

Stein ICP for Uncertainty Estimation in Point Cloud Matching

Fahira Afzal Maken^{1,*}, Fabio Ramos^{1,2} and Lionel Ott³

Abstract—Quantification of uncertainty in point cloud matching is critical in many tasks such as pose estimation, sensor fusion, and grasping. Iterative closest point (ICP) is a commonly used pose estimation algorithm which provides a point estimate of the transformation between two point clouds. There are many sources of uncertainty in this process that may arise due to sensor noise, ambiguous environment, initial condition, and occlusion. However, for safety critical problems such as autonomous driving, a point estimate of the pose transformation is not sufficient as it does not provide information about the multiple solutions. Current probabilistic ICP methods usually do not capture all sources of uncertainty and may provide unreliable transformation estimates which can have a detrimental effect in state estimation or decision making tasks that use this information. In this work we propose a new algorithm to align two point clouds that can precisely estimate the uncertainty of ICP’s transformation parameters. We develop a Stein variational inference framework with gradient based optimization of ICP’s cost function. The method provides a non-parametric estimate of the transformation, can model complex multi-modal distributions, and can be effectively parallelized on a GPU. Experiments using 3D kinect data as well as sparse indoor/outdoor LiDAR data show that our method is capable of efficiently producing accurate pose uncertainty estimates.

Index Terms—Perception for Grasping and Manipulation, Localization, Range sensing, Probabilistic Inference

I. INTRODUCTION

POINT cloud registration plays a fundamental role in many robotics and computer vision tasks such as localization and mapping [19], autonomous navigation [50], pose estimation [41], surgical guidance [33], and augmented reality [38] to name a few. Iterative closest point (ICP) [8] is the gold standard registration algorithm which estimates a relative transformation between two point clouds. Given an initial estimate, ICP minimizes the Euclidean distance between pairs of matching points from both point clouds in an iterative manner.

The ICP pose alignment process is adversely affected by different sources of error and uncertainty. These include initial pose uncertainty, sensor noise, partial overlap, multiple local minima of the cost function, and under-constrained or ill-posed

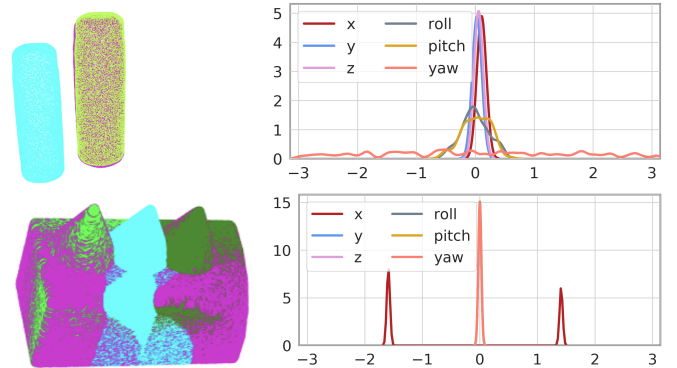


Fig. 1: Our method estimates ICP pose distributions (right) for the solution (green) when a source cloud (cyan) is aligned to a reference cloud (magenta). Flat yaw distribution (top right) characterizes the rotational ambiguity of the shape. For the bottom figure, our method estimates two possible modes for x as shown in the corresponding kernel density estimate on the right which a deterministic ICP algorithm cannot capture.

cases, *e.g.* long featureless corridors or rotational symmetric objects such as bottles that admit infinite solutions. In the alignment of an under-constrained object where there is ambiguity in the matching, uncertainty relates to the geometry of the object. This type of uncertainty is intrinsic to the problem and known as *epistemic*. For example, in Figure 1 (top), the broad yaw distribution captures the rotational symmetry of the bottle. For a rectangular building block (bottom left), source cloud (cyan) can match to any of the two sides of the reference cloud (magenta), thus giving bi-modal x distribution (right) which a deterministic ICP algorithm cannot capture.

The quantification of the transformation uncertainty is crucial for many other tasks, especially those requiring robust pose estimates of an autonomous vehicle in an urban environment. In order to accurately localize a vehicle, a single source of pose information is usually not enough. Pose information from different sensors are usually probabilistically fused together [12], [52] in a state estimation framework such as simultaneous localization and mapping (SLAM) [4]. In such a framework, ICP pose estimates are fused with odometry measurements or GPS observations using Extended Kalman Filter (EKF) [16], [5], particle filter [47] or GraphSLAM [10] which require information about the uncertainty of pose parameters. This uncertainty estimate encompasses information of the reliability of sensory data and aids in improving the estimates of the pose parameters. In the localization and mapping tasks, this uncertainty information can improve loop closure detection and yield better maps. In the task of scene reconstruction, the amount of uncertainty in the transformation

Manuscript received: July, 20, 2021; Revised November, 11, 2021; Accepted December, 13, 2021.

This paper was recommended for publication by Editor Javier Civera upon evaluation of the Associate Editor and Reviewers’ comments.

This work was supported by the Australian Government Research Training Program (RTP) Scholarship.

* Corresponding author: fafz3958@uni.sydney.edu.au

¹ School of Computer Science, The University of Sydney, Australia

² NVIDIA, USA

³ ETH Zürich, Switzerland

Digital Object Identifier (DOI): see top of this page.

parameters can indicate a tangential drift that may occur in the registration of two plane objects.

Bayesian methods [39] can be used to incorporate uncertainty into ICP pose estimates. These methods provide a probabilistic framework in which Bayes' rule is used to obtain a posterior distribution given a prior distribution and a likelihood function. However, Bayesian methods tend to be computationally intractable for complicated likelihood functions and high-dimensional problems. To address the tractability issue, modern Bayesian methods rely on approximate techniques such as Markov Chain Monte Carlo (MCMC) [40] and variational inference (VI) [56] which can scale to large problems using stochastic gradient descent (SGD) [46].

MCMC methods approximate the intractable posterior distribution by drawing samples from the prior and likelihood functions. In contrast, VI transforms the problem into an optimization process that reduces the Kullback-Leibler (KL) divergence between the intractable posterior distribution and an analytically tractable variational distribution. MCMC will converge to the underlying distribution but can be slow in practice. VI on the other hand is typically faster but the expressiveness of the variational distribution is restricted and may not capture the complexity of the true posterior distribution.

In this paper we formulate a Stein variational gradient descent (SVGD) [30] method to approximate the posterior distribution of ICP pose parameters. SVGD combines the accuracy and flexibility of MCMC with the speed of VI using gradient descent. SVGD approximates the intractable posterior distribution with a non-parametric representation given by a set of particles. These particles are optimized with a functional gradient descent of KL divergence and provide a reliable uncertainty estimation by incorporating a *repulsive* term that prevents particles from clustering together and allows them to capture a wide range of complex distributions. These particles can be updated jointly and benefit from GPU parallelization. **Contribution:** The main contribution of this paper is a non-parametric point cloud registration method based on Stein variational gradient descent to estimate ICP's pose uncertainty. Our method exploits GPU-computation for increased speed and captures both epistemic and aleatory uncertainty [27] intrinsic in the matching problems. A python implementation of the method is publicly available¹.

II. RELATED WORK

There exists a large number of point cloud registration methods as reviewed in [42], [48]. They can be broadly categorized as algorithms that provide point estimates and probabilistic outputs.

A. Point-Based ICP Algorithms

Point estimate registration algorithms produce a single estimate of the transformation between two point clouds. These algorithms mostly vary in how they select points to perform data association, which error metric they use, and what optimization technique they employ for their cost function [48].

ICP algorithms which vary in the selection of points include those using random sampling [36] and uniform sub-sampling [35]. Other ICP methods have different cost functions. Notable choices are point-to-point [8] point-to-plane [14] and plane-to-plane [49]. The ICP cost function can be optimized using either closed form solutions *e.g.* singular value decomposition (SVD) [3], and quaternions [22] or other methods *e.g.* Levenberg-Marquardt [20], simulated annealing [31], and stochastic gradient descent [1].

Another line of research deals with eliminating outliers due to partial overlap, sensor noise and other complexities. This includes the use of a trimmed square cost function to estimate the optimal transformation [15], and rejecting pairs with distance larger than a certain threshold [49], [36].

An alternative class of registration technique reformulates ICP as an alignment of two Gaussian mixtures where a statistical discrepancy measure is minimized between the two mixtures [18]. In this case Expectation Maximization [17] can be used to provide robust solutions *i.e.* [21]. Learning-based point cloud registration methods employ neural networks to provide pose parameters [2], [54]. These methods rely on the data used during training and may not necessarily generalize well to an unseen dataset.

B. Uncertainty-Based ICP Algorithms

This class of ICP provides uncertainty estimates for the transformation parameters. The uncertainty of the transformation can be estimated using closed form solutions [9], [7], [11], [13], and sampling-based approaches [7], [24], [34]. Closed form solutions require the Hessian of the cost function. Among closed form solutions such as in [9], [7], [11], errors due to sensor noise are not properly captured as pointed out in [13], [44]. [13] considers the effect of sensor noise in the covariance structure but [32], [37] indicate that it is overoptimistic and hence not suitable for sensor fusion. Closed-form solutions are efficient but do not capture the uncertainty of the data association into account in the covariance estimation.

Sampling based approaches such as [7], [24] provide high quality samples of pose parameters but are computationally expensive. A recently proposed scalable probabilistic ICP method, Bayesian ICP [34], employs stochastic gradient Langevin dynamics (SGLD) [55] to provide samples of the posterior distribution in an online manner. However, due to it still being an MCMC style method it suffers from the inherent limitations and challenges of parallelizing any MCMC method.

Other recently proposed methods such as [32], [29] take initial pose estimates into account in the pose uncertainty prediction. [32] combines a closed-form solution to incorporate sensor noise with the unscented transform [25] to capture other sources of uncertainty. This method relies on an accurate covariance of the initial pose uncertainty to produce reliable uncertainty estimation which is not always available.

Our method can take prior information into account to estimate the posterior distribution over transformation parameters. When a prior is available, our method leverages this information to achieve faster and more accurate convergence. In comparison to Bayesian ICP which uses MCMC sampling,

¹<https://bitbucket.org/fafz/stein-icp/src/master/>

our method employs Stein particles that can be propagated in parallel through Stein's gradients, directly modeling the interactions between the particles and promoting diversity in the solution. This is particularly important in multi-modal problems as we demonstrate in a few examples. Additionally, our method is directly amenable to parallelization, therefore suitable to GPU computation.

III. PRELIMINARIES

In this section we first give an overview of standard ICP [8], followed by a description of stochastic gradient descent ICP (SGD-ICP) [1] which constitutes a core component of our method.

A. Standard ICP

Standard ICP aligns two input point clouds by iteratively performing the following two steps until convergence:

- 1) Establish the point pairings on the basis of minimum Euclidean distance using an initial guess $\theta = \{x, y, z, \text{roll}, \text{pitch}, \text{yaw}\}$. The commonly used point-to-point distance metric to select the corresponding pair of points is expressed as follows:

$$\text{point-to-point}(\mathbf{s}_i', \mathcal{R}) = \min_{\mathbf{r} \in \mathcal{R}} \|\mathbf{r} - \mathbf{s}_i'\|, \quad (1)$$

where $\mathbf{s}_i, \mathbf{r}_i \in \mathbb{R}^3$ are points in 3D space belonging to a source cloud $\mathcal{S} = \{\mathbf{s}_i\}_{i=1}^N$ and a reference cloud $\mathcal{R} = \{\mathbf{r}_i\}_{i=1}^M$ respectively. $\mathbf{s}_i' = (R\mathbf{s}_i + \mathbf{u})$ is a transformed point in the source cloud, $\mathbf{u} \in \mathbb{R}^{3 \times 1}$ is a translation vector comprising of $\theta_{1:3}$ and $R \in \mathbb{R}^{3 \times 3}$ is a rotation matrix parametrized by $\theta_{4:6}$.

- 2) Find the transformation that minimizes a loss function defined by the Euclidean distance between each pair of corresponding points. The point-to-point cost function has the following form:

$$\operatorname{argmin}_{\theta} \mathcal{L}(\theta) = \frac{1}{N} \sum_i \|(R\mathbf{s}_i + \mathbf{u}) - \mathbf{r}_i\|^2. \quad (2)$$

B. Stochastic Gradient Descent ICP

SGD-ICP [1] optimizes the ICP cost function (2) using stochastic gradient descent (SGD) [46]. Instead of using all available points, SGD-ICP samples a mini-batch \mathcal{M} of m points from the source cloud to find corresponding points from the reference cloud. SGD-ICP computes gradients of the cost function (2) to update the transformation parameters θ with the following update rule:

$$\theta^{t+1} = \theta^t - \eta A \bar{g}(\theta^t, \mathcal{M}^t), \quad (3)$$

where θ^t and θ^{t+1} are the values of the pose parameters at the current and next iteration respectively. The matrix $A \in \mathbb{R}^{6 \times 6}$ acts as a pre-conditioner for the gradients and the learning rate η dictates how quickly parameter values change. The average

gradients \bar{g} of the parameter vector θ with respect to the loss \mathcal{L} are computed as follows:

$$\bar{g}(\theta_{1:3}^t, \mathcal{M}^t) = \frac{1}{m} \sum_i \left((R^t \mathbf{s}_i + \mathbf{u}^t) - \mathbf{r}_i \right) \frac{\partial \mathbf{u}^t}{\partial \theta_{1:3}^t}, \quad (4)$$

$$\bar{g}(\theta_{4:6}^t, \mathcal{M}^t) = \frac{1}{m} \sum_i \left((R^t \mathbf{s}_i + \mathbf{u}^t) - \mathbf{r}_i \right) \frac{\partial R^t}{\partial \theta_{4:6}^t} \mathbf{s}_i, \quad (5)$$

where $\bar{g}(\theta_{1:3}^t, \mathcal{M}^t)$ are the gradients of the translation components and $\bar{g}(\theta_{4:6}^t, \mathcal{M}^t)$ are the rotational gradients.

The quality of the estimated pose of SGD-ICP is the same as that of standard ICP but it is computationally more efficient due to the mini-batch formulation [1]. In the next section we employ the mini-batch gradient-based formulation of SGD-ICP within Stein variational inference to derive Stein ICP.

IV. STEIN ICP

Our proposed method, Stein ICP, makes the connection between the gradient-based stochastic optimization of an ICP cost function and Stein variational inference. Specifically, Stein ICP utilizes the SGD-ICP [1] gradients within the Stein variational gradient descent (SVGD) framework [30]. SVGD approximates an intractable but differentiable posterior distribution $p(\mathbf{x})$ by constructing a non-parametric variational distribution represented by a set of K particles $\{\mathbf{x}_j\}_{j=0}^K$. These particles are updated iteratively by an update rule of the following form:

$$\mathbf{x}_j \leftarrow \mathbf{x}_j + \eta \phi(\mathbf{x}_j) \quad \forall j = 1, \dots, K, \quad (6)$$

where η is a small step size, and $\phi: \mathbb{R}^d \rightarrow \mathbb{R}^d$ characterizes a perturbation direction that gives the steepest descent. ϕ should be chosen to push-forward the particle distribution closer to the target. The optimal value of the perturbation direction ϕ can be obtained from the following expression:

$$\phi^* = \operatorname{argmax}_{\phi \in \mathcal{B}} \left(-\frac{d}{d\eta} \text{KL}(q_{[\eta\phi]} \parallel p)|_{\eta=0} \right), \quad (7)$$

where $q_{[\eta\phi]}$ is specified nonparametrically by the updated particles, $\mathbf{x}' = \mathbf{x} + \eta \phi(\mathbf{x})$. ϕ^* is selected from a function set, \mathcal{B} , to minimize the KL divergence between the particle distribution and the target [53].

To obtain tractable and flexible solutions, SVGD chooses \mathcal{B} to be in the unit ball of a vector-valued reproducing kernel Hilbert space (RKHS) $\mathcal{H}^d = \mathcal{H} \times \dots \times \mathcal{H}$, where \mathcal{H}^d is an RKHS formed by scalar-valued functions associated with a positive definite kernel $k(\mathbf{x}, \mathbf{x}')$, that is, $\mathcal{B} = \{\phi \in \mathcal{H}^d: \|\phi\|_{\mathcal{H}^d} \leq 1\}$. Formally, an RKHS is defined as a Hilbert space \mathcal{H}^d with inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}^d}$ and norm $\|\cdot\|_{\mathcal{H}^d}$ satisfying the reproducing property, $\forall f \in \mathcal{H}^d, f(\mathbf{x}) = \langle f, k(\mathbf{x}, \cdot) \rangle_{\mathcal{H}^d}$. A key observation to solve Eq. 7 is that the optimization is a linear functional of ϕ that connects to the Stein operator \mathcal{A} in the following equation [30],

$$-\frac{d}{d\eta} \text{KL}(q_{[\eta\phi]} \parallel p)|_{\eta=0} = \mathbb{E}_{\mathbf{x} \sim q} [\text{trace}(\mathcal{A} \phi(\mathbf{x}))], \quad (8)$$

with

$$\mathcal{A} \phi(\mathbf{x}) = \nabla_{\mathbf{x}} \log p(\mathbf{x}) \phi(\mathbf{x})^\top + \nabla_{\mathbf{x}} \phi(\mathbf{x}). \quad (9)$$

By Stein's identity in Eq. 8, we obtain $\frac{d}{d\eta} \text{KL}(q[\eta\phi] \parallel p) = 0$ when $p = q$. Following from Eq. 9, the gradient of the log posterior $\nabla_{\mathbf{x}} \log p(\mathbf{x}|D)$ can be expressed as:

$$\nabla_{\mathbf{x}|D} \log p(\mathbf{x}) = \sum_{i=1}^N \nabla_{\mathbf{x}} \log p(d_i|\mathbf{x}) + \nabla_{\mathbf{x}} \log p(\mathbf{x}), \quad (10)$$

where $\nabla_{\mathbf{x}} \log p(d_i|\mathbf{x})$ and $\nabla_{\mathbf{x}} \log p(\mathbf{x})$ are the gradient of the log likelihood function, and the gradient of the log of the prior with respect to \mathbf{x} respectively, for a dataset $D = \{d_i\}_{i=1}^N$ of N data points. Note that the Stein operator does not depend on the normalization $p(D)$, also known as the marginal likelihood, which makes the computation of the posterior significantly easier. This makes SVGD a powerful tool for inference of intractable distributions. The Stein variational gradient descent algorithm follows from a closed form solution of (7), as shown in [30], and given by:

$$\phi^*(\cdot) = \mathbb{E}_{\mathbf{x} \sim q} [\nabla_{\mathbf{x}} \log p(\mathbf{x}) k(\mathbf{x}, \cdot) + \nabla_{\mathbf{x}} k(\mathbf{x}, \cdot)]. \quad (11)$$

In the above we omit the dependence on the data D to simplify notation. Equation (11) provides the optimal update direction for the particles within \mathcal{H}^d . In practice, a set of particles $\{\mathbf{x}_j\}_{j=1}^K$ approximates q yielding the following update rule:

$$\mathbf{x}_j \leftarrow \mathbf{x}_j + \eta \hat{\phi}^*(\mathbf{x}_j), \quad (12)$$

where $\hat{\phi}^*(\mathbf{x})$ is an approximate steepest direction given by

$$\hat{\phi}^*(\mathbf{x}) = \sum_{j=1}^K [\nabla \log p(\mathbf{x}_j) k(\mathbf{x}_j, \mathbf{x}) + \nabla_{\mathbf{x}_j} k(\mathbf{x}_j, \mathbf{x})]. \quad (13)$$

The first gradient term in (13) is weighted by a kernel function, which smooths the gradients, and acts as the steepest direction for the log probability. The second term is the gradient of the kernel function. It can be viewed as a repulsive force which induces the spread among the particles and prevents them from collapsing to the local modes of the log probability.

In SGD-ICP $-\nabla_{\mathbf{x}} \log p(D|\mathbf{x})$, in Eq. 10, is expressed as the gradients of the cost function in (4) and (5). $\nabla_{\mathbf{x}} \log p(\mathbf{x})$ represents the gradient of the log of Gaussian priors for translations and von Mises priors for rotations [34]:

$$\nabla_{\theta} \log p(\theta_{1:3}) = -(\theta_{1:3} - \mu_{1:3})/\sigma_{1:3}, \quad (14)$$

$$\nabla_{\theta} \log p(\theta_{4:6}) = -\kappa_{4:6} \sin(\theta_{4:6} - \mu_{4:6}), \quad (15)$$

where $\mu_{1:6}$ represents the mean while $\sigma_{1:3}$, and $\sigma_{4:6} = 1/\kappa_{4:6}$ represent the variance of the prior distributions of the translation and rotation components respectively.

In Stein ICP, we replace the gradients in (13) with the mini-batch gradients of the ICP cost function shown in (4) and (5), along with the gradients of priors expressed in (14) and (15). By doing this, Stein variational gradients are independently obtained for translations ($\theta_{1:3}$) and rotations ($\theta_{4:6}$), for all the particles in (16), which are then updated using (12).

$$\hat{\phi}^*(\theta) = \sum_{j=1}^K \left[-\left(N\bar{g}(\theta_j, \mathcal{M}) + \nabla_{\theta} \log p(\theta_j) \right) k(\theta_j, \theta) + \nabla_{\theta} k(\theta_j, \theta) \right]. \quad (16)$$

In SVGD the kernel plays a critical role in weighing the gradients and dispersing the particles. In Stein ICP, for translation parameters we use an RBF kernel $k(\theta_{1:3}, \theta'_{1:3}) = \exp(-\frac{1}{h} \|\theta_{1:3} - \theta'_{1:3}\|_2^2)$, where h is the bandwidth of the kernel. For the rotational parameters, we use a modified version of the RBF kernel as the Euclidean distance is not directly applicable to angles. In particular we use the following modified RBF kernel for rotations:

$$k(\theta_{4:6}, \theta'_{4:6}) = \exp \left[-\frac{1}{h} \left(\arctan 2 \left(\sin(\theta_{4:6} - \theta'_{4:6}), \cos(\theta_{4:6} - \theta'_{4:6}) \right) \right)^2 \right]. \quad (17)$$

In summary, Stein ICP begins by initializing a set of K particles randomly. In general the implementation of Stein variational inference does not depend on the initial distribution [30]. In practice, the particles can be initialized using prior knowledge, *e.g.* readings from an inertial measurement unit or from a previous solution. Each particle represents the transformation of a mini-batch sampled from a source cloud \mathcal{S} to a reference cloud \mathcal{R} , producing K transformed mini-batches. Next for all points in each transformed mini-batch, corresponding closest points from a reference cloud are sought and stored in pairs using (1). Then mean gradients are estimated for all the matching pairs belonging to each of the K particles using (4) and (5). Next, Stein variational gradients are obtained independently for translations and rotations using (16) which are then used to update each particle with (12). This procedure is repeated for T iterations producing K particles representing the posterior distribution.

V. EXPERIMENTS

In the experiments we demonstrate the ability of Stein ICP to obtain high quality pose distributions while its formulation leads to significant computational gains when using a GPU.

To gain an intuitive insight into the behaviour of Stein ICP, we use objects with obvious distributions from an RGB-D dataset [28]. For the quantitative evaluations we use a challenging point cloud alignment dataset [43] which contains different sequences recorded in both structured and unstructured environments. We provide comparisons with three other uncertainty aware ICP methods, namely Bayesian ICP [34], Cov-3D ICP [32], and Closed-form ICP [13].

As there is no ground truth distribution available for the point cloud alignment we compute this ground truth information using Monte Carlo sampling. To obtain these Monte Carlo samples in a timely manner, we use SGD-ICP, which produces solutions equivalent to standard ICP methods in quality but much more quickly [1]. We run 1000 instances of SGD-ICP with initial transformation estimates sampled from ± 1 m and ± 0.1745 rad for translations and rotations respectively. The resulting samples form the ground truth distribution. As a measure of the quality of the distribution estimated by a method we compute the Kullback-Leibler (KL) divergence and overlapping coefficient (OVL) [23] between the obtained distribution and the ground truth distribution. A value of 0 and

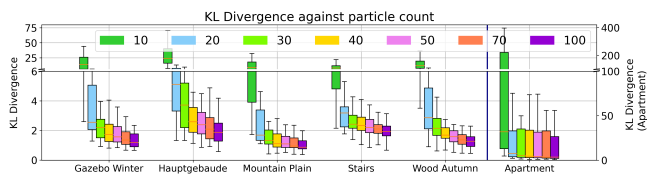


Fig. 2: KL Divergence against particle count. Overall, the best performance is achieved with 100 particles. Note however that there is little improvement in KL divergence between 40 and 100 particles.

1 for *OVL* indicates no-overlap and full overlap between the two probability distributions respectively.

In all experiments, Stein ICP uses Adam [26] as the optimizer with a step size of 0.03 and batch size of 150 for the RGB-D data and a step size of 0.01 and batch size of 300 for the Challenging dataset [43]. The bandwidth parameter required by SVGD is chosen using the median heuristic as described in [30]. Furthermore, uniform priors are assumed as no prior knowledge is available. Bayesian ICP generates 1000 samples and uses a step size of 0.008. Cov-3D and Closed-form ICP use point-to-plane SGD-ICP estimated pose to compute the covariance matrix using the default parameter setting of the author’s implementation [32].

A. Particle count

We begin by evaluating the effect the number of particles has on the pose distribution quality on the Challenging dataset. Figure 2 shows the KL divergence between the Stein ICP and ground truth distribution using the point-to-point error metric as a function of the particle count. The right Y-axis is for the *Apartment* sequence only due to the difference of the scale. Overall the KL divergence for both translations (top) and rotations (bottom) improves with a larger number of particles. However, the improvement in KL divergence between 40 and 100 particles is minor compared to the change from 10 to 40 particles. In the subsequent experiments we will use 100 particles as this reliably achieves high quality results.

B. Distribution Quality Analysis

We start with a qualitative evaluation of the distribution estimated by Stein ICP on two examples from the RGB-D dataset. In Figure 3 we show the source (cyan), reference (magenta), and aligned (green) clouds for a *bowl* (left) and a *mug* (right). The kernel density estimates (KDE) of both the ground truth samples (left) and Stein ICP samples (right) are shown in Figure 4. We can see how for the *bowl* (top), which has a rotational symmetry around yaw, the estimates are peaked with the exception of yaw which spans the full $[-\pi, \pi]$ range, both in the ground truth distribution and Stein ICP. By contrast the distributions for the *mug* (bottom) are peaked for all parameters. As the handle of the *mug* clearly identifies the single correct solution, this uni-modal distribution is what we would expect. These results demonstrate that the distributions recovered by Stein ICP match the ground truth distributions and that both distributions agree with our intuitions about object alignment ambiguity and uniqueness.

Next, we perform a quantitative comparison between our method and the state-of-the-art probabilistic ICP methods




Fig. 3: Source (cyan), reference (magenta) and aligned cloud (green) of the *bowl* (left) and the *mug* (right). The *bowl* is rotationally symmetric around yaw, while handle of the *mug* constrains the yaw estimate.

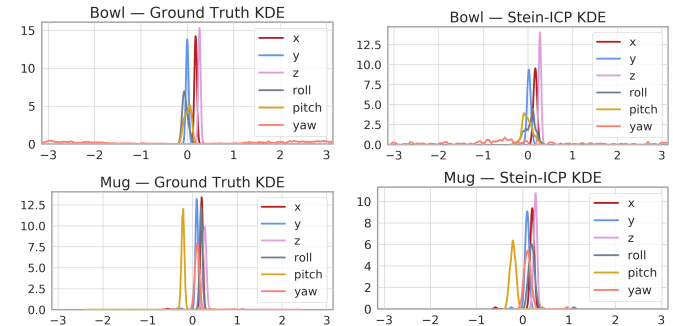


Fig. 4: Visualization of the pose distribution estimated by a kernel density estimator (KDE) for the ground truth and Stein ICP samples. The broad yaw distribution correctly captures the rotational symmetry of the round *bowl* (top). This is in contrast to the peaked yaw distribution of the *mug* (bottom) which has no symmetries.

namely *Bayesian ICP* [34], a multi-chain Bayesian ICP—denoted by *Bayesian-multi*—this is the result of running 10 chains of Bayesian ICP and compounding the results, *COV-3D ICP* [32], and a popular benchmark *Closed-form ICP* [13] on the Challenging dataset. The results are summarised in Table I which shows the median of the KL divergence and *OVL* as well as the 10th and 90th quantile in the brackets. The divergence and *OVL* are computed over all consecutive pairs in each sequence with all methods using the point-to-plane error metric.

From those numbers we can see that our method outperforms both *COV-3D* and *Closed-form ICP* on all of the sequences. The large values of KL divergence and small values of *OVL* for *COV-3D* and *Closed-form ICP* indicate that these methods fail to provide correct uncertainty estimates. The *Closed-form ICP* method only provides reliable estimates in situations where uncertainty arises solely due to sensor noise [32], which is not the case in many of the scenes in the Challenging dataset. The *COV-3D* method relies on an accurate initial pose uncertainty estimation to provide a correct pose covariance. This causes the algorithm to under- or over-estimate the pose covariance when the initial covariance estimate is either inexact or not available [32].

When comparing Stein ICP with *Bayesian ICP* and *Bayesian-multi* we observe that in most scenes Stein ICP performs better. The only exception is *Apartment* where *Bayesian ICP* performs better. As the KL divergence is computed based on the covariance of the samples, the fact that Stein ICP only produces 100 samples while Bayesian ICP produces 1000 samples can have an impact on the KL divergence. However, overall, both methods produce results of comparable quality. Similar results can be observed for the *OVL* measure, where

TABLE I: Quality Comparison of the Stein ICP against other methods using median KL Divergence and overlapping coefficient (OVL) with a 10th and a 90th quantile written in the brackets.

Sequence	Method	KL Div.	OVL
Apartment	Ours	5.7 [1.2, 4e ²]	0.7 [0.4, 0.9]
	Bayesian-multi	4.2 [0.8, 1e ²]	0.7 [0.3, 0.9]
	Bayesian ICP	4.6 [1.6, 2e ²]	0.7 [0.4, 0.8]
	COV-3D ICP	2e ² [17.2, 2e ⁴]	0.3 [0.1, 0.5]
	Closed-form ICP	1e ⁴ [8e ² , 9e ⁶]	0.0 [0.0, 0.0]
Gazebo Winter	Ours	1.1 [0.5, 3e ²]	0.9 [0.4, 0.9]
	Bayesian-multi	2.9 [1.0, 4e ²]	0.8 [0.4, 0.9]
	Bayesian ICP	3.0 [1.2, 3e ²]	0.8 [0.4, 0.9]
	COV-3D ICP	42.4 [7.7, 1e ³]	0.3 [0.1, 0.4]
	Closed-form ICP	1e ⁵ [4e ⁴ , 1e ⁷]	0.0 [0.0, 0.0]
Hauptgebäude	Ours	0.6 [0.3, 17.6]	0.9 [0.8, 0.9]
	Bayesian-multi	1.9 [1.3, 13.2]	0.8 [0.7, 0.8]
	Bayesian ICP	1.8 [1.3, 15.7]	0.8 [0.7, 0.8]
	COV-3D ICP	2e ² [55.5, 8e ²]	0.2 [0.1, 0.3]
	Closed-form ICP	4e ⁴ [1e ⁴ , 1e ⁵]	0.0 [0.0, 0.0]
Mountain Plain	Ours	2.1 [1.6, 12.7]	0.7 [0.6, 0.9]
	Bayesian-multi	62.0 [1.9, 5e ²]	0.5 [0.3, 0.8]
	Bayesian ICP	88.9 [4.3, 3e ²]	0.5 [0.3, 0.7]
	COV-3D ICP	16.1 [6.4, 69.4]	0.3 [0.2, 0.5]
	Closed-form ICP	4e ⁴ [1e ⁴ , 1e ⁵]	0.0 [0.0, 0.0]
Stairs	Ours	2.6 [0.9, 3.7]	0.7 [0.5, 0.8]
	Bayesian-multi	1.9 [1.0, 5.7]	0.8 [0.5, 0.9]
	Bayesian ICP	2.1 [0.9, 4.9]	0.7 [0.5, 0.9]
	COV-3D ICP	16.1 [11.3, 1e ²]	0.3 [0.2, 0.4]
	Closed-form ICP	4e ³ [6e ² , 8e ⁴]	0.0 [0.0, 0.0]
Wood Autumn	Ours	0.6 [0.4, 1.7]	0.9 [0.8, 0.9]
	Bayesian-multi	1.2 [0.8, 2.5]	0.9 [0.8, 0.9]
	Bayesian ICP	1.2 [0.7, 3.2]	0.9 [0.6, 0.9]
	COV-3D ICP	8.1 [7.1, 9.4]	0.4 [0.3, 0.5]
	Closed-form ICP	4e ⁴ [2e ⁴ , 1e ⁵]	0.0 [0.0, 0.0]

Stein ICP performs better overall.

C. Odometry

In this section we use Stein ICP to compute odometry estimates by propagating the frame-to-frame expectations of the ICP transformation distributions to obtain a global pose estimate. These results are compared with the odometry estimate obtained by Bayesian ICP. As our method provides samples of the transformation parameters and can capture multi-modal distributions, the overall expectation and covariance may not represent the underlying correct pose.

For a sequence of l scans the individual ICP pose estimates are combined as $T_E = T_0^1 T_1^2 \dots T_{l-1}^l$. Where T_{l-1}^l is a 4×4 homogeneous transformation matrix that captures the mean rotation and translation between scan l and $l-1$. A 4th order approximation [6] is used to combine the individual scan-to-scan covariances. The resulting trajectories for *Gazebo Winter*, *Stairs*, and *Newer College dataset* [45] are overlaid onto the ground plane to visualize the estimated trajectories in Figure 5. The lines display the trajectories and the ellipses show the 3- σ (95%) confidence sets for uncertainty estimation of Stein ICP and Bayesian ICP.

The plots in Figure 5 show that the mean trajectory obtained from Stein ICP and Bayesian ICP estimates are similar and track the ground truth trajectory well. The covariance ellipses estimated using Bayesian ICP are larger and less certain than those of Stein ICP. As the odometry estimates are obtained by accumulating transformations small differences in errors can

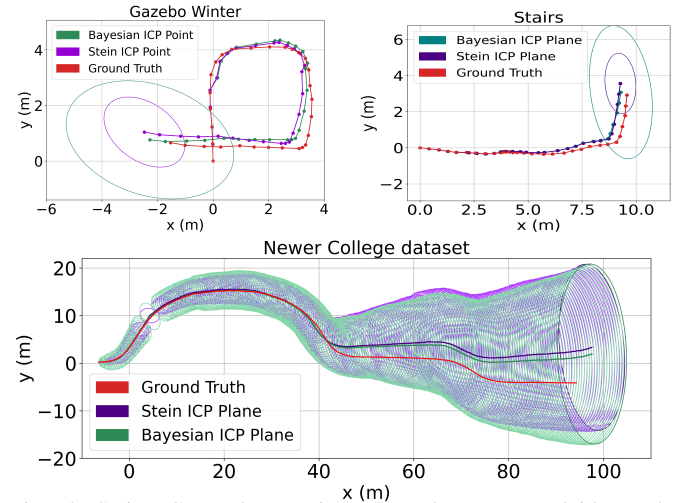


Fig. 5: Stein ICP and Bayesian ICP odometry overlaid on the ground truth trajectory for *Gazebo Winter*, *Stairs*, and *Newer College dataset* [45]. The ellipses represent the 95% confidence set of the compound covariance at the last pose of odometry for each method for sequences and along the entire trajectory for [45].

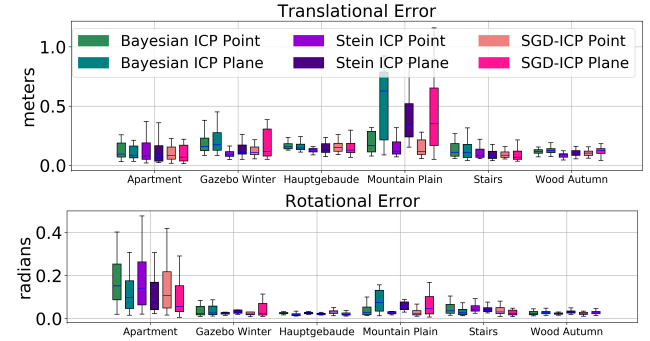


Fig. 6: Translational (top) and rotational (bottom) error distributions over entire trajectories of Stein ICP in comparison to Bayesian ICP and SGD-ICP method on the Challenging dataset. Overall, Stein ICP gives equivalent performance to that of SGD-ICP indicating that Stein ICP does not suffer from a reduction in quality despite providing uncertainty estimates.

add up over the trajectory, amplifying the small differences between Stein ICP and Bayesian ICP reported in Table I. Overall both methods provide good odometry trajectories and covariance estimates. The ability to estimate odometry trajectories including uncertainty is crucial in many robotic applications such as autonomous navigation and SLAM. More generally any method reliant on pose estimation can use the uncertainty provided by these methods for decision making.

D. Impact of Distribution Estimation on Mean Pose Estimation Quality

Estimating the full distribution over pose parameters is useful. However, if a single solution exists what, if any, is the price in pose estimation accuracy being paid in comparison to a standard point estimation method? To answer this question we examine the trajectory estimation errors for Stein ICP, Bayesian ICP, and SGD-ICP all using both the point-to-point and the point-to-plane error metric. To measure the local accuracy of the trajectory estimation we compute the *relative pose error* [51] (RPE) for the SGD-ICP solution and

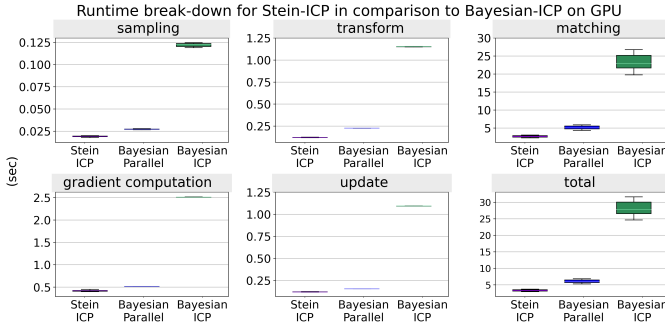


Fig. 7: Break-down of the runtime (sec) on the GPU. The times correspond to mini-batch sampling (top left), transforming point cloud with the updated transformation time (top middle), matching time (top right), gradient time (bottom left), and update time (bottom middle). Bottom right shows the total time. For Stein ICP, the gradients computational time includes the kernel estimation and repulsive force. These results demonstrate that Stein ICP is more efficient than both versions of Bayesian ICP.

each sample of the distributions produced by Stein ICP and Bayesian ICP. The RPE is defined as follows:

$$E = (\hat{T}_i^{-1} \hat{T}_{i+\Delta})^{-1} (T_i^{-1} T_{i+\Delta}), \quad (18)$$

where $(\hat{T}_i^{-1} \hat{T}_{i+\Delta})$ is the ground truth transformation and $(T_i^{-1} T_{i+\Delta})$ is the ICP transformation between scan i and $i+\Delta$, with $\Delta = 1$ here.

The translational error is computed using the Euclidean distance of the translation components of E while the rotational errors are computed using absolute angular differences between the orientation of the ground truth and the estimated pose. For sample based methods we obtain a single error by averaging together the per sample errors. In Figure 6 we show the translational (top) and rotational (bottom) error distributions over all the consecutive scan pairs of each used sequence independently.

From this we can see that the estimates of the distribution based methods are not any worse than those of a point estimation method. The choice of error metric has a bigger impact on the results than whether or not a distribution is estimated. This means that we do not give up any accuracy as far as transformation estimation goes with Stein ICP.

E. Runtime

Finally, we compare the run-time of Stein ICP, Bayesian ICP and Bayesian-multi. Traditionally, obtaining uncertainty estimates for ICP methods was computationally expensive, especially if accurate but slow Markov Chain Monte Carlo based methods were used. Runtime efficiency is one of the main advantages of Stein ICP due to its parallelism and leverage of GPUs. In the following, we present a break down of the runtime by the individual components of Stein ICP compared to Bayesian ICP. To this end, we use the *Gazebo Winter* dataset and run Stein ICP for 100 iterations using 100 particles. In Bayesian ICP, each iteration provides us with a sample. However we need to run the method over a certain number of iterations before the Markov Chain becomes stable; a process known as burn-in. To this end, we run Bayesian ICP for 1100 iterations to collect 1000 sequential samples after discarding initial 100 burn-in samples. We also

run 10 parallel chains of Bayesian-multi where each chain collects 100 samples with a burn-in period of 100 iterations to produce a distribution of equivalent quality. All computations are performed on a desktop computer with an NVidia Titan-V GPU.

Figure 7 presents the runtime of the various components along the Y-Axis. Even though Stein ICP is run with 100 particles, which is roughly equivalent to running 100 independent SGD-ICP instances, Stein ICP takes $\frac{1}{8^{th}}$ of the time Bayesian ICP requires for each ICP component. This is because Stein ICP’s particles can trivially be parallelized on a GPU. In contrast, Bayesian ICP produces samples under the Markov assumption, *i.e.* every sample depends on the previous one, therefore it cannot reliably exploit GPU parallelism to generate samples. Even parallel Bayesian ICP (if run for efficiency) consumes relatively more time. This is due to the burn-in period necessary for each parallel chain of Bayesian ICP to ensure convergence to the true distribution. Stein ICP ability to perform computations in parallel on a per particle basis means that the GPU implementation can do more work for 100 particles in roughly the same time as required by a single particle. This results in over 5 times speedup compared to Bayesian ICP. Both Cov-3D and Closed-form methods require at least solving a point-to-plane ICP first (2.62 ± 0.19 sec on the GPU) in addition to the time required to compute the covariances. In addition, Cov-3D requires 12 registration steps for the unscented transform.

VI. CONCLUSION

In this paper we devised Stein ICP, a probabilistic generalization of the popular ICP that provides accurate posterior pose distributions. Stein ICP exploits the stochastic optimization nature of SGD-ICP combining it with the recent Stein variational gradient descent (SVGD) algorithm to approximate posterior distributions of transformation parameters between two point clouds using particles. The independent nature of these particles makes Stein ICP suitable for GPU parallelization. Extensive experiments using both RGB-D as well as LiDAR data demonstrate the capability of our method in providing high-quality pose parameter distributions in only a few seconds of computation. In contrast to Markov chain Monte Carlo solutions to probabilistic point-cloud matching, Stein ICP propagates gradients of particles in parallel, connecting gradient descent on the objective function with a repulsive term that ensures particles are aligned and capture the tail of the underlying distribution. This opens up new possibilities to leverage calibrated uncertainty estimation in applications requiring robust point cloud matching such as simultaneous localization and mapping or 3D reconstruction. Another exciting direction on using the uncertainty provided by the method is for grasping unknown objects based on a list of known objects with associated grasping positions. The uncertainty can be used to indicate which of the known objects is a better match and which direction (with less uncertainty) to grasp the unknown object.

REFERENCES

- [1] F. Afzal Maken, F. Ramos, and L. Ott. Speeding up iterative closest point using stochastic gradient descent. In *IEEE ICRA*, 2019. 2, 3, 4
- [2] Y. Aoki, H. Goforth, R.A. Srivatsan, and S. Lucey. Pointnetlk: Robust and efficient point cloud registration using pointnet. *CVPR*, 2019. 2
- [3] K. S. Arun, T. S. Huang, and S. D. Blostein. Least-squares fitting of two 3-d point sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (5):698–700, 9 1987. 2
- [4] M. Barczyk and S. Bonnabel. Towards realistic covariance estimation of ICP-based Kinect V1 scan matching: The 1D case. *Proceedings of the American Control Conference*, 2017. 1
- [5] M. Barczyk, S. Bonnabel, J.E. Deschaud, and F. Goulette. Invariant ekf design for scan matching-aided localization. *IEEE Transactions on Control Systems Technology*, 23, no. 6:2440–2448, November 2015. 1
- [6] T. D. Barfoot and P. T. Furgale. Associating uncertainty with three-dimensional poses for use in estimation problems. *IEEE Transactions on Robotics*, 30(3):679–693, 2014. 6
- [7] O. Bengtsson and A.J. Baerwaldt. Robot localization based on scan-matching—estimating the covariance matrix for the icp algorithm. *Robotics and Autonomous Systems*, 44(1):29 – 40, 2003. 2
- [8] P. Besl and N.D. McKay. Method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1992. 1, 2, 3
- [9] P. Biber and W. Strasser. The normal distributions transform: a new approach to laser scan matching. In *IROS*, pages 2743–2748, 2003. 2
- [10] D. Borrmann, J. Elseberg, K. Lingemann, A. Nüchter, and J. Hertzberg. Globally consistent 3d mapping with scan matching. *Robotics and Autonomous Systems*, 56(2):130 – 142, 2008. 1
- [11] M. Bosse and R. Zlot. Map matching and data association for large-scale two-dimensional laser scan-based SLAM. *The International Journal of Robotics Research.*, 27(6):667–691, 2008. 2
- [12] V. Caglioti, F. Mainardi, M. Pilu, and D. Sorrenti. Improving pose estimation using image, sensor and model uncertainty. In *BMVC*, 1994. 1
- [13] A. Censi. An accurate closed-form estimate of ICP’s covariance. In *IEEE ICRA*, pages 3167–3172, 2007. 2, 4, 5
- [14] Y. Chen and G. Medioni. Object modeling by registration of multiple range images. *Image Vision Computing*, 10(3):145–155, 1992. 2
- [15] D. Chetverikov, D. Svirko, D. Stepanov, and P. Krsek. The trimmed iterative closest point algorithm. *International Conference on Pattern Recognition*, 2002. 2
- [16] D. M. Cole and P. M. Newman. Using laser range data for 3d slam in outdoor environments. In *Proceedings of IEEE International Conference on Robotics and Automation*, May 2006. 1
- [17] 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):1–38, 1977. 2
- [18] B. Eckart, K. Kim, and J. Kautz. Hgmr: Hierarchical gaussian mixtures for adaptive 3d registration. In *ECCV*, 2018. 2
- [19] S. Izadi; D. Kim; O. Hilliges; D. Molyneaux; R. Newcombe; P. Kohli; J. Shotton; S. et al. Kinectfusion: real-time 3d reconstruction and interaction using a moving depth camera. In *Proceedings of ACM symposium on User interface software and technology*, 2011. 1
- [20] A. Fitzgibbon. Robust registration of 2D and 3D point sets. *Image and Vision Computing*, 21:1145–1153, 2003. 2
- [21] W. Gao and R. Tedrake. Filterreg: Robust and efficient probabilistic point-set registration using gaussian filter and twist parameterization. In *Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [22] B.K. P. Horn. Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America A*, 4, 1987. 2
- [23] Henry F. Inman and Edwin L. Bradley. The overlapping coefficient as a measure of agreement between probability distributions and point estimation of the overlap of two normal densities. *Communications in Statistics-theory and Methods*, 18:3851–3874, 1989. 4
- [24] T. M. Iversen, A. G. Buch, and D. Kraft. Prediction of ICP pose uncertainties using Monte Carlo simulation with synthetic depth images. In *International Conference on Intelligent Robots and Systems*, 2017. 2
- [25] S. Julier, J. Uhlmann, and H. F. Durrant-Whyte. A new method for the nonlinear transformation of means and covariances in filters and estimators. *IEEE Transactions on Automatic Control*, 2000. 2
- [26] D. Kingma and J. Ba. Adam: a method for stochastic optimization. In *International Conference on Learning Representations*, 2015. 5
- [27] Armen Der Kiureghian and Ove Ditlevsen. Aleatory or epistemic? does it matter? *Structural Safety*, 2009. 2
- [28] K. Lai, L. Bo, X. Ren, and D. Fox. A large-scale hierarchical multi-view rgb-d object dataset. In *IEEE ICRA*, 2011. 4
- [29] D. Landry, F. Pomerleau, and P. Giguère. Cello-3d: Estimating the covariance of icp in the real world. In *ICRA*, 2018. 2
- [30] Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose bayesian inference algorithm. In *NIPS*, 2016. 2, 3, 4, 5
- [31] J. Luck, C. Little, and W. Hoff. Registration of range data using a hybrid simulated annealing and iterative closest point algorithm. In *IEEE International Conference on Robotics and Automation*, 2000. 2
- [32] A. Barrau M. Brossard, S. Bonnabel. A new approach to 3d icp covariance estimation. *IEEE Robotics and Automation Letters*, 2020. 2, 4, 5
- [33] B. Ma and R. Ellis. Surface-based registration with a particle filter. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, page 566–573, 2004. 1
- [34] F. Afzal Maken, F. Ramos, and L. Ott. Estimating motion uncertainty with bayesian icp. *ICRA*, 2020. 2, 4, 5
- [35] T. Masuda. Generation of geometric model by registration and integration of multiple range images. In *Proceedings of International Conference on 3-D Digital Imaging and Modeling*, 3DIM, 2001. 2
- [36] T. Masuda and N. Yokoya. Robust method for registration and segmentation of multiple range images. *Computer Vision and Image Understanding*, 1995. 2
- [37] E. Mendes, P. Koch, and S. Lacroix. Icp-based pose-graph slam. *IEEE International Symposium on Safety, Security, and Rescue Robotics*, 2016. 2
- [38] W. Ming-Long. An augmented reality system using improved-iterative closest point algorithm for on-patient medical image visualization. *Sensors (Basel, Switzerland)*, 2018. 1
- [39] R. M. Neal. *Bayesian Learning for Neural Networks*. Springer-Verlag, Berlin, Heidelberg, 1996. 2
- [40] R. M. Neal. MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2010. 2
- [41] I.K. Park, M. Germann, M D. Breitenstein, and HPfister. Fast and automatic object pose estimation for range images on the gpu. *Machine Vision and Applications*, 21:749–766, 08/2010 2010. 1
- [42] F. Pomerleau, F. Colas, and R. Siegwart. A review of point cloud registration algorithms for mobile robotics. *Foundations and Trends in Robotics*, 4(1):1–104, May 2015. 2
- [43] F. Pomerleau, M. Liu, F. Colas, and R. Siegwart. Challenging data sets for point cloud registration algorithms. *The International Journal of Robotics Research*, 31:1705 – 1711, 2012. 4, 5
- [44] S. M. Prakhya, L. Bingbing, Y. Rui, and W. Lin. Covariance of ICP Algorithm A Closed-form Estimate of 3D ICP Covariance. In *International Conference on Machine Vision Applications*, 2015. 2
- [45] M. Ramezani, Y. Wang, M. Camurri, D. Wisth, M. Mattamala, and M. Fallon. The newer college dataset: Handheld lidar, inertial and vision with ground truth. In *IROS*, 2020. 6
- [46] H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3):400–407, 09 1951. 2, 3
- [47] J. Röwekämper, C. Sprunk, G D Tipaldi, C. Stachniss, P. Pfaff, and W. Burgard. On the position accuracy of mobile robot localization based on particle filters combined with scan matching. In *International Conference on Intelligent Robots and Systems*, 2012. 1
- [48] S. Rusinkiewicz and M. Levoy. Efficient variants of the ICP algorithm. In *International Conference on 3-D Digital Imaging and Modeling*, 2001. 2
- [49] A. Segal, D. Haehnel, and S. Thrun. Generalized-ICP. *Robotics: Science and Systems*, 2009. 2
- [50] S. Shi, Zheng You, K. Zhao, Z. Wang, C. Ouyang, and Y. Cao. A 6-dof navigation method based on iterative closest imaging point algorithm. *Scientific Reports*, 7, 2017. 1
- [51] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. A benchmark for the evaluation of RGB-D SLAM systems. *IEEE International Conference on Intelligent Robots and Systems*, 2012. 6
- [52] J. M. Tur. Onto computing the uncertainty for the odometry pose estimate of a mobile robot. *IEEE Conference on Emerging Technologies and Factory Automation (EFTA)*, pages 1340–1345, 2007. 1
- [53] D. Wang, Z. Tang, C. Bajaj, and Q. Liu. Stein variational gradient descent with matrix-valued kernels. In *NIPS*, 2019. 3
- [54] Y. Wang and J. Solomon. Deep closest point: Learning representations for point cloud registration. *ICCV*, 2019. 2
- [55] M. Welling and Y. W. Teh. Bayesian learning via stochastic gradient langevin dynamics. In *ICML*, 2011. 2
- [56] C. Zhang, J. Butepage, H. Kjellstrom, and S. Mandt. Advances in variational inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 2