

BO-ICP: Initialization of Iterative Closest Point Based on Bayesian Optimization

Harel Biggie, Andrew Beathard, and Christoffer Heckman*

Abstract—Typical algorithms for point cloud registration such as Iterative Closest Point (ICP) require a favorable initial transform estimate between two point clouds in order to perform a successful registration. State-of-the-art methods for choosing this starting condition rely on stochastic sampling or global optimization techniques such as branch and bound. In this work, we present a new method based on Bayesian optimization for finding the critical initial ICP transform. We provide three different configurations for our method which highlights the versatility of the algorithm to both find rapid results and refine them in situations where more runtime is available such as offline map building. Experiments are run on popular data sets and we show that our approach outperforms state-of-the-art methods when given similar computation time. Furthermore, it is compatible with other improvements to ICP, as it focuses solely on the selection of an initial transform, a starting point for all ICP-based methods.

I. INTRODUCTION

Point cloud registration or the process of estimating a transformation between two point clouds is paramount to modern robotic applications such as 3D scene reconstruction and decision-making [1]. The point cloud registration problem is easy to solve if the point correspondences are known. However, in practical applications, the point cloud registration problem is more difficult due to the unknown correspondence of points obtained from sensors such as lidar and stereo cameras. Various registration algorithms to simultaneously estimate the transform and correspondence between two point clouds have been proposed over the last two decades, and Iterative Closest Point (ICP) based methods [2] have largely dominated the field. ICP-based algorithms aim to minimize the point-to-point correspondence distance between a fixed reference cloud R and a floating target cloud T using an initial transform T_0 which in turn provides an estimate of the optimal transform between the two clouds. The results of ICP-based algorithms are highly dependent on the initial transform choice T_0 .

At its core, the ICP algorithm works as follows:

- 1) using an initial transform estimate T_0 , transform the target cloud T into the coordinate frame of reference cloud R ;
- 2) for each point in target cloud T find the closest corresponding point in the reference cloud R ;

This work was supported through the DARPA Subterranean Challenge cooperative agreement HR0011-18-2-0043, the National Science Foundation #1764092, and USDA-NIFA #2021-67021-33450.

*All authors are with the Autonomous Robotics and Perception Group at the University of Colorado Boulder, Boulder Colorado, USA. Corresponding author: christoffer.heckman@colorado.edu




Fig. 1: A visualization showing a 2D projection of the optimization in BO-ICP on the XY plane. Two different iteration counts of BO-ICP are shown: 10 iterations (top) and 60 iterations (bottom). In each example, a 3D surrogate function is shown below a color map of the acquisition function where red indicates areas of least utility and blue indicates areas of most utility. Squares highlight key alignment features in the scene.

- 3) using an objective (point-to-point distance or similar) find the best transform T_i between T and R ; and
- 4) repeat the process until convergence using T_i as the new T_0 .

Despite being a highly generalized algorithm, hundreds of variations of the ICP framework exist, e.g, matching points to planes [3] and colors [4]. Each of these algorithms was tailored to improve results for data available in specific instances such as RGB information. While many of these algorithms are capable of producing exceptional results given an initial transform estimate T_0 , they are prone to falling into local minima. A poor initial estimate can lead to vastly different results because ICP is an inherently non-convex problem that terminates when the point-to-point correspondence distance stops improving. Finding a globally optimal solution to the ICP problem is an ongoing discipline of research in both the computer vision and robotics communities.

A wide variety of methods have been proposed to either increase the accuracy of T_0 through external instrumentation [1] or the use of feature matching and alignment [5]. Other methods have utilized global optimization techniques such as branch and bound to determine the optimal transform in certain situations such as scan to model matching [6]. In situations where outliers, dynamic obstacles, or noisy sensor measurements exist, these global methods cannot guarantee an optimal solution [7]. The lack of consistently reliable methods for global registration motivates research targeted at more robust initialization techniques.




Fig. 2: Example alignment of source cloud 000739 (orange) to target cloud 000740 (blue) of sequence 10 from KITTI, with key features highlighted in pink. Due to the inaccuracy of ground truth transforms, GO-ICP and BO-ICP alignment metrics are evaluated on mean point to point distance.

In this paper, we outline a framework based on Bayesian optimization (BO) [8], a global optimization method, to systematically compute T_0 , and find this approach produces more accurate alignments than state-of-the-art methods. It is fundamentally compatible with variants of ICP that address point correspondences, the weighting of correspondences, and other methods that focus on adjusting how the objective is constructed.

Bayesian optimization is a global optimization scheme that has been commonly used to optimize neural network training parameters [9]. In practice, BO aims to find the minima or maxima of an *objective function* and is effective in situations where the objective is either complex, noisy, or expensive to calculate, such as in the case of minimizing the point-to-point correspondence distance. BO utilizes a computationally efficient probabilistic model of the objective function (see Figure 1) that is inexpensive to evaluate. As the point cloud registration problem as formulated by ICP is inherently a non-convex problem that requires expensive iterations for each initial T_0 estimate, it is an ideal candidate for BO. We demonstrate that our approach outperforms exhaustive searches for this initial estimate, as well as other “globally optimal” methods.

In this work, we present an open-source Bayesian optimization-based method (BO-ICP) for determining the crucial initial condition for ICP problems. This method optimizes the initial condition T_0 for point cloud registration in a principled manner that is grounded in statistics and is shown to outperform state-of-the-art methods both computationally and in accuracy. Our proposed method is not dependent on prior contextual knowledge. A common approach in such scenarios is to apply random guesses to the initial conditions; we refine this approach through a principled estimation of the objective function, the powerful engine that drives BO. Furthermore, our method generalizes to many of the variants of ICP [1] as it does not alter the internal operation of the ICP algorithm itself. Finally, our approach is straightforward to integrate into existing point cloud alignment systems, as almost all of them require some initial guess to be provided; our technique provides a principled technique for adjusting this guess.

II. RELATED WORKS

An abundance of high-quality prior work has been produced on point cloud registration. As our method directly contributes to initialization procedures for ICP and global

point cloud registration, we will focus our review on this facet of the field and refer the reader to the following surveys for a more broad view of the field [1], [10], [11].

Initial Transform Estimation: A wide variety of methods generally falling into three categories have been developed for estimating the required T_0 in ICP methods [1]. *External* based methods rely on obtaining the estimate from an outside source such as a human in the loop [12], Inertial Measurement Unit (IMU), or GPS [13], [14], [15]. Relying on external sensors limits the application space of these registration methods and prevents the generalization of various solutions. *Parameter cascade* based methods generally modify parameters such as the resolution of the point cloud through multiple iterations in order to achieve faster results. These methods commonly operate on a coarse-to-fine strategy where computationally results are obtained by down-sampling the point cloud and later refined with a better initial guess [16] or by sectioning off the point cloud into local regions [17]. *Registration Cascade* based methods are a class of initialization methods that rely on a separate external process to estimate the initial transformation. These approaches generally use histograms [18], [19] or keyframe-based approaches [20] to determine an initial T_0 .

Optimization Based Approaches: Other approaches have attempted to use optimization techniques such as the Levenberg–Marquardt algorithm [21] to perform point cloud registration [22]. The authors found that in certain situations the algorithm converged faster than traditional ICP based methods. Gaussian mixture models (GMMs) have also widely been explored to locally refine the initial transform [23], [24]. While these methods are capable of optimizing the estimate locally, they still require an initialization for T_0 .

Global Methods: More recent methods have used global optimization techniques such as branch and bound to solve the 3D registration problem [25], [6], [26]. A mathematical guarantee of convergence to the optimal solution was presented in GO-ICP [6] using a nested branch and bound approach. GO-ICP computes inner and outer bounds for both the $\text{SO}(3)$ rotation and $\text{SE}(3)$ translation spaces using an uncertainty radius approach. The method works remarkably well for model to scan alignments and small object to scene matching. However, it requires that both input clouds are scaled into the same unit cube and is not robust to low overlap conditions [7]. [27] presents a RANSAC-based geometric feature extraction method that is used to compute the correspondence inside of a fast point matching




Fig. 3: Statistical summaries for the point to point distance in KITTI (a), translational error (b), rotational error (c) in TUM. The vertical box edges signify the first and third quartiles; the whiskers extend from the box edges to the minimum/maximum value in either direction, up to a max distance of $1.5 \times$ the interquartile range, outliers are not shown. Means are shown in red and medians are shown in blue. GT represents ground truth for the KITTI datasets and A - E are the different configurations for BO-ICP.

histogram. This method has the potential for long run times due to the need to evaluate a potential infinite number of feature extractions. Fast Global Registration [5] overcomes some of these limitations using a combination of Fast Point Feature Histograms and an efficient optimization approach based on Black-Rangarajan duality [28]. Algorithms based on 4PCS [29], [30] perform global registration by utilizing intelligent selections of congruent points. More recently, [31] presents a truncated least squares approach combined with a maximal clique inlier selection procedure to solve the global registration problem in a certifiable manner that is robust to outliers.

III. METHODOLOGY

Consider the problem of aligning a target point cloud S to a reference cloud R through a rigid geometric transformation $T \in \mathbb{SE}(3)$, which we parametrize by $[x, y, z, \theta, \psi, \phi]$, with θ , ψ , and ϕ representing roll, pitch, and yaw respectively. In ICP, this is accomplished through the minimization of an *objective function* which is calculated over the two point sets; see step three of the ICP algorithm described in Section I. A challenge here is that the initial condition to ICP T_0 is arbitrary, and has a significant impact on the convergence and results of the final alignment. In order to address this, we develop a technique here that principally identifies T_g , a global transform for use as T_0 in ICP, using BO a global optimization technique. Note that BO can fundamentally be applied to any objective function used inside of ICP, e.g., point-to-point [2], point-to-plane [3], color [4]; we focus on its application towards optimizing the T_0 for the point-to-point ICP [2] for this discussion.

Point-To-Point ICP: For the point-to-point formulation of ICP, let S represent the source cloud that is being aligned to the reference cloud R using a transformation T . The algorithm works by finding the set of corresponding points K between each of the point clouds S and R transformed by T , $K = \{(s, r)\}$. Next, the algorithm aims to minimize an objective function representing the point-to-point correspondence distance between the points (Equation 1).

$$E(T_i) = \sum_{(s, r)} \|r - T_i s\|^2 \quad (1)$$

Algorithm 1: BO-ICP: For full 6DOF search. The algorithm can be done with rotations $[\theta, \psi, \phi]$ and then translations $[x, y, z]$. Where the identity is used as the initial translation and the optimized rotation is used to find the optimized translation.

Data: Source cloud S , Reference cloud R , bounds b , Number of random steps N_r , Number of iterations N

Result: Transformation between the two clouds T_f

Populate $\mu(T)$ for N_r random T_i generated from

$[x_i, y_i, z_i, \theta_i, \psi_i, \phi_i]$ within bounds b ;

while $i < N$ **do**

$T_i \leftarrow \text{max utility } [x_i, y_i, z_i, \theta_i, \psi_i, \phi_i] \text{ from}$

$EI(x)$;

$E(T_i) \leftarrow \text{Point To Point ICP}(T_i, S, R)$;

Update $\mu(T_i)$ with $E(T_i)$;

end

$T_g \leftarrow \text{at min } E(T_i)$;

The process is repeated with T_i becoming the newly minimized transform obtained from each iteration i . This iterative process is highly susceptible to converging to a local minimum which can result in erroneous alignments, especially with a poor initial alignment. However, trying all possible alignments is largely considered an NP-hard problem [25] and which motivates the need for a more computationally efficient approach.

Bayesian Optimization: The crux of BO is an efficient approximation based on Bayes Rule of the objective function; this is known as the *surrogate function* which we call $\mu(T)$. Next, an *acquisition function* is used to select the next best sample from the posterior search space modeled by the $\mu(T)$. Performing global optimization on ICP based methods is an ideal use case for BO due to the inherent non-convex nature of the functions and the expensive computations required at each iteration. By estimating the ICP function using a *surrogate function* BO is able to efficiently choose the next best guess to minimize the uncertainty based on the *acquisition function*, in comparison to evaluating $E(T_i)$ at each iteration.

Typically the *surrogate function* is modeled using Gaussian processes which can be used to directly infer a distribution over a “black box” function. Gaussian processes use




Fig. 4: Example alignment of source cloud timestamp 1311868250.071665 (orange) to target cloud timestamp 1311868250.142995 (blue) of sequence fr2/desk of the TUM odometry dataset. Pink squares highlight key alignment features.

prior observations to learn a posterior over a given function. We make an assumption here that the ICP function can be modeled as a random variable distributed over a Gaussian.

Once the *surrogate function* is modeled one can evaluate the utility of sampling at a new point using a variety of probabilistic models. For the purposes of BO-ICP we use the expected improvement (EI) metric. Let y^* represent the minimum value observed from the objective function, x represent the modeled parameter, and y represent the possible new observations. Then $p(y|x)$ is representative of the surrogate model which is modeled as a Gaussian. The EI can then be written as:

$$EI_{y^*}(x) = \int_{-\infty}^{\infty} \max(y^* - y, 0) p(y|x) dy \quad (2)$$

BO-ICP can be run in two variations, a full six degrees of freedom (6DOF) search in T (shown in Algorithm 1), and a more practical nested approach for computational efficiency. In all cases, the bounds b of the problem should be chosen based on the range of the sensor being used and the size of the scenes. As shown in [6] it is possible to first search for the optimal rotation parameters $([\theta, \psi, \phi])$ and then the translation parameters $[x, y, z]$. While we note BO does not share the same theoretical guarantee behind this nested approach we find that in practice we are able to leverage this technique to help reduce the number of hyperparameters required by the *surrogate function*. In the nested variant, we first run BO on only the rotation parameters, and then we pass the computed rotation that minimizes the objective through BO again to find the optimal translation. For both the translation and rotation searches we sample a pseudo-random set of points to build an initial estimate of the *surrogate function*. We note that generally, sampling in $[\theta, \psi, \phi]$ is susceptible to gimbal lock which we overcome by centering our axis around the sample and applying a change of basis. Sampling in terms of quaternions was also explored but the conversion of sampled points to a uniform quaternion as described in [32] breaks the normal assumption present in the GP. In practice, this is not strictly required but greatly aids in the ability of the *acquisition function* to make initially informed guesses.

Determining the termination criteria for BO-ICP is difficult in practice without knowing any prior information about the sensor. ICP will locally terminate when the point-to-point distance becomes stationary. Given enough iterations, BO will also reach a point where the objective stops improving

but in practice, the optimal value will be found long before this occurs. As such, we select T_0 which best minimizes $E(T)$ after running the algorithm for a set number of iterations N . Additionally, as described in Section IV we conduct experiments to find values for N which work well in practice.

IV. EXPERIMENTS

We compare the nested variant of BO-ICP against two other global methods: a brute-force pyramid-based grid search (PY) and GO-ICP [6] which, to the best of our knowledge, is the most successful global ICP variant that does not rely on selecting informative geometries or feature point extraction. Experiments are conducted from the KITTI Odometry dataset [33], [34] and the TUM Vision Handheld SLAM RGB-D dataset [35] which enables us to test the robustness of our method to realistic noise. The KITTI dataset was chosen to highlight the effectiveness of our method in the presence of fast-moving scenes and dynamic obstacles. The TUM dataset was collected using a structured light sensor which allows us to evaluate our dataset on point clouds with more noise in comparison to the lidar-based clouds of the KITTI dataset. Unlike the inaccurate GPS-based ground truth in the KITTI dataset (see Section V for more), the TUM dataset provides ground truth obtained from a motion capture system which allows us to conduct an accurate analysis of the absolute error in transformations.

Pyramid-based searches are one way of addressing local minima in ICP. In our comparisons, we follow a nested approach where we first search for rotation and then translation. An initial grid of size of $12 \times 12 \times 12$ (1728 samples) is populated using a radius-based nearest neighbor search to estimate the fit of the initial transform. New grids of size $6 \times 6 \times 6$ are then created from the top 10 results using a fixed delta value from the parameter and the process is repeated for a total of 3 times. At completion, ICP is performed on the top 10 values and the top result is selected. The resulting rotation values from the first search are used to find the optimal translational component using the same procedure.

Intervals of testing data are selected from the KITTI and TUM datasets using a pseudo-random initialization procedure. From the selected intervals, reference R and target T point clouds are paired together for testing using a conservative cloud overlap estimation from [7]. We use this metric to ensure there is sufficient overlap between the two clouds (70% on KITTI and 60% on TUM) for ICP to

be performed. Once that threshold is surpassed, the target cloud T becomes the new reference cloud R . This method produces 248 testing pairs from the KITTI dataset and 325 testing pairs from the TUM dataset on which we analyze the efficacy of the proposed method.

Config	Random Samples	Iterations	Voxel Size
A	10	20	0.7
B	20	35	0.7
C	30	60	0.6

TABLE I: Summary of the testing configurations for BO-ICP. Configuration A is designed to emphasize the ability of the algorithm to quickly find reasonable alignments while B serves as a direct comparison to GO-ICP, and C shows that BO-ICP is capable of further refining measurements for map building.

For each alignment pair, we select r pseudo-random initialization points and then run Bayesian Optimization for N iterations to find the best rotation estimate. We select 3 different configurations for BO-ICP which highlight the algorithm’s versatility:

- Configuration A illustrates the potential of the method to run as part of a loop closure mechanism in the context of SLAM or when realigning maps in the online navigational setting.
- Configuration B serves as a comparison between BO-ICP, GO-ICP, and the pyramid algorithm when allowed to run for similar durations.
- Configuration C emphasizes BO-ICP’s ability to continuously improve estimates when given more computation time which can be utilized in batch processing large offline maps.

Parameters were chosen to show trends between the different methodologies rather than directly comparing their implementations. We note that ICP implementation in GO-ICP is single-threaded while the other methods are multi-threaded. However, on profiling, the code the bulk of the runtime occurs in the generation of the distance transformation required for the algorithm. For each alignment pair, we optimize over the full rotation space and the following bounds for translation: $x \in [-4, 4]$, $y \in [-2, 2]$, $z \in [-1, 1]$. The bounds were selected based on the velocity of the vehicle in the KITTI dataset.

Summaries of the important parameters for each configuration can be found in Table I including the voxel-based down-sampling approach used for efficiency. We note that the voxel size and the number of iterations have the most significant impact on the performance of BO-ICP. Other parameters required minimal adjustments and the full parameter set as well as our open source implementation can be found at https://github.com/arpg/bo_icp. The provided implementation also provides configurations for the full 6DOF variant of BO-ICP. However, experiments for this variant showed that the full version of BO-ICP quadrupled the runtime without any noticeable improvement in results.

Our Python implementation leverages the Open3D Point Cloud Library [36] and an open source Bayesian optimization package [37]. Our implementation utilizes multi-

threading for ICP but not in the BO process. We compare our implementation to the C++ implementation of GO-ICP provided by the authors [6] using the default configuration except for a trim fraction of 0.3 to prevent failures from outliers in the data [38]. All presented experimental results were run on an Intel i7-8850H CPU with 32GB of RAM.

V. RESULTS

	Runtime				Error	
	KITTI		TUM		Trans (m)	Rot (rad)
	μ (s)	σ (s)	μ (s)	σ (s)		
A	11.3	1.3	12.3	0.8	0.42	0.033
B	22.3	2.0	22.5	1.7	0.38	0.028
C	43.1	2.3	40.0	3.7	0.36	0.028
GO	25.1	0.6	26.0	0.9	0.37	0.074
PY	28.4	6	11.4	8.0	0.88	0.039

TABLE II: Runtime comparisons for all methods are shown for both the KITTI and TUM datasets along with the mean rotational and translational errors on the TUM dataset.

Figure 2 provides a representative view of a random scene pairing from the KITTI dataset. It is important to note that the KITTI ground truth data is obtained by GPS and IMU measurements and is therefore not sufficiently accurate for rotational and translational alignment analysis¹. This misalignment is highlighted in Figure 2a. We instead provide a statistical distribution in Figure 3a and of the mean point-to-point distance of sample pairs as a metric to evaluate the quality of registration.

We observe that BO-ICP configuration B outperforms GO-ICP and the pyramid approach both with respect to runtime (Table II) and performance and alignment precision. In Figure 3a, we compare mean point-to-point distances of all BO-ICP test configurations as described in Table I against the pyramid approach, GO-ICP, and dataset-provided ground truth. Figure 3a shows that B and C have the highest percentage of mean point-to-point readings within 0.5m, while GO-ICP and BO-ICP A have a high frequency of outliers measuring greater than 2m point to point.




Fig. 5: Histogram showing the translational error from the TUM dataset

Additionally, we conduct a one-way ANOVA [39] test on BO-ICP configuration B (22.3s runtime which is most similar to GO-ICP) versus GO-ICP which produces an f -value of 13.4 and a p -value of 1.9×10^{-6} , confirming the significance of these results.

¹http://www.cvlibs.net/datasets/kitti/eval_odometry.php

Quantitative results for the TUM dataset in the form of translational and rotational errors are summarized in Figures 3b and 3c. Unlike KITTI, the TUM dataset was captured with a high tolerance ground truth trajectory taken from a motion capture system. Thus we are able to compare GO-ICP and BO-ICP transformation values directly against a known accurate transform.

Rotational and translational alignment errors of all BO-ICP test configurations are weighed against the pyramid approach and GO-ICP in Figures 3b, 3c, 5, and 6. From these results, we conclude that configurations *A*, *B*, and *C* on average (mean) outperform GO-ICP in translational error and all configurations of BO-ICP outperform the other methods in rotational error. In Figures 3b and 3c we observe that GO-ICP has a slightly lower median than all configurations of BO-ICP for translational error but BO-ICP has lower medians in rotational error. The brute force-based pyramid approach is the least accurate of all methods. BO-ICP configuration *B* performed exceptionally well in rotational alignment, with a significantly lower mean and quartile range than GO-ICP while running in a similar amount of time as shown in Table II.

For clarity, in Figures 5 and 6 we have truncated measurements beyond a certain rotational and translational error and consolidated all removed outliers into the maximum bin. These figures demonstrate that *B* and *C* have the highest percentage of readings within 0.1 radians of rotational error and that GO-ICP has a considerable accumulation of maximum-level outliers for rotation. From these figures, we also see that configuration *C* is able to reduce the spread of outliers in comparison to all other methods.




Fig. 6: Histogram showing the translational error on the TUM dataset.

VI. DISCUSSION

From the statistical results presented in Figures 3a, 3b, 3c it is clear that both tests *B* and *C* of BO-ICP outperform GO-ICP on average in terms of accuracy. Perhaps more remarkably, in test *A* BO-ICP provides similar alignments on average in 50% less runtime than GO-ICP, as seen in Table II. Test *A* highlights the ability of BO to quickly and effectively approximate the true distribution of the point-to-point ICP objective function. Meanwhile tests *B* and *C* emphasize that BO-ICP is able to further refine its previously more accurate

estimates given the ability to run for a longer duration in some instances.

Taking a closer look at the distributions of alignment errors in Figures 5 and 6 it becomes apparent that BO-ICP is much more adept at handling outliers in comparison to GO-ICP. For instance, BO-ICP is able to find initial transformations that align key features in the scene rather than large flat surfaces such as the ground. This is highlighted in Figure 2 where we see GO-ICP has a more refined alignment of the ground but BO-ICP aligned the post in the scene. For the TUM dataset specifically, most of the failures for both BO-ICP and GO-ICP occurred on the floor sequence. This sequence presents a challenge due to the lack of clear distinguishing features in the geometry. As with any ICP method, we found that tuning the voxel size parameter can significantly aid in convergence when there are limited distinguishing geometric features. We note that the testing configuration for GO-ICP used Trimmed ICP [38] to prevent outright failures in alignment from data outliers. Additional tests were conducted where GO-ICP did not use trimmed ICP but in many cases, as noted by the authors the algorithm outright failed and as such, we found the results did not provide a meaningful comparison. Contrast, the probabilistic nature of BO-ICP makes it more robust to individual outliers in the point cloud.

From Figure 3a it is clear that test *C* provides both the lowest median error as well as the smallest distribution spread. In contrast to GO-ICP which relies on an underlying uncertainty radius assumption to compute the search bounds, BO-ICP exploits a probabilistic model that requires significantly less assumptions to model and therefore is more robust to dynamic obstacles and outliers. While this does prevent BO-ICP from having a clear termination criterion, it is evident that practically, the algorithm works better than state of the art methods which do claim optimality under specific assumptions. Additionally, BO-ICP is probabilistically grounded which provides a confidence metric for the alignment that can be used to estimate the overall quality of the registration.

VII. CONCLUSIONS AND FUTURE WORK

We demonstrated that BO-ICP is able to solve the global point cloud registration problem both more accurately and efficiently than state-of-the-art methods that do not rely on feature extraction. Additionally, the underlying BO process is resilient to noisy sensors and moving obstacles and is capable of reasoning over 6DOF without a nested approach. Potential future work includes building an online mapping framework using BO-ICP. The random initialization component of the algorithm lends well to seeding the optimization using a priori knowledge such as the motion of a robot as the initial sample points. We hypothesize that using such a seed will only further speed up the convergence time for BO-ICP. Furthermore, it is noteworthy that this addition for identifying good initializations is compatible with many modern improvements to ICP as it only alters the initialization step.

REFERENCES

- [1] F. Pomerleau, F. Colas, and R. Siegwart, "A review of point cloud registration algorithms for mobile robotics," *Foundations and Trends in Robotics*, vol. 4, no. 1, pp. 1–104, 2015.
- [2] P. J. Besl and N. D. McKay, "Method for registration of 3-d shapes," in *Sensor fusion IV: control paradigms and data structures*, vol. 1611, pp. 586–606, Spie, 1992.
- [3] Y. Chen and G. Medioni, "Object modelling by registration of multiple range images," *Image and vision computing*, vol. 10, no. 3, pp. 145–155, 1992.
- [4] J. Park, Q.-Y. Zhou, and V. Koltun, "Colored point cloud registration revisited," in *Proceedings of the IEEE international conference on computer vision*, pp. 143–152, 2017.
- [5] Q.-Y. Zhou, J. Park, and V. Koltun, "Fast global registration," in *European conference on computer vision*, pp. 766–782, Springer, 2016.
- [6] J. Yang, H. Li, and Y. Jia, "Go-icp: Solving 3d registration efficiently and globally optimally," in *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1457–1464, 2013.
- [7] J. Stechschulte, N. Ahmed, and C. Heckman, "Robust low-overlap 3-d point cloud registration for outlier rejection," in *2019 International Conference on Robotics and Automation (ICRA)*, pp. 7143–7149, IEEE, 2019.
- [8] J. Mockus, V. Tiesis, and A. Zilinskas, "The application of bayesian methods for seeking the extremum," *Towards global optimization*, vol. 2, no. 117-129, p. 2, 1978.
- [9] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," *Advances in neural information processing systems*, vol. 25, 2012.
- [10] G. K. Tam, Z.-Q. Cheng, Y.-K. Lai, F. C. Langbein, Y. Liu, D. Marshall, R. R. Martin, X.-F. Sun, and P. L. Rosin, "Registration of 3d point clouds and meshes: A survey from rigid to nonrigid," *IEEE transactions on visualization and computer graphics*, vol. 19, no. 7, pp. 1199–1217, 2012.
- [11] F. Pomerleau, F. Colas, R. Siegwart, and S. Magnenat, "Comparing icp variants on real-world data sets," *Autonomous Robots*, vol. 34, no. 3, pp. 133–148, 2013.
- [12] G. Godin, M. Rioux, and R. Baribeau, "Three-dimensional registration using range and intensity information," in *Videometrics III*, vol. 2350, pp. 279–290, International Society for Optics and Photonics, 1994.
- [13] J. Diebel, K. Reutersward, S. Thrun, J. Davis, and R. Gupta, "Simultaneous localization and mapping with active stereo vision," in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566)*, vol. 4, pp. 3436–3443, IEEE, 2004.
- [14] S. Druon, M.-J. Aldon, and A. Crosnier, "Color constrained icp for registration of large unstructured 3d color data sets," in *2006 IEEE International Conference on Information Acquisition*, pp. 249–255, IEEE, 2006.
- [15] Y. Gao, S. Liu, M. M. Atia, and A. Noureldin, "Ins/gps/lidar integrated navigation system for urban and indoor environments using hybrid scan matching algorithm," *Sensors*, vol. 15, no. 9, pp. 23286–23302, 2015.
- [16] T. Jost and H. Hugli, "A multi-resolution scheme icp algorithm for fast shape registration," in *Proceedings. First International Symposium on 3D Data Processing Visualization and Transmission*, pp. 540–543, IEEE, 2002.
- [17] M. Bosse and R. Zlot, "Continuous 3d scan-matching with a spinning 2d laser," in *2009 IEEE International Conference on Robotics and Automation*, pp. 4312–4319, IEEE, 2009.
- [18] 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*, vol. 27, no. 6, pp. 667–691, 2008.
- [19] R. B. Rusu, N. Blodow, and M. Beetz, "Fast point feature histograms (fpfh) for 3d registration," in *2009 IEEE international conference on robotics and automation*, pp. 3212–3217, IEEE, 2009.
- [20] C.-L. Tsai, C.-Y. Li, G. Yang, and K.-S. Lin, "The edge-driven dual-bootstrap iterative closest point algorithm for registration of multimodal fluorescein angiogram sequence," *IEEE transactions on medical imaging*, vol. 29, no. 3, pp. 636–649, 2009.
- [21] J. J. Moré, "The levenberg-marquardt algorithm: implementation and theory," in *Numerical analysis*, pp. 105–116, Springer, 1978.
- [22] A. W. Fitzgibbon, "Robust registration of 2d and 3d point sets," *Image and vision computing*, vol. 21, no. 13-14, pp. 1145–1153, 2003.
- [23] B. Jian and B. C. Vemuri, "A robust algorithm for point set registration using mixture of gaussians," in *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, vol. 2, pp. 1246–1251, IEEE, 2005.
- [24] 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*, pp. 4292–4300, 2015.
- [25] O. Enqvist, K. Josephson, and F. Kahl, "Optimal correspondences from pairwise constraints," in *2009 IEEE 12th international conference on computer vision*, pp. 1295–1302, IEEE, 2009.
- [26] J. Yang, H. Li, D. Campbell, and Y. Jia, "Go-icp: A globally optimal solution to 3d icp point-set registration," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 11, pp. 2241–2254, 2015.
- [27] S. Choi, Q.-Y. Zhou, and V. Koltun, "Robust reconstruction of indoor scenes," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5556–5565, 2015.
- [28] M. J. Black and A. Rangarajan, "On the unification of line processes, outlier rejection, and robust statistics with applications in early vision," *International journal of computer vision*, vol. 19, no. 1, pp. 57–91, 1996.
- [29] D. Aiger, N. J. Mitra, and D. Cohen-Or, "4-points congruent sets for robust pairwise surface registration," in *ACM SIGGRAPH 2008 papers*, pp. 1–10, ACM, 2008.
- [30] N. Mellado, D. Aiger, and N. J. Mitra, "Super 4pcs fast global pointcloud registration via smart indexing," in *Computer graphics forum*, vol. 33, 5, pp. 205–215, Wiley Online Library, 2014.
- [31] H. Yang, J. Shi, and L. Carlone, "Teaser: Fast and certifiable point cloud registration," *IEEE Transactions on Robotics*, vol. 37, no. 2, pp. 314–333, 2020.
- [32] K. Shoemake III, "6—uniform random rotations in: Kirk d, editor. graphics gems iii (ibm version)," 1992.
- [33] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [34] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *International Journal of Robotics Research (IJRR)*, 2013.
- [35] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of rgb-d slam systems," in *Proc. of the International Conference on Intelligent Robot Systems (IROS)*, Oct. 2012.
- [36] Q.-Y. Zhou, J. Park, and V. Koltun, "Open3D: A modern library for 3D data processing," *arXiv:1801.09847*, 2018.
- [37] F. Nogueira, "Bayesian Optimization: Open source constrained global optimization tool for Python," 2014–.
- [38] D. Chetverikov, D. Svirko, D. Stepanov, and P. Krsek, "The trimmed iterative closest point algorithm," in *Object recognition supported by user interaction for service robots*, vol. 3, pp. 545–548, IEEE, 2002.
- [39] B. L. Welch, "The generalization of 'student's' problem when several different population variances are involved," *Biometrika*, vol. 34, no. 1-2, pp. 28–35, 1947.