat, *i.e.*, the surface variation is small, a large point-to-plane penalization is added. The GMM component will be plane-like and points will be mainly attracted towards the local plane instead of the centroid (Fig. 1(a)). If the local surface has large curvature or the normal estimation is noisy, *i.e.*, the surface variation is large, the point-to-plane penalization will be small (Fig. 1(b)). In this case, the GMM component will be sphere-like and robustness is maintained – the registration will not be misled by those ineffectively estimated surface normals. In addition, we propose a method to estimate the weight placed on the outlier component of the GMM (Sec. 3.2). Furthermore, we introduce the idea of confidence filtering (CF) which allows flexibility in reallocating weights of GMM components and outlier weights based on measurement confidence of points (Sec. 3.3). The EM algorithm [4, 12] is used to solve the MLE problem. To make our method efficient, we reformulate the computation in the E step into parallelizable matrix manipulations which can be efficiently computed on a GPU (Sec. 3.4). In the M step, we propose an unconstrained optimization method on a matrix Lie group and find the optimum with the efficient Newton’s method (Sec. 3.5).

Our motivation is to make the GMM-based description of the point cloud geometry richer and more faithful to the “true” probabilistic distribution by incorporating local surface geometry, afforded by the surface normal and variation, with the original point position. The optimum of the objective function in the EM process will be closer to the “true” transformation compared to that using the isotropic covariances. Thus, our method is able to improve the registration accuracy while maintaining the robustness of probabilistic methods. Results show that the proposed method outperforms state-of-the-art algorithms in terms of accuracy and robustness on a wide variety of dataset captured by range scanners, RGBD cameras, and LiDARs.

2. Related Works

Point cloud registration has been extensively studied in the literature. In this review, we limit the scope to GMM-based probabilistic methods. One popular formulation of GMM-based methods builds a GMM distribution on one point set and softly assigns correspondence between each pair of points in the two sets probabilistically [13, 14, 16, 20, 24, 25, 27, 29, 32, 33]. Our method falls into this category. The most related works are CPD [29] and FilterReg [14]. CPD [29] derives the EM procedure and proposes a close-form solution in the M step. Recently, FilterReg [14] introduces Gaussian filter and twist parameterization to improve the efficiency of the CPD framework. Point-to-point FilterReg [14]

is similar to CPD [29]. However, both methods ignore local surface geometry and place an isotropic covariance around every point in the first point cloud to construct the GMM. Point-to-plane FilterReg [14] tries to account for the surface normal by substituting the point-to-point penalization with point-to-plane penalization in the M step. This leads to various problems including rank deficiency of the covariance and inconsistency between the probabilistic model in the E and M step. A more technical analysis can be found in Sec. 3.1. ECMPR [20] and [26] update the whole covariance matrix in the M step. Although it allows general anisotropic covariances, no local geometric information is considered. GMM-Tree [13] partitions points into clusters and builds a hierarchy of Gaussian mixtures for the clusters. The covariances, estimated directly by all the points in the cluster instead of local surface geometry, are anisotropic. But it does not account for outliers in its probabilistic modeling, making it not robust to outliers and missing points. In [27, 33], hybrid mixture models (*e.g.*, Gaussian + von-Mises-Fisher) are used to model position and orientation uncertainty. The complicated mixture models introduce heavy computation. In contrast, our method incorporates surface normals in a simple GMM adaptively, making it robust to ineffective surface normal estimation.

Another popular line of GMM-based methods build GMM probability distribution on both the source and the target sets [6, 22, 37, 38]. Point cloud registration is formulated as a problem of minimizing the difference between the two distributions. [22] uses KL-divergence to parameterize the distribution difference while [38] uses squared L2 norm. To improve the efficiency, voxelization [37] and Support Vector Machine [6] are used to represent the probability representation in a compact and efficient way.

3. Methods

3.1. Formulation of Probabilistic Model

The alignment of two point clouds is formulated as an MLE problem. We fix one point cloud (target set) and align the other one (source set) to it. Unlike CPD [29] which places the GMM on the source set, the GMM in our method is built on the target set, which is similar to FilterReg [14] (Fig. 1). Rigid transformation is applied on the source set and the transformed source set is regarded as the observation of the GMM. The goal is to find the optimal transformation which makes the observation best fit the GMM. Through out the paper, we adopt the following notation:

- N - number of points in the source set
- M - number of points in the target set
- $\mathbf{X} \in \mathbb{R}_{N \times 3}$ - source set where the n -th row \mathbf{x}_n^T is the position of the n -th point in the set
- $\mathbf{Y} \in \mathbb{R}_{M \times 3}$ - target set where the m -th row \mathbf{y}_m^T is the position of the m -th point in the set
- $\mathbf{N} \in \mathbb{R}_{M \times 3}$ - surface normal matrix where the m -th row \mathbf{n}_m^T is the unit normal vector of the corresponding m -th point in \mathbf{Y} ,
- \mathbf{I} - identity matrix

Following [13, 14, 29], the probability density function of the GMM is given as:

$$p(\mathbf{x}_n) = w_n p_o(\mathbf{x}_n) + (1 - w_n) \sum_{m=1}^M \pi(m) p(\mathbf{x}_n|m) \quad (1)$$

where $\pi(m)$ is the prior probability of \mathbf{x}_n being assigned to the m -th component; $p(\mathbf{x}_n|m)$ is the corresponding component density function. Besides the M Gaussian components centered at each point in the target set, an outlier component is added to the GMM to make the algorithm robust to outliers and missing points. $w_n \in [0, 1)$ is the outlier weight of \mathbf{x}_n and p_o is the outlier probability density.

In [14, 29], the outlier probability density $p_o(\mathbf{x}_n) = \frac{1}{N}$. As pointed out in [18], it does not satisfy a property of probability density functions: the integral of $p_o(\mathbf{x}_n)$ over \mathbb{R}^3 equals 1. In this paper, we follow [18] and define a working space $\mathbb{V} \subset \mathbb{R}^3$ whose volume equals V . \mathbb{V} represents a solid body which encapsulates the target set. $p_o(\mathbf{x}_n) = 1/V$ if $\mathbf{x}_n \in \mathbb{V}$ and $p_o(\mathbf{x}_n) = 0$ otherwise. In [14, 29], $w_n = w_0$ for $n = 1, 2, \dots, N$; $\pi(m) = \frac{1}{M}$ for $m = 1, 2, \dots, M$. Both w_n and $\pi(m)$ are uniform across all points in the source and target set, respectively. In Sec. 3.2, we introduce a method to estimate w_0 from the outlier ratio of the point cloud. In Sec. 3.3, we propose a method which allows flexibility to assign different w_n and $\pi(m)$ to different \mathbf{x}_n and \mathbf{y}_m respectively according to their measurement confidence.

The key novelty of our method lies in our design of $p(\mathbf{x}_n|m)$ which incorporates both point-to-point and point-to-plane penalization into the probabilistic modeling. Surface normals are easily obtainable features of point cloud data [19, 31] and have been successfully utilized in many ICP variants [3, 7, 35, 36] to improve the registration performance. Point-to-plane extension of FilterReg [14] tries to extend the CPD [29] framework by substituting the point-to-point distance in the original objective function with the point-to-plane distance in the M step without further modifying the E step. This will lead to rank deficiency of the covariances, thus resulting in an ill-posed optimization in the M step. More importantly, the underlying probabilistic model in the E and M step are inconsistent. Neither convergence nor optimality is guaranteed for the EM process [40]. In this paper, we propose a component density function as follows:

$$p(\mathbf{x}_n|m) = c_m \exp\left(-\frac{1}{2}(\mathbf{x}_n - \mathbf{y}_m)^T \Sigma_m^{-1} (\mathbf{x}_n - \mathbf{y}_m)\right) \Sigma_m^{-1} = \frac{1}{\sigma^2} (\alpha_m \mathbf{n}_m \mathbf{n}_m^T + \mathbf{I}) \quad (2)$$

c_m is a normalizing constant. Σ_m is a full-rank anisotropic covariance matrix of which the inverse encodes information of the local surface geometry. σ^2 is the covariance multiplier. It is initialized autonomously similar to CPD [29]. The identity matrix \mathbf{I} in Σ_m^{-1} penalizes the point-to-point distance between \mathbf{x}_n and \mathbf{y}_m similar to [14, 29]. The rank one projection matrix $\mathbf{n}_m \mathbf{n}_m^T$ penalizes the point-to-plane distance between \mathbf{x}_n and the local plane defined by \mathbf{n}_m at \mathbf{y}_m . These two matrices mutually benefit each other. $\mathbf{n}_m \mathbf{n}_m^T$ enhances the

accuracy provided by the point-to-point penalization. Meanwhile, \mathbf{I} serves as a regularizer, pulling the problem away from ill-posedness. The penalization coefficient α_m is set adaptively to add different levels of point-to-plane penalization depending on how flat the local surface is. This is realized by evaluating the surface variation κ_m [31] within the neighborhood of \mathbf{y}_m . Surface variation $\kappa \in [0, \frac{1}{3}]$ [31] is the ratio of the smallest eigenvalue to the trace of the covariance matrix of a local neighborhood (note that this covariance matrix is not Σ_m). It is computed via eigenanalysis. It describes the degree of deviation of points from the tangent plane. Intuitively, the flatter the local surface, the smaller the curvature, the smaller the surface variation. α_m is chosen as a modified logistic sigmoid function:

$$\alpha_m = \frac{1 - \exp(\lambda(3 - \frac{1}{\kappa_m}))}{1 + \exp(\lambda(3 - \frac{1}{\kappa_m}))} \alpha_{\max} \quad (3)$$

where α_{\max} is the upper bound; λ is a coefficient governing the sensitivity of α_m to $\frac{1}{\kappa_m}$. $\alpha_m \rightarrow 0$ when $\kappa_m \rightarrow \frac{1}{3}$ (large surface variation, non-flat local surface geometry); $\alpha_m \rightarrow \alpha_{\max}$ when $\kappa_m \rightarrow 0$ (small surface variation, flat local surface geometry). By making α_m adaptive to local surface geometry, our algorithm obtains several desirable properties: (1) large point-to-plane penalization is only applied to points around which the local surface can be well approximated by a plane; (2) little directional preference is imposed when the curvature is large or effective normal estimation is unavailable, *e.g.*, corrupted by noise.

We denote $g(\mathbf{x}_n)$ as the image of the point \mathbf{x}_n transformed by the rigid transformation g . Given the GMM and the observation, the negative log-likelihood function is:

$$L(g, \sigma^2) = - \sum_{n=1}^N \log(p(g(\mathbf{x}_n))) \quad (4)$$

Our goal is to find the optimal rigid transformation which minimizes L . However, directly solving this optimization problem is intractable. Thus, the EM algorithm [4, 12] is applied to find the solution. The correspondence between the points in \mathbf{X} and \mathbf{Y} is treated as a latent variable.

E step: We use the transformation g_{old} from the previous iteration to compute the posterior correspondence probability between each pair of point in \mathbf{X} and \mathbf{Y} via Bayes' rule. The correspondence probability is stored in a matrix $\mathbf{P} \in \mathbb{R}^{M \times N}$ where the (m, n) -th entry \mathbf{P}_{mn} represents the posterior probability of the point \mathbf{x}_n being assigned to the m -th component:

$$\mathbf{P}_{mn} = \frac{(1 - w_n) \pi(m) p(g_{old}(\mathbf{x}_n)|m)}{w_n p_o + (1 - w_n) \sum_m \pi(m) p(g_{old}(\mathbf{x}_n)|m)} \quad (5)$$

The computation of \mathbf{P} can be reformulated into efficient and parallelizable matrix manipulations (Sec. 3.4).

M step: The rigid transformation g and covariance multiplier σ^2 are inferred by minimizing the following [4, 29] (see Supplementary Material for details):

$$Q(g, \sigma^2) = - \sum_{n=1}^N \sum_{m=1}^M \mathbf{P}_{mn} \log(\pi(m) p(g(\mathbf{x}_n)|m)) \quad (6)$$

In the case where the covariance matrix Σ_m is isotropic, a closed form solution of Eq. (6) can be found [29]. However, this is not the case when Σ_m is anisotropic [20]. Thus, we propose an unconstrained optimization method on a matrix Lie group to solve the optimization in the M step in Sec. 3.5.

3.2. Outlier Weight Estimation

In [14, 29], the uniform outlier weight $w_0 \in [0, 1)$ is manually tuned by the user. However, w_0 is merely a hyperparameter which is different from the outlier ratio η - the ratio of outliers to the number of points in the source set. This makes it hard to choose w_0 in different tasks intuitively. Here, we introduce a method to determine w_0 from η using the fact that w_0 and η are probabilistically related. In this way, the user only needs to set the outlier ratio η , which is more intuitive than setting w_0 . When the point clouds are well aligned by the “true” transformation g^* , the expectation of the outlier number $N_{outlier}$ should equal the true number of outliers ηN :

$$E(N_{outlier}) = \sum_{n=1}^N P(M+1|g^*(\mathbf{x}_n)) = \eta N \quad (7)$$

where

$$P(M+1|g^*(\mathbf{x}_n)) = \frac{\frac{w_0}{1-w_0} p_o}{\frac{w_0}{1-w_0} p_o + \sum_m \pi(m) p(g^*(\mathbf{x}_n)|m)}$$

is the posterior probability of the point \mathbf{x}_n being an outlier. w_0 can then be solved from Eq. (7). However, g^* is not available *a priori*. Thus, from Eq. (7), we can only derive the upper bound of w_0 , *i.e.*, $w_0 \leq w_{\max} = \frac{\eta V \sum_m \pi(m) c_m}{(1-\eta) + \eta V \sum_m \pi(m) c_m} \in [0, 1)$, by setting $p(g^*(\mathbf{x}_n)|m)$ as its upper bound c_m (Eq. (2)). To make the algorithm robust, generally a large w_0 is desired. However, if $w_0 > w_{\max}$, probabilistically, more points are regarded as outliers than the truth. The inlier density has the risk of being submerged by the outlier density. And as $w_0 \rightarrow 1$, the density of the GMM will be primarily dominated by the outlier component. The algorithm tends to get stuck at a local minimum. Therefore, we set $w_0 = w_{\max}$ in practice to strike a balance.

3.3. Confidence Filtering

Most GMM-based methods [13, 14, 20, 22, 27, 29] weigh all the GMM components uniformly with $\pi(m) = \frac{1}{M}$. In addition, the outlier weight $w_n = w_0$ is uniform across all points in the source set. This setting is reasonable when the measurement error induced by the sensor is uniform within its range, *e.g.*, range scanners. And thus the measurement confidence of all the points is equal in general. However, point clouds acquired from consumer-grade RGBD cameras are subject to bias and noise which grows with the increment of depth [1, 17]. The larger the depth of the point is, the less confident (or more uncertain) the measurement is. The algorithm should pay more attention to the alignment of points measured with high confidence. In this paper, we define $\phi(\mathbf{x}) = e_{\min}/e(\mathbf{x})$ as the measurement confidence of a

point \mathbf{x} , where $e(\mathbf{x})$ is the error model which provides the statistical measurement error of a point measured at \mathbf{x} [17]; e_{\min} is the minimum measurement error within the sensor range. Error analysis and modeling of different sensors can be found in [10, 17, 21]. The confidence filtering (CF) modifies $\pi(m)$ and w_n as:

$$\pi(m) = \frac{\phi(\mathbf{y}_m)}{\sum_M \phi(\mathbf{y}_m)}, \quad w_n = 1 - (1 - w_0)\phi(\mathbf{x}_n) \quad (8)$$

The component prior probability $\pi(m) \in (0, 1)$ decreases as $\phi(\mathbf{y}_m)$ decreases. The reallocation of $\pi(m)$ reduces the attraction force of a component generated by a point with less confidence and vice versa. The outlier weight $w_n \in [w_0, 1]$ increases as $\phi(\mathbf{x}_n)$ decreases. w_0 is the baseline outlier weight determined in Sec. 3.2. As the measurement of \mathbf{x}_n becomes less confident, we would rather categorize it as an “outlier” to lower its contribution to the registration. Furthermore, points with confidence lower than a threshold can be truncated. This not only improves the accuracy but also reduces the runtime as shown in Sec. 4.3. Note that the uniformly weighing setting is a special case where $\phi(\mathbf{x}) = 1$ for any \mathbf{x} .

3.4. E Step: Efficient Computation of Correspondence Matrix

The introduction of the correspondence matrix \mathbf{P} provides GMM-based methods with an inherent advantage on robustness. However, computing \mathbf{P} is the most computationally expensive part of the algorithm: the correspondence probability of every pair of points in \mathbf{X} and \mathbf{Y} needs to be computed. We reformulate this process into simple matrix manipulations which can be efficiently computed on a GPU.

Plugging w_n and $\pi(m)$ defined in the previous sections and Eq. (2) into Eq. (5), the (m, n) -th entry of the correspondence matrix \mathbf{P} becomes:

$$\mathbf{P}_{mn} = \frac{\mathbf{K}_{mn}}{\sum_{m=1}^M \mathbf{K}_{mn} + \gamma} \quad (9)$$

where $\gamma = (2\pi\sigma^2)^{3/2}/V$ and \mathbf{K}_{mn} is the (m, n) -th entry of a matrix $\mathbf{K} \in \mathbb{R}_{M \times N}$:

$$\mathbf{K}_{mn} = \frac{1 - w_n}{w_n} \pi(m) \exp\left(-\frac{1}{2} \|g(\mathbf{x}_n) - \mathbf{y}_m\|_{\Sigma_m^{-1}}^2\right)$$

Exploiting the special formulation of Σ_m , the computation of \mathbf{K} can be recast into fundamental matrix manipulations:

$$\mathbf{K} = \mathbf{C} \odot \underline{\exp}\left(-\frac{1}{2\sigma^2}(\mathbf{D}\mathbf{A} + \mathbf{B})\right) \quad (10)$$

in which \odot is the element-wise product of matrices; $\underline{\exp}(\cdot)$ is the element-wise exponential; the (m, n) -th entry of \mathbf{A}

$$\mathbf{A}_{mn} = ((\mathbf{N}\mathbf{R}\mathbf{X}^T + \mathbf{N}\mathbf{t}\mathbf{1}^T - \mathbf{s}\mathbf{1}^T)_{mn})^2$$

$$\mathbf{B} = \mathbf{Q} + \mathbf{t}^T \mathbf{t}\mathbf{I} + 2(\mathbf{1}\mathbf{t}^T \mathbf{R}\mathbf{X}^T - \mathbf{Y}\mathbf{R}\mathbf{X}^T - \mathbf{Y}\mathbf{t}\mathbf{1}^T)$$

$$\mathbf{C} \in \mathbb{R}_{M \times N}, \quad \mathbf{C}_{mn} = \frac{1 - w_n}{w_n} \pi(m)$$

$$\mathbf{D} = \text{diag}(\alpha_1, \alpha_2, \dots, \alpha_m)$$

in which $\mathbf{1}$ is the column vector with all ones,

$$\mathbf{s} = [\mathbf{y}_1^T \mathbf{n}_1 \quad \mathbf{y}_2^T \mathbf{n}_2 \quad \dots \quad \mathbf{y}_M^T \mathbf{n}_M]^T$$

$$\mathbf{Q} \in \mathbb{R}_{M \times N}, \mathbf{Q}_{mn} = \|\mathbf{x}_n\|_2^2 + \|\mathbf{y}_m\|_2^2$$

are terms independent of transformation $g = (\mathbf{R}, \mathbf{t})$, and thus can be pre-computed.

3.5. M Step: Optimization On Matrix Lie Group

In the M step, our goal is to infer the optimal transformation and variance multiplier σ^2 by minimizing Q given by Eq. (6). We propose an unconstrained optimization method to solve this optimization problem. Instead of parameterizing the rigid transformation, our method optimizes on the matrix Lie group via Newton’s method. The gradient and the Hessian matrix of a function on matrix Lie group are defined with respect to the basis in Lie algebra [9]. We denote $\mathbf{E}_1, \mathbf{E}_2, \dots, \mathbf{E}_6$ as the basis matrices of $se(3)$ Lie algebra. The gradient $\nabla Q = [E_1^T Q \quad E_2^T Q \quad \dots \quad E_6^T Q]^T \in \mathbb{R}^6$. The Hessian $\mathbf{H} \in \mathbb{R}_{6 \times 6}$ is asymmetric in which the (i, j) -th entry $\mathbf{H}_{ij} = E_i^T E_j^T Q$. E_i^T is the right derivative operators with respect to the i -th basis matrix \mathbf{E}_i . Both the gradient and the Hessian can be derived in closed-form:

$$E_i^T Q \doteq \frac{d}{dt} Q(g \circ \exp(t\mathbf{E}_i))|_{t=0}$$

$$= 2 \sum_{m=1}^M \sum_{n=1}^N \mathbf{P}_{mn} (\tilde{\mathbf{g}}\tilde{\mathbf{x}}_n - \tilde{\mathbf{y}}_m)^T \tilde{\Sigma}_m^{-1} \tilde{\mathbf{g}} \mathbf{E}_i \tilde{\mathbf{x}}_n \quad (11)$$

$$E_i^T E_j^T Q = 2 \sum_{m=1}^M \sum_{n=1}^N \mathbf{P}_{mn} (\tilde{\mathbf{x}}_n^T \mathbf{E}_j^T \tilde{\mathbf{g}}^T \tilde{\Sigma}_m^{-1} \tilde{\mathbf{g}} \mathbf{E}_i \tilde{\mathbf{x}}_n$$

$$+ (\tilde{\mathbf{g}}\tilde{\mathbf{x}}_n - \tilde{\mathbf{y}}_m)^T \tilde{\Sigma}_m^{-1} \tilde{\mathbf{g}} \mathbf{E}_j \mathbf{E}_i \tilde{\mathbf{x}}_n) \quad (12)$$

\tilde{g} , $\tilde{\mathbf{x}}_n$, and $\tilde{\mathbf{y}}_m$ are the homogeneous form of g , \mathbf{x}_n , and \mathbf{y}_m , respectively. $\tilde{\Sigma}_m^{-1} = \Sigma_m^{-1} \oplus 0$ is the augmented inverse covariance matrix where \oplus denotes the direct sum. Detailed derivations can be found in the Supplementary Material. It is worth noting that this method can be easily extended to other types of registration of which the transformation can be represented by a matrix Lie group, *e.g.*, affine transformation, by replacing the basis matrices in Lie algebra. The transformation is updated by Newton’s method until convergence:

$$\tilde{g}_{i+1} = \tilde{g}_i \circ \exp\left(\frac{1}{2}(\mathbf{H} + \mathbf{H}^T)^{-1} \nabla Q\right)^\wedge \quad (13)$$

where \wedge is the operation which maps the vector to the corresponding matrix Lie algebra [28]. After g gets updated, σ^2 is inferred by setting the corresponding partial derivative of Eq. (6) to zero. Unlike ECMPR [20] in which the whole covariance matrix is updated, only the covariance multiplier σ^2 gets updated in our approach. As a result, while the volume of the covariance Σ_m is updated in every iteration, the local geometric structure is preserved.

4. Results

We conduct a wide variety of experiments on public 3D point cloud datasets, including both object and environment dataset, to show the robustness and accuracy of our method. We implement our method on Matlab. We evaluate on a computer running Intel Core i9-9900K with an NVIDIA GeForce RTX 2080Ti GPU. In all the experiments, the code of the baseline methods are either provided directly by the authors or taken from the popular open source library with various performance optimization. The parameters for all the baseline methods are either provided by the authors or the softwares or well-tuned by ourselves if not provided. CF is only used in the experiment on the RGBD dataset (Sec. 4.3). The error model $e(\mathbf{x})$ of the sensor (Asus Xtion Pro), which is a quadratic function of the depth of \mathbf{x} , is sourced from [17]. More details on the experiment settings can be found in the Supplementary Material.

4.1. Registration on Data with Outliers and Noise

We follow [14, 29] and test on laser range data by adding outliers and noise, respectively. For both experiments, we use the Bunny dataset in Stanford 3D scanning repository [11]. We downsample the original point cloud to about 3500 points and rotate it about a random axis by 50 degrees. We compare with 4 baseline method: CPD [29], FilterReg [14], ECMPR [20], and TrICP [8]. CPD and FilterReg represent GMM-based algorithm with isotropic covariances; ECMPR represents GMM-based algorithm with anisotropic covariances; TrICP, a widely used robust ICP variant, represents the ICP methods. Following [14], we use the corresponding point-to-point distance to measure the registration error:

$$\text{error}(g) = \frac{1}{N} \sum_{i=1}^N \|g(\mathbf{x}_i) - g_{gt}(\mathbf{x}_i)\|_2 \quad (14)$$

where \mathbf{x}_i is the i -th point in the source set; g and g_{gt} are the transformations of the registration and the ground truth, respectively. All methods are exhaustively run with 100 iterations to ensure convergence. All the results are statistical: errors are averaged values of 30 independent runs.

In the outlier experiment, we add different ratios (the number of outliers to the number of original points) of Gaussian outliers. Results are shown in Fig. 2. All the GMM-based methods significantly outperform TrICP when there are outliers present, where CPD, FilterReg, and the proposed method are the best. However, with the increase of the outlier ratio, the error of CPD and FilterReg increases while the performance of the proposed method maintains and outperforms CPD and FilterReg (Fig. 2(b)).

In the noise experiment, we corrupt each point in the source and target set by adding Gaussian noise with different std. Results are shown in Fig. 3. All the GMM-based methods significantly outperform TrICP when there is noise. CPD, FilterReg, and the proposed methods have similar performance among which the proposed method has a small advantage (Fig. 3(b)). We also show the averaged penalization

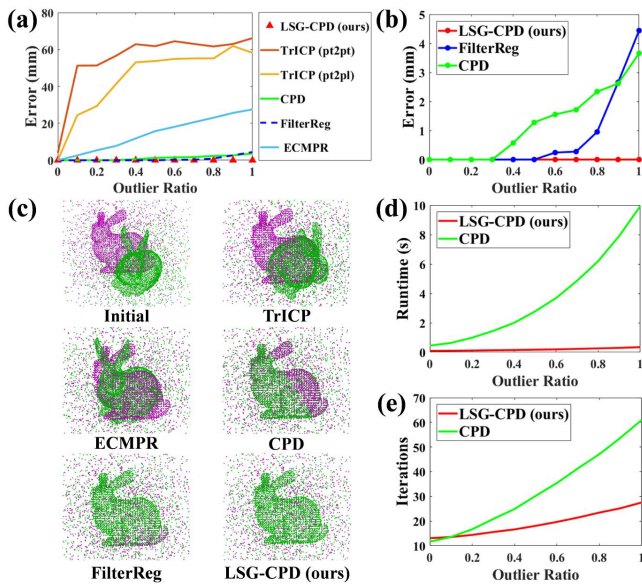


Figure 2. Outlier results. (a) Errors (Eq. (14)) of different methods. (b) Enlarged image showing errors of CPD, FilterReg, and LSG-CPD. (c) Initialization and results of different methods on data with 0.5 outlier ratio. TrICP point-to-plane result is shown as the result for TrICP. The images are cropped for better visualization. (d) Average runtime of CPD and LSG-CPD on data with different outlier ratios. (e) Averaged iterations of CPD and LSG-CPD on data with different outlier ratios.

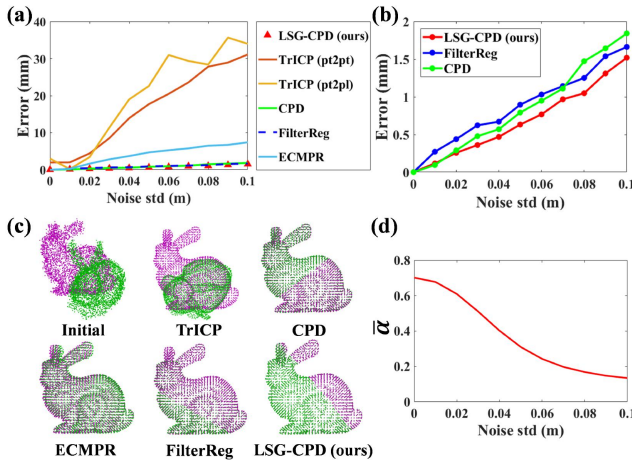


Figure 3. Noise results. (a) Errors (Eq. (14)) of different methods. (b) Enlarged image showing errors of CPD, FilterReg, and LSG-CPD. (c) Initialization and results of different methods on data with 0.03m noise level. TrICP point-to-plane result is shown as the result for TrICP. Denoised point clouds are used for better visualization of registration result. (d) Averaged penalization coefficient $\bar{\alpha}$ (Eq. (2)) of all points in the target set with respect to the noise level.

coefficient $\bar{\alpha}$ at different noise levels (Fig. 3(d)). As the noise level increases, $\bar{\alpha}$ decreases. This is because without noise, the local surface will look like a plane and thus α will be relatively large; when the points are corrupted, the local plane will also be corrupted and normal estimation becomes less effective, resulting in small α as we desire.

We compare the speed with CPD on data with different outlier ratios. In this experiment, we terminate the algorithm when the rotation error is less than 8° and the translation error

Method	Dragon	Happy	Armadillo
Initial	5.12	5.41	5.23
TrICP	1.87	1.16	0.66
FilterReg	1.13	2.13	0.63
MATrICP	1.18	1.29	0.53
EMPMR	1.39	1.07	0.50
LSG-CPD (ours)	0.97	1.08	0.43

Table 1. Object Modeling Errors (Eq. (14), unit mm).

is less than 1cm. Fig. 2(d) and (e) summarize the result. Our method is 5 times faster than CPD (89ms vs 450ms) when there are no outliers and 29 times faster (346ms vs 9940ms) when the outlier ratio is 1. Generally, our method takes less iterations to converge. And the time for each iteration is less than that of CPD. Averaging across all the experiment, our method is 20 times faster than CPD.

For FilterReg, we use the point-to-point version. We tested the point-to-plane version extensively but find it doesn't give a good performance. We think the reasons could be the singular covariance and the inconsistency of the probabilistic model in the E and M step, as mentioned in Sec. 3.1, making it not robust to bad initialization and/or outliers and noise. In the following experiments, if not mentioned, FilterReg is tested with the point-to-point version.

4.2. Object Modeling on Range Datasets

In this experiment, we test on object modeling from laser range data of objects captured from different viewpoints. We evaluate on three datasets from the Stanford 3D scanning repository [11], *i.e.*, Dragon Stand, Happy Stand, and Armadillo. Each consisting of 12 or 15 sequential scans of an object with evenly spaced view angles. We compare our method against two pairwise methods, *i.e.*, FilterReg [14] and TrICP [8], and two multi-view methods, *i.e.*, MATrICP [23] and EMPMR [43]. Object modeling for pairwise methods makes use of the spatial information provided by the scanning sequence. Point clouds that are spatially adjacent to each other are first pairwise registered. Then, all the point clouds are aligned to the first one for whole object modeling. Following [43], we perturb the point clouds from its ground truth pose with a random pose (± 0.05 rad in rotation and ± 0.01 m in translation). We downsample the point cloud of each scan with voxelized grids, resulting in about 3500 points for each view. For each dataset, we carry out 20 independent experiments. We use the averaged corresponding point-to-point distance (Eq. (14)) as the evaluation metric.

Qualitative results are shown in Table 1. Fig. 4 shows the cross-section results of different methods. The proposed method outperforms the two pairwise baseline methods. The chain-like reconstruction method used for pairwise methods has a disadvantage of error accumulation. The error in each pairwise registration will propagate and accumulate along the chain when aligned to the base point cloud. On the other hand, multi-view methods can simultaneously process all the point clouds and take advantage of techniques such as motion averaging [23]. Even with such a significant disadvan-

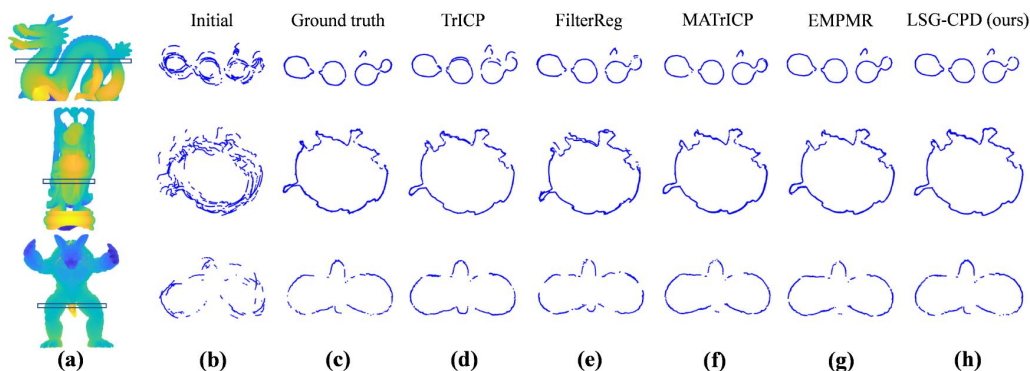


Figure 4. Multiview cross-section results where the corresponding regions are indicated in (a). Thinner curves indicate better registration. (a) shows the whole object modeling results of the proposed method; (b) and (c) show the initial and ground truth; the rest show the cross-section results of different methods.

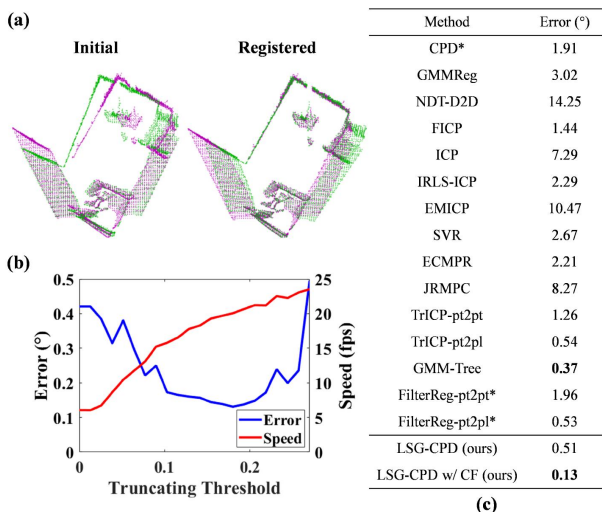


Figure 5. Stanford Lounge Result. (a) Examples of registration result. (b) Angle error and speed versus truncating threshold. (c) Registration results of various methods. The results of the baseline methods with a star (*) are sourced from [14]; the results of other baseline methods are sourced from [13].

tage, compared with state-of-the-art multi-view methods, our method outperforms them on two datasets and is on par with them on the third.

4.3. Registration on RGBD Dataset

In this section, we follow the setting in [13] and test on the Stanford Lounge dataset [42] which contains range data of an indoor scene captured with an RGBD camera. We register every 5-th frame for the first 400 frames. Each frame is downsampled to about 5000 points. Following [13], we use the average Euler angle deviation from the ground truth to measure the rotation error between each pair.

The results are shown in Fig. 5. Our method achieves the state-of-the-art without CF. With CF, the accuracy improves greatly and outperforms the state-of-the-art. The angle error is more than 2 times lower than the state-of-the-art. Additionally, CF greatly accelerates the registration by truncating points with lower accuracy confidence. The registration frame rate increases from 5 FPS to 19 FPS. We also experiment on truncating points at different confidence

thresholds (Fig. 5(b)). The registration error decreases from 0.51° to 0.42° by just adding CF without any truncation. As expected, the registration speed increases as the truncating threshold increases. However, if the threshold is too large, too many points, including the accurately measured points, are truncated. The point cloud will contain insufficient information for registration. If the threshold is too small, *i.e.*, too few points are truncated, the point cloud contains many noisy data.

4.4. Registration on LiDAR Dataset

In this experiment, we evaluate on the Kitti dataset [15]. We test on the 07 sequence of the odometry benchmark, which contains a drive for about 700 meters. We down-sample the point clouds to about 5000 points with voxelized grids. We register together every other frame pairwise for the whole sequence, resulting in 550 pairs of point clouds for registration. All the frames are aligned to the first frame to reconstruct the whole sequence. We evaluate the averaged relative and absolute rotation and translation errors. We also compute the rotation and translation errors of the last frame. The translation error is the Euclidean distance between the registration result and the ground truth. The rotation error $\Delta\theta$ is the angle of the axis-angle parameterization of the relative rotation between the ground truth and the registration result. The results are shown in Fig. 6 and Table 2. Our method outperforms all baseline methods on every evaluation metric. Remarkably, even without loop closure, our proposed method is able to achieve very small translation and rotation errors in the last frame of the whole sequence. All the ICP-based methods except GICP [35] have relatively large errors. For FilterReg [14], we use the point-to-plane version as the point-to-point version does not give a good result. FilterReg [14] and the proposed method outperform CPD [29].

5. Discussion & Future Works

Extensive experiments on various kinds of data show that our method is accurate and robust. As shown in Sec. 4.1, ICP-based methods require good initialization and are not robust to outliers and noise. CPD and point-to-point FilterReg, although robust to outliers, noise, and missing points,

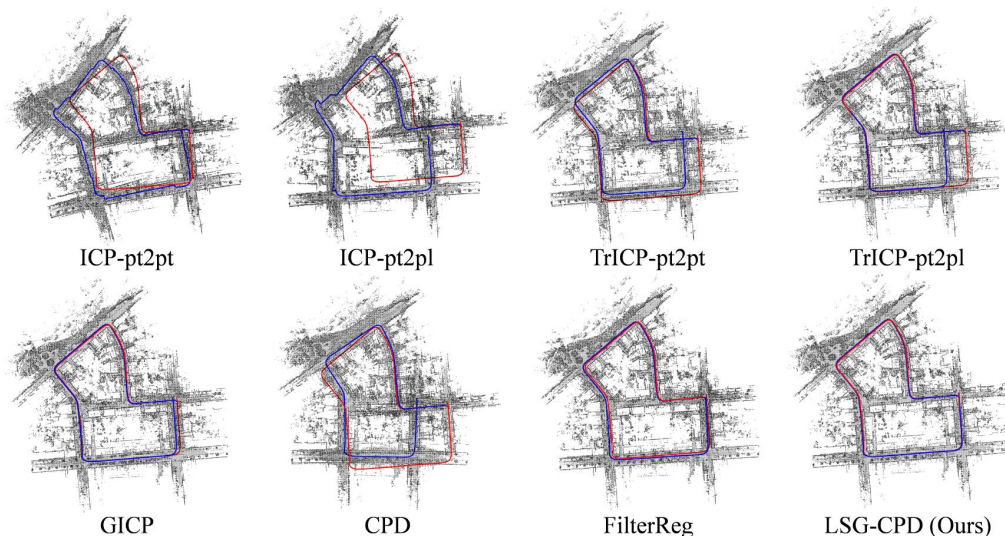


Figure 6. Kitti Odometry Benchmark 07 Sequence Result Visualization. Red indicates the ground truth; blue indicates the registration result.

Method	Rel. Rot. ($^{\circ}$)	Rel. Tran. (m)	Abs. Rot. ($^{\circ}$)	Abs. Tran. (m)	Last Rot. ($^{\circ}$)	Last Tran. (m)
ICP-pt2pt	0.283 (3.690)	0.214 (5.966)	9.476 (17.641)	17.296 (36.379)	7.100	16.131
ICP-pt2pl	0.475 (8.683)	0.187 (14.320)	7.986 (16.496)	48.422 (70.712)	6.526	58.902
TrICP-pt2pt	0.123 (0.689)	0.117 (2.324)	1.606 (2.476)	10.446 (24.881)	1.587	24.449
TrICP-pt2pl	0.071 (0.607)	0.080 (2.604)	1.365 (3.150)	8.914 (33.113)	3.073	33.094
GICP	0.070 (0.787)	0.035 (2.108)	1.181 (1.799)	3.052 (6.939)	1.772	6.593
CPD	0.167 (1.100)	0.190 (2.701)	5.091 (7.900)	21.696 (56.690)	6.353	47.359
FilterReg	0.091 (0.615)	0.029 (0.099)	1.995 (7.016)	3.812 (5.666)	6.984	5.373
LSG-CPD (ours)	0.062 (0.591)	0.022 (0.090)	0.812 (1.419)	1.069 (1.830)	1.377	0.382

Table 2. Kitti Odometry Benchmark 07 Sequence Result. We show the relative rotation error (Rel. Rot.), relative translation error (Rel. Tran.), absolute rotation error (Abs. Rot.), absolute translation error (Abs. Tran.), last frame rotation error (Last Rot.), and last frame translation error (Last Tran.) of different methods. In the first four columns of result, the value outside the parenthesis is the averaged error of the 550 pairs; the value inside is the maximum.

do not perform well on the Kitti dataset. We hypothesize this is because LiDAR data contain many large planes. The isotropic covariance in CPD and point-to-point FilterReg are not able to accurately represent the geometry of the data, thus leading to bad performance. Point-to-plane FilterReg performs well on the Stanford Lounge and Kitti dataset in which the relative pose between adjacent frames for registration are close to each other. However, it is sensitive to bad initialization and are not robust to outliers and noise as discussed in Sec. 3.1 and 4.1. Our method is robust to outliers, noise, and missing points and accurate. The reason for the robustness is twofold. The probabilistic modeling which includes the outlier component make it robust to outliers and missing points like other GMM-based methods [14, 29]. The covariance adaptively adjusts its shape according to the local surface geometry, making it insusceptible to ineffectively estimated surface normals. The reason for the accuracy is also credit to the local surface geometry-based covariance. By using the surface variation [31] to adjust the point-to-plane penalization in the covariance, our method can accurately represent the geometry of a wide variety of data, including data of different scales, with/without large planes. Future works include extending the current framework to non-rigid

registration and applying the method to more computer vision tasks, *e.g.*, pose estimation and SLAM.

6. Conclusions

We present a novel GMM-based probabilistic registration method which outperforms state-of-the-art algorithms in terms of accuracy and robustness on various datasets. Our method makes use of the surface normal and surface variation [31] to create anisotropic covariances in the GMM which faithfully represents the geometry of the point cloud. We show that the computation can be reformulated into simple matrix manipulations which can be efficiently computed on a GPU. In addition, we propose an efficient optimization method to solve for the optimal transformation on a matrix Lie group in the EM process. Thanks to the efficient computation and optimization, the proposed method is about 20 times faster than modern implementations of CPD [29].

Acknowledgements This work was supported by NUS Startup grants R-265-000-665-133, R-265-000-665-731, Faculty Board account C-265-000-071-001, MOE grant R-265-000-655-114, and JHU internal funds.

References

- [1] F. Basso, E. Menegatti, and A. Pretto. Robust intrinsic and extrinsic calibration of rgb-d cameras. *IEEE Transactions on Robotics*, 34(5):1315–1332, 2018.
- [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] S. Billings and R. Taylor. Iterative most likely oriented point registration. In Polina Golland, Nobuhiko Hata, Christian Barillot, Joachim Hornegger, and Robert Howe, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2014*, pages 178–185, Cham, 2014. Springer International Publishing.
- [4] C. M. Bishop. *Pattern recognition and machine learning*. springer, 2006.
- [5] Jesus Briales and Javier Gonzalez-Jimenez. Convex global 3d registration with lagrangian duality. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [6] D. Campbell and L. Petersson. An adaptive data representation for robust point-set registration and merging. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4292–4300, 2015.
- [7] Y. Chen and G. Medioni. Object modelling by registration of multiple range images. *Image and vision computing*, 10(3):145–155, 1992.
- [8] D. Chetverikov, D. Svirko, D. Stepanov, and P. Krsek. The trimmed iterative closest point algorithm. In *Pattern Recognition, International Conference on*, volume 3, pages 30545–30545, 2002.
- [9] G. S. Chirikjian. *Stochastic Models, Information Theory, and Lie Groups, Volume 2*. 01 2012.
- [10] J. A. Christian and S. Cryan. A survey of lidar technology and its use in spacecraft relative navigation. In *AIAA Guidance, Navigation, and Control (GNC) Conference*, page 4641, 2013.
- [11] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312, 1996.
- [12] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B*, 39:1–38, 1977.
- [13] B. Eckart, K. Kim, and J. Kautz. Hgmr: Hierarchical gaussian mixtures for adaptive 3d registration. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 705–721, 2018.
- [14] W. Gao and R. Tedrake. Filterreg: Robust and efficient probabilistic point-set registration using gaussian filter and twist parameterization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11095–11104, 2019.
- [15] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361. IEEE, 2012.
- [16] S. Granger and X. Pennec. Multi-scale em-icp: A fast and robust approach for surface registration. In *European Conference on Computer Vision*, pages 418–432. Springer, 2002.
- [17] G. Halmetschlager-Funek, M. Suchi, M. Kampel, and M. Vincze. An empirical evaluation of ten depth cameras: Bias, precision, lateral noise, different lighting conditions and materials, and multiple sensor setups in indoor environments. *IEEE Robotics Automation Magazine*, 26(1):67–77, 2019.
- [18] O. Hirose. A bayesian formulation of coherent point drift. *IEEE transactions on pattern analysis and machine intelligence*, 2020.
- [19] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. In *Proceedings of the 19th annual conference on computer graphics and interactive techniques*, pages 71–78, 1992.
- [20] R. Horaud, F. Forbes, M. Yguel, G. Dewaele, and J. Zhang. Rigid and articulated point registration with expectation conditional maximization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(3):587–602, 2010.
- [21] R. Horaud, M. Hansard, G. Evangelidis, and C. M enier. An overview of depth cameras and range scanners based on time-of-flight technologies. *Machine vision and applications*, 27(7):1005–1020, 2016.
- [22] B. Jian and B. C. Vemuri. Robust point set registration using gaussian mixture models. *IEEE transactions on pattern analysis and machine intelligence*, 33(8):1633–1645, 2010.
- [23] Z. Li, J. Zhu, K. Lan, C. Li, and C. Fang. Improved techniques for multi-view registration with motion averaging. In *2014 2nd International Conference on 3D Vision*, volume 1, pages 713–719. IEEE, 2014.
- [24] B. Luo and E. R. Hancock. A unified framework for alignment and correspondence. *Computer Vision and Image Understanding*, 92(1):26–55, 2003.
- [25] G. McNeill and S. Vijayakumar. A probabilistic approach to robust shape matching. In *2006 International Conference on Image Processing*, pages 937–940. IEEE, 2006.
- [26] Z. Min, J. Pan, A. Zhang, and M. Q.-H. Meng. Robust non-rigid point set registration algorithm considering anisotropic uncertainties based on coherent point drift. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2019.
- [27] Z. Min, J. Wang, and M. Q.-H. Meng. Robust generalized point cloud registration with orientational data based on expectation maximization. *IEEE Transactions on Automation Science and Engineering*, 17(1):207–221, 2019.
- [28] R. M. Murray, Z. Li, and S. S. Sastry. *A mathematical introduction to robotic manipulation*. CRC press, 1994.
- [29] A. Myronenko and X. Song. Point set registration: Coherent point drift. *IEEE transactions on pattern analysis and machine intelligence*, 32(12):2262–2275, 2010.
- [30] A. N uchter, K. Lingemann, J. Hertzberg, and H. Surmann. 6d slam—3d mapping outdoor environments. *Journal of Field Robotics*, 24(8-9):699–722, 2007.
- [31] M. Pauly, M. Gross, and L. P. Kobbelt. Efficient simplification of point-sampled surfaces. In *IEEE Visualization, 2002. VIS 2002.*, pages 163–170, 2002.
- [32] A. Rangarajan, H. Chui, E. Mjolsness, S. Pappu, L. Davachi, P. Goldman-Rakic, and J. Duncan. A robust point-matching algorithm for autoradiograph alignment. *Medical image analysis*, 1(4):379–398, 1997.
- [33] N. Ravikumar, A. Gooya, A. F. Frangi, and Z. A. Taylor. Generalised coherent point drift for group-wise registration of multi-dimensional point sets. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 309–316. Springer, 2017.
- [34] T. Schmidt, R. A. Newcombe, and D. Fox. Dart: Dense articulated real-time tracking. In *Proceedings of Robotics: Science and Systems*, 2014.

- [35] A. Segal, D. Haehnel, and S. Thrun. Generalized-icp. In *Proceedings of Robotics: Science and Systems*, Seattle, USA, June 2009.
- [36] J. Serafin and G. Grisetti. Nicp: Dense normal based point cloud registration. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 742–749, 2015.
- [37] T. Stoyanov, M. Magnusson, H. Andreasson, and A. J. Lilienthal. Fast and accurate scan registration through minimization of the distance between compact 3d ndt representations. *The International Journal of Robotics Research*, 31(12):1377–1393, 2012.
- [38] W. Tabib, C. O’Meadhra, and N. Michael. On-manifold gmm registration. *IEEE Robotics and Automation Letters*, 3(4):3805–3812, 2018.
- [39] T. Whelan, S. Leutenegger, R. Salas-Moreno, B. Glocker, and A. Davison. Elasticfusion: Dense slam without a pose graph. In *Proceedings of Robotics: Science and Systems*, 2015.
- [40] C. F. Wu. On the convergence properties of the em algorithm. *Ann. Statist.*, 11(1):95–103, 03 1983.
- [41] H. Yang, J. Shi, and L. Carlone. Teaser: Fast and certifiable point cloud registration. *IEEE Transactions on Robotics*, 37(2):314–333, 2021.
- [42] Q.-Y. Zhou and V. Koltun. Dense scene reconstruction with points of interest. *ACM Transactions on Graphics (ToG)*, 32(4):1–8, 2013.
- [43] J. Zhu, R. Guo, Z. Li, J. Zhang, and S. Pang. Registration of multi-view point sets under the perspective of expectation-maximization. *IEEE Transactions on Image Processing*, 29:9176–9189, 2020.