Quatro++: Robust Global Registration Exploiting Ground Segmentation for Loop Closing in LiDAR SLAM

Journal Title XX(X):1-26 @The Author(s) 2022 Reprints and permission: sagepub.co.uk/journalsPermissions.nav DOI: 10.1177/ToBeAssigned www.sagepub.com/ SAGE

Hyungtae Lim¹, Beomsoo Kim², Daebeom Kim³, Eungchang Mason Lee¹, and Hyun Myung^{1*}

Abstract

Global registration is a fundamental task that estimates the relative pose between two viewpoints of 3D point clouds. However, there are two issues that degrade the performance of global registration in LiDAR SLAM: one is the sparsity issue and the other is degeneracy. The sparsity issue is caused by the sparse characteristics of the 3D point cloud measurements in a mechanically spinning LiDAR sensor. The degeneracy issue sometimes occurs because the outlier-rejection methods reject too many correspondences, leaving less than three inliers. These two issues have become more severe as the pose discrepancy between the two viewpoints of 3D point clouds becomes greater. To tackle these problems, we propose a robust global registration framework, called *Quatro++*. Extending our previous work that solely focused on the global registration itself, we address the robust global registration. Through the experiments, we demonstrate that our proposed method shows a higher success rate than the state-of-the-art global registration methods, overcoming the sparsity and degeneracy issues. In addition, we show that ground segmentation significantly helps to increase the success rate for the ground vehicles. Finally, we apply our proposed method to the loop closing module in LiDAR SLAM and confirm that the quality of the loop constraints is improved, showing more precise mapping results. Therefore, the experimental evidence corroborated the suitability of our method as an initial alignment in the loop closing. Our code is available at https://guatro-plusplus.github.io.

Keywords

LiDAR SLAM, Global Registration, 3D Point Cloud, Ground Segmentation

1. Introduction

3D point cloud registration is a fundamental task that estimates relative pose between two viewpoints of 3D point clouds, namely source and target, in robotics and computer vision fields. Using the characteristic of being able to estimate the relative pose between two point clouds, 3D point cloud registration is widely exploited in various applications such as relocalization (Du et al. 2020), ego-motion estimation (Koide et al. 2021), object recognition (Chua and Jarvis 1996; Ashbrook et al. 1998; Belongie et al. 2002; Cho and Gai 2014).

Furthermore, these point cloud registration methods are key components in light detection and ranging (LiDAR) sensor-based simultaneous localization and mapping (SLAM), which is the integral part of building a map while localizing a robot or an autonomous vehicle itself. In particular, point cloud registration is not only used in odometry, but also exploited in the loop closing procedure to obtain constraints for pose graph optimization (PGO) in graph-based SLAM.

In general, the graph-based SLAM mainly consists of three parts (Thrun and Montemerlo 2006; Grisetti et al. 2010; Qin et al. 2018; Shan et al. 2020): a) odometry, which estimates relative pose between the consecutive frames, b) loop detection, which achieves data association between two non-consecutive frames, and c) loop closing, which estimates relative pose between the non-consecutive frames searched by the loop detection module. The results through these three parts are represented in a graph structure, i.e. nodes and vertices. Then, the results are taken as inputs of PGO. The graph-based SLAM reduces global trajectory errors through this procedure, building a precise and globally consistent map.

In recent years, as the demand for LiDAR-based SLAM has increased, novel and accurate LiDAR-based SLAM frameworks have been proposed. However, these novel methods usually focus on improving the odometry or loop detection modules. Once the performance of loop detection is improved, the quality of constraints from the loop closing is also increased because the relation between loop detection and loop closing is not completely independent (Chen et al. 2022). However, the methods for increasing the quality of loop constraints in the loop closing step is still less actively studied.

Furthermore, many LiDAR-based SLAM methods only utilize iterative closest point (ICP) or its variants, which are

Corresponding author:

¹School of Electrical Engineering at KAIST (Korea Advanced Institute of Science and Technology), Daejeon, 34141, Republic of Korea

²Department of Artificial Intelligence, Hanyang University, Seoul, 04763, Republic of Korea

³Robotics Program at KAIST, Daejeon, 34141, Republic of Korea

^{*}Prof. Hyun Myung is with School of Electrical Engineering and KI-R at KAIST, Daejeon, 34141, Republic of Korea.

However, the local registration searches the point pairs by using nearest neighbor search, so the local registration-based pose estimation is only valid if two point clouds are mostly overlapped and the pose discrepancy between the two point clouds is small (empirically, within 2 m and 10° (Kim et al. 2019; Lim et al. 2020)). Otherwise, the estimates of the local registration are highly likely get stuck in the local minima. For this reason, even though loop detection finds the correct loop candidate, the loop candidate with a large pose discrepancy can be rejected owing to the narrow convergence region. Consequently, these false negative loops potentially hinder more accurate PGO results.

To tackle this problem, global registration methods, whose performance is less affected by the initial pose difference between the two point clouds, can be utilized as a coarse alignment. That is, the initial guess using the global registration reduces the pose difference in a coarse manner. This reduction of the pose discrepancy helps local registration methods converge into a global optimum. As a result, the loop candidates with large pose discrepancies can also be exploited as precise loop constraints.

However, there are still two issues that degrade the performance of the global registration in 3D LiDAR SLAM: one is the sparsity issue (Lim et al. 2021b) and the other is degeneracy (Lim et al. 2022). First, the sparsity issue occurs in 3D point cloud measurements when using a mechanically spinning LiDAR sensor: the farther the range from the origin goes, the dramatically more sparse the density of the 3D point cloud becomes. This decreasing density degrades the expressibility of the feature descriptors. That is, even if the same place is observed, the descriptor has a different value due to the density difference between the two point clouds. Consequently, the sparsity issue induces lots of false matching, increasing the ratio of the outliers within the estimated correspondences.

Second, degeneracy sometimes occurs owing to the effect of outlier rejection algorithms. Here, degeneracy refers to the case when fewer inliers remain than the degree of freedom (DoF). The degeneracy is occasionally observed when the well-known outlier rejection methods, such as maximum clique inlier selection (MCIS) (Shi et al. 2021b; Sun 2021b) or weight update in graduated nonconvexity (GNC) (Zhou et al. 2016; Yang et al. 2020b), are employed. These methods unintentionally prune too many estimated correspondences because these methods usually aim to achieve high recall of outlier rejection; thus, sometimes less than three inliers are left after the pruning step. As a result, the degeneracy triggers tilted or flipped results (see Section 7.3). These two issues become more severe as the pose difference between the two viewpoints of 3D point clouds becomes greater.

In our earlier work, i.e. Quatro (Lim et al. 2022), we demonstrated that Quatro overcomes the aforementioned problems by leveraging the decoupling-based method that reduces the minimum number of required correspondences for the estimation from three to one. However, Quatro was still not exploited as a loop closing module in SLAM; thus, the influence of robust global registration on the performance of graph-based SLAM was still not closely examined.

Therefore, we propose a robust global registration framework, called Quatro++, and analyze the effect of the global registration from the loop closing-centric perspective. As shown in Fig. 1, Quatro++ is a superset of Quatro that exploits ground segmentation to be more robust against outlier pairs and degeneracy issues. In Quatro++, we introduce ground segmentation to allow global registration to be more robust against distant cases by filtering ground points, which are most likely to be featureless, before the pose estimation. Even though many researchers empirically use ground segmentation as a preprocessing step (Kim and Kim 2018; Yue et al. 2019; Komorowski 2021), the effect of ground segmentation on the performance of global registration has seldom been quantitatively analyzed. Thus, we validate that ground segmentation substantially increases the success rate of global registration even in degenerate environments, such as corridor-like scenes, or distant cases. Finally, as an extension of our previous work (Lim et al. 2022), we additionally compare our approach with stateof-the-art learning-based approaches (Shi et al. 2021a; Cattaneo et al. 2022) and show that our approach achieves substantially higher performance than other approaches in loop closing level.

In summary, our ultimate goal is to achieve robust global registration when spurious correspondences caused by the degradation of feature extraction and matching are given. Unlike other studies that focus more on improving the expressibility of feature descriptors to increase the quality of correspondences (Ao et al. 2021; Shi et al. 2021a; Chen et al. 2022; Cattaneo et al. 2022; Yin et al. 2023), what we aim for is robustly estimating the relative pose even though undesirable imprecise correspondences are given. The contribution of this paper is fourfold:

- We propose a novel global registration framework, Quatro++, that robustly achieves an initial alignment, overcoming the effect of the gross outliers and degeneracy issue as well.
- We demonstrate that employing ground segmentation makes global registration more robust and reduces the computational cost, allowing more accurate and fast loop closing for LiDAR SLAM.
- · Our proposed method was analyzed from various perspectives, showing a superior performance compared with state-of-the-art methods, including deep learningbased approaches. By doing so, we demonstrate the ease of exploiting our Quatro++ because our method is a learning-free approach.
- By integrating our approach with various LiDAR SLAM frameworks, we showed that our Quatro++ successfully improves the performance of SLAM by providing more precise loop constraints as an initial alignment.

2. Related Works

2.1. Local Registration

As mentioned earlier, point cloud registration is mainly classified into two categories depending on how to find correspondences between the two point clouds: one is the local registration (Besl and McKay 1992; Rusinkiewicz and



Fig. 1. Overview of our global registration method, called *Quatro++*. As an extension of our previous work (Lim et al. 2022), Quatro++ consists of five parts. The red, green, and blue lines of the raw and pruned correspondences (shown inside dotted boxes), i.e. A_{raw} and A, respectively, denote the outliers, true inliers, and quasi-inliers (see Section 4.7), respectively. (a) The preprocessing step using ground segmentation (gray points denote the estimated ground points). (b) Correspondence estimation by using descriptor extraction and matching, which can be interchangeable with other methods. (c) Graph-based outlier pruning to initially reject outlier correspondences. (d) Quasi-SO(3) estimation based on graduated non-convexity (GNC). (e) Component-wise translation estimation. (best viewed in color).

Levoy 2001; Chetverikov et al. 2002; Segal et al. 2009; Pomerleau et al. 2013; Koide et al. 2021; Dellenbach et al. 2022) and the other is the global registration (Fischler and Bolles 1981; Dong et al. 2017; Yang et al. 2015; Zhou et al. 2016; Yang et al. 2020b; Bernreiter et al. 2021).

First, the local registration methods heavily rely on the nearest neighbor search (Greenspan and Godin 2001), which assumes that a target point correlates with the source point closest to the target point. ICP (Besl and McKay 1992) and its subsequent studies are renowned local registration methods. Unfortunately, the assumption above is invalid once two point clouds are far apart and slightly overlapped. For this reason, ICP variants in loop closing situations often fail to estimate relative pose when the pose discrepancy between the two point clouds is large.

2.2. Global Registration

Unlike these local registration methods, global registration is more likely to be appropriate to estimate the relative pose in the case where the pose discrepancy is large. Accordingly, the output of the global registration can be used as an initial guess, allowing the estimate of the local registration to successfully converge into the global minimum. Note that the two types of global registration methods exist: a) correspondence-based (Fischler and Bolles 1981; Yang et al. 2015; Zhou et al. 2016; Dong et al. 2017; Tzoumas et al. 2019; Yang et al. 2020b; Lin et al. 2022) and b) correspondence-free methods (Rouhani and Sappa 2011; Brown et al. 2019; Bernreiter et al. 2021). In this study, we place more emphasis on the correspondence-based methods.

2.3. Correspondence-Based Global Registration

The correspondence-based methods utilize feature extraction and matching to obtain correspondences between the two point clouds. However, the outliers inevitably occur within the putative correspondences. To tolerate the effect of outliers on pose estimation, numerous researchers have studied outlier-robust global registration methods. Typical examples are random sample consensus (RANSAC) proposed by Fischler and Bolles (1981), and its variants (Papazov et al. 2012; Chum et al. 2003; Choi et al. 1997; Schnabel et al. 2007). However, the RANSAC-based methods are likely to be brittle with high outlier rates. Empirically, these methods often fail if the ratio of outliers is over 50 % (Tzoumas et al. 2019).

The other examples are branch-and-bound (BnB)-based methods (Lawler and Wood 1966; Olsson et al. 2008; Hartley and Kahl 2009; Pan et al. 2019). These methods have an advantage in that they guarantee theoretical optimality. Unfortunately, these BnB-based methods are too slow to be exploited in real-world applications, e.g. it takes over 50 sec for a single registration (Lei et al. 2017).

2.4. Fast and Outlier-Robust Global Registration

Unlike the RANSAC-variants or BnB-based methods, GNC has been introduced to achieve outlier-robust and fast registrations. The GNC-based methods estimate pose while rejecting the outlier measurements simultaneously (Black and Rangarajan 1996; Zhou et al. 2016), overcoming up to 70-80% of outliers and providing faster speed compared with the previous works. Consequently, GNC has been widely used to robustly estimate the relative pose against the gross outliers (Tzoumas et al. 2019; Yang et al. 2020a,b; Sun 2021a; Song et al. 2022).

In addition, the certifiable methods that guarantee their optimality in polynomial-time have been proposed, such as semidefinite programming and sums of squares relaxationbased methods (Vandenberghe and Boyd 1996; Briales and Gonzalez-Jimenez 2017; Maron et al. 2016; Carlone et al. 2016).

2.5. Geometry-Aware Traditional Outlier Rejection Modules

Further, some researchers have studied novel traditional outlier rejection modules. To check the geometric consistency and thus prune spurious correspondences effectively, these outlier rejection methods leverage spectral techniques (Leordeanu and Hebert 2005), maximum clique (Shi et al. 2021b; Sun 2021b), voting schemes (Glent Buch et al. 2014; Yang et al. 2019), or game theory (Rodolà et al. 2013).

2.6. Deep Learning-Based Registration

As deep learning era has come, deep learning-based methods have shown remarkable achievements. The research directions can be categorized into three groups: a) improving the quality of feature extraction and matching (Zeng et al. 2017; Yew and Lee 2018; Deng et al. 2018; Gojcic et al. 2019), b) end-to-end registration (Wang and Solomon 2019; Gojcic et al. 2020), c) learning-based outlier rejection techniques within the deep learning architecture (Choy et al. 2020; Pais et al. 2020; Bai et al. 2021). Among them, the last group shares our philosophy that outliers are always inevitable.

However, these data-driven methods are too fitted into the situations included in the training dataset; thus, the learningbased methods can result in catastrophic failure if they encounter untrained situations or the data captured by other sensor configuration are taken as inputs. As we pursue a versatile method that is applicable in any environment, the learning-based methods are thus beyond our scope.

3. Preliminaries

In this section, we present the notation of variables and the problem definition of global registration. In addition, the challenges of global registration using sparse point clouds captured by a 3D LiDAR sensor are explained. In particular, we focus more on the mechanically spinning 3D LiDAR sensor because its uniform beam forming triggers a severe sparsity issue (Lim et al. 2021b), potentially causing more outlier correspondences.

3.1. Notation

First of all, source and target clouds are denoted as **P** and **Q**, respectively, and let us define their coordinate frames as \mathcal{P} and \mathcal{Q} , respectively. These point clouds consist of a number of cloud points, which can be expressed as $\mathbf{P} = {\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_N}$ and $\mathbf{Q} = {\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_M}$. Each point is expressed in the Cartesian coordinates, so the elements of \mathbf{p}_i $(1 \le i \le N)$ and \mathbf{q}_j $(1 \le j \le M)$ are comprised of x, y, and z values, respectively.

Next, we aim for correspondence-based global registration, so the correspondences set, \mathcal{A} , is given by feature extraction and matching; \mathcal{A} consists of the indices pairs, (i, j), which indicate that the *i*-th point in **P** and the *j*-th point in **Q** are matched to each other. In general, it is natural that \mathcal{A} inevitably includes inherent outlier pairs due to the false matching (Rusu et al. 2009; Yew and Lee 2018; Tzoumas et al. 2019). Accordingly, \mathcal{A} can be classified into the outlier set, \mathcal{O} , and inlier set, \mathcal{I} , such that $\mathcal{O} \cup \mathcal{I} = \mathcal{A}$. Note that these two sets are disjoint, i.e. $\mathcal{O} \cap \mathcal{I} = \emptyset$.

Then, the relative rotation matrix and translation vector between \mathcal{P} and \mathcal{Q} are denoted as $\mathbf{R} \in SO(3)$ and $\mathbf{t} \in \mathbb{R}^3$, respectively. Finally, the relationship between the *i*-th point in \mathbf{P} and the *j*-th point in \mathbf{Q} is defined as follows:

$$\mathbf{q}_j = \mathbf{R}\mathbf{p}_i + \mathbf{t} + \boldsymbol{\epsilon}_{ij} \tag{1}$$

where $\epsilon_{ij} \in \mathbb{R}^3$ denotes the unknown measurement noise. That is, ϵ_{ij} depends on whether the pair (i, j) is a true inlier or not; ϵ_{ij} is Gaussian noise if $(i, j) \in \mathcal{I}$ and is irregular error, otherwise.

3.2. Problem Definition

Based on the notations above, the objective function of correspondence-based global registration can be defined as follows:

$$\hat{\mathbf{R}}, \hat{\mathbf{t}} = \operatorname*{argmin}_{\mathbf{R} \in \mathrm{SO}(3), \mathbf{t} \in \mathbb{R}^3} \sum_{(i,j) \in \mathcal{A}} \rho \Big(r(\mathbf{q}_j - \mathbf{R}\mathbf{p}_i - \mathbf{t}) \Big)$$
(2)

where $\rho(\cdot)$ denotes a robust loss function, such as Cauchy or Huber norm (Tzoumas et al. 2019), intended to suppress undesirable large errors caused by false correspondences; $r(\cdot)$ denotes the squared residual function, i.e. $|\cdot|^2$ for the scalar and $||\cdot||^2$ for the vector.

3.3. Challenge of 3D Point Cloud Registration in Distant and Partially Overlapped Cases

Unlike object-level point cloud data whose points are evenly distributed, such as bunny dataset (Curless and Levoy 1996), a 3D point cloud captured by a spinning LiDAR sensor has somewhat different characteristics. That is, as the distance from the origin of a point cloud becomes longer, the density dramatically decreases, which is referred to as sparsity issue (Lim et al. 2021b). The sparsity issue directly has a negative impact on the quality of feature descriptors and thus induces false matching results.

To be more concrete, most feature descriptor methods, including deep learning-based methods, utilize the geometrical inter-relations between each point and its neighboring points to generate descriptors. However, as the density becomes sparse, these neighboring points are insufficient to describe the actual geometrical inter-relations because the distance between the two points becomes large. As a result, the expressibility of descriptors becomes worse. Finally, if these degraded descriptors are taken as the inputs of descriptor matching, it will deteriorate the matching performance, reducing the number of actual inliers while increasing the ratio of outliers, as shown in Fig. 2.

Based on these observations, (2) is redefined to achieve robust global registration as follows:

$$\hat{\mathbf{R}}, \hat{\mathbf{t}} = \operatorname*{argmin}_{\mathbf{R} \in \mathrm{SO}(3), \mathbf{t} \in \mathbb{R}^3} \sum_{(i,j) \in \mathcal{A} \setminus \hat{\mathcal{O}}} \rho \Big(r(\mathbf{q}_j - \mathbf{R}\mathbf{p}_i - \mathbf{t}) \Big) \quad (3)$$

where \hat{O} denotes the estimated outlier pairs. In short, our ultimate goal is to robustly estimate relative pose although



Fig. 2. (L-R): Feature matching results between frames 939 as a source cloud (cyan), and 4,206, 4,204 and 4,195 as target clouds (yellow), respectively, for Seq. 02 in the KITTI dataset. The farther the pose discrepancy is, the worse quality of the estimated correspondences is, showing an increase in the ratio of outliers (red lines) and a decrease in the number of actual inliers (green lines).

imprecise correspondences that contain the outliers are given.

4. Quatro++: Quasi-SE(3) Estimation Leveraging Ground Segmentation

4.1. System Overview

Our proposed method is based on two premises: a) most terrestrial mobile robots and autonomous driving vehicles inevitably come into contact with the ground and thus b) the yaw rotation is dominant in relative rotation motion, which are also presumed in Scaramuzza (2011) and Kim et al. (2019).

Based on these assumptions, our Quatro++ mainly consists of five parts, as shown in Fig. 1. First, our proposed method begins with ground segmentation (Lim et al. 2021b) to filter ground points within the source and target clouds, respectively, because these points are usually geometrically featureless compared with the non-ground points (Shan and Englot 2018). Empirically, it was found that ground removal improves the quality of feature matching and thus dramatically increases the success rate of global registration (see Sections 7.1 and 7.2). Second, feature extraction and matching are performed to estimate correspondences.

However, as mentioned in Section 3.3, the potential outliers are unavoidable; thus, the outliers are initially filtered out by using MCIS as the third step. Next, the relative rotation and translation are estimated by the decoupling estimation method. Accordingly, in the fourth step, quasi-SO(3) estimation is performed by utilizing the concept of GNC to estimate the relative rotation in an alternating optimization manner, while rejecting the potential outliers. Fifth, component-wise translation estimation (COTE) is exploited to estimate the relative x, y, and z, respectively, based on the concept of consensus set (Li 2009).

The methodologies and rationale for each module in Quatro++ are explained in the following subsections.

4.2. Fast and Robust Ground Segmentation as a Preprocessing Step

First, ground segmentation is performed by using Patchwork (Lim et al. 2021b) to extract non-ground points from the source and target cloud, respectively, as illustrated in Fig. 3(a). Patchwork is a region-wise ground segmentation method that firstly divides all the cloud points into multiple



Fig. 3. (a) Illustration of Patchwork (Lim et al. 2021b). CZM, R-GPF, and GLE are abbreviations of concentric zone model, region-wise ground plane fitting, and ground likelihood estimation, respectively. (b)-(c) Visualization of the ground segmentation output in Seq. 00 in the KITTI dataset for comparison. The green and red colors denote the estimated ground and non-ground points, respectively (best viewed in color). (b) Estimated ground points by using ground segmentation in LeGO-LOAM (Shan and Englot 2018) and (c) by using Patchwork (Lim et al. 2021b). Empirically, it was shown that Patchwork shows more robust segmentation with fewer false positives and false negatives (best viewed in color).

subsets based on the concentric zone model (CZM). Next, region-wise ground plane fitting (R-GPF) is used to output the central point, normal vector, and flatness for each subregion. The flatness indicates the variance of points in the subset along the normal vector. Finally, ground likelihood estimation (GLE) is used to check whether the estimated ground plane is likely to be actual ground or not. Note that the Patchwork is utilized because it shows consistent and precise ground segmentation performance compared with other well-known ground segmentation methods (Seo et al. 2022), as presented in Figs. 3(b) and 3(c).

In summary, ground segmentation has two major advantages. First, it increases the matching accuracy. In fact, the ground points usually produce lots of false matching because ground points have more ambiguous geometrical characteristics than non-ground points. Therefore, by rejecting these geometrically indistinguishable points beforehand, it is shown that ground segmentation significantly increases the success rate of the global registration methods in the distant cases (see Sections 7.1 and 7.2).



Fig. 4. Visual description of the procedure of Quatro when the frames 2,328 and 3,268 for Seq. 00 in the KITTI dataset are given as source (cyan) and target (yellow), respectively, which is a distant and partially overlapped case. (a) Estimated correspondences, A_{raw} , which include potential outliers owing to false matching. The left-top text represents the pose discrepancy. (b) Filtered correspondences by using MCIS-heuristic, which is a graph-based outlier pruning method. (c)-(e) Procedure of Quasi-SO(3) estimation based on graduated non-convexity (GNC). (c) On the initialization step, A is transformed into translation invariant measurements (TIMs) all of whose weights are set to one (brown color). (d) Next, GNC-based optimization is performed to estimate relative rotation while rejecting outliers by reweighting the weights. (e) The problem is that less than three pairs are occasionally left, which is highlighted by red circles, by assigning near-zero values to the weights of correspondences. Despite this degeneracy, our quasi-SO(3) estimation robustly outputs the relative rotation because its minimum required DoF is one. (f)-(i) Examples of component-wise translation estimation (COTE) to estimate the relative x value. (f) Boundary interval set ${}^{1}\mathcal{E}$, where the brackets [and] denote the lower and upper bounds, respectively. (g) Cardinality of the g-th consensus set, $|^{1}\mathcal{I}_{g}|$. (h) Each value of the optimal solutions, i.e. $^{1}\hat{\mathbf{t}}_{g}$, which is associated with each $^{1}\mathcal{I}_{g}$. Note that $^{1}\hat{\mathbf{t}}_{14}$ is not assigned because $^{1}\mathcal{I}_{14} = \varnothing$. (i) Residual of the objective function for all correspondences when ${}^{1}\mathbf{t}_{g}$ is given. The residual becomes lowest when g = 8, thus ${}^{1}\hat{\mathbf{t}}_{8}$ is selected as the final solution for the 1st element, i.e. ${}^{1}\hat{t} \leftarrow {}^{1}\hat{t}_{8}$ (best viewed in color).

Second, rejecting these ground points in advance reduces the computational cost of the subsequent extraction and matching procedure because almost $40 \sim 60 \%$ of points in a point cloud are from the ground once the point clouds are captured by the ground vehicles (Lim et al. 2021b). As a result, ground segmentation usually reduces the number of cloud points.

4.3. Feature Extraction and Matching

Next, feature extraction and matching are employed to estimate correspondences. In this study, fast point feature histogram (FPFH) is exploited to generate descriptors (Rusu et al. 2009). Of course, any relevant feature extraction and matching method can be used, yet we chose FPFH-based matching to guarantee the generality without any training procedure and to emphasize that global registration must be robust even though imprecise correspondences are given. For these reasons, improving feature extraction and matching itself is beyond our scope.

The procedure to estimate correspondences is summarized as follows. Given the estimated non-ground points of source and target, voxel sampling is firstly employed with voxel size ν . Then, FPFH is used by taking a voxel sampled cloud as an input to generate the feature descriptor of each point with the radius for normal estimation, r_{normal} , and that for FPFH descriptor, r_{FPFH} . Finally, the reciprocal test (Zhou et al. 2016) is performed to obtain A_{raw} , as represented in Fig. 4(a). More details about the parameters are presented in Section 6.3.

Next, MCIS (Rossi et al. 2015) is applied to reject outlier

4.4. Graph-Based Outlier Rejection

correspondences in advance. Given A_{raw} as an input, MCIS outputs the filtered correspondences, \mathcal{A} (see Fig. 4(b)) However, finding the exact maximal clique is NP-complete and its complexity increases in exponential time, so it is likely to be a bottleneck in the pipeline and thus occasionally takes lots of time. For this reason, we modify it to heuristically find a maximal clique, named MCISheuristic to reduce computational burden. That is, we estimate the maximal clique found within a threshold time as the optimal result. More details about MCIS can be found in Yang et al. (2020b) and Yin et al. (2023).

4.5. Overview of Decoupling Rotation and Translation Estimation

After initially pruning outlier correspondences, the relative rotation and translation are estimated, respectively. Sharing the philosophy of Horn (1987) and Arun et al. (1987) that $\hat{\mathbf{R}}$ and $\hat{\mathbf{t}}$ can be estimated in a decoupled manner, the relative rotation is firstly estimated in the translation-invariant space, followed by translation estimation. Note that this decoupling estimation strongly assumes that the relative rotation is successfully estimated (Yang et al. 2020b).

Here, a brief explanation of the decoupling method is presented (Lim et al. 2022). Let us assume that A is a tuple that has an order. Then, to cancel out the effect of relative translation, two consecutive pairs (i, j) corresponding to the *n*-th element of \mathcal{A} and (i', j') corresponding to the (n + 1)th element of \mathcal{A} are subtracted in a chain form, which are expressed as $\alpha_k = \mathbf{p}_{i'} - \mathbf{p}_i$ and $\boldsymbol{\beta}_k = \mathbf{q}_{j'} - \mathbf{q}_j$,

respectively. By doing so, the number of correspondences is preserved, preventing an increase in the computational cost. Note that the last element is made by the subtraction between the last and first pair in A.

 α_k and β_k are referred to as translation invariant measurements (TIMs) that satisfy $\beta_k = \mathbf{R}\alpha_k + \epsilon_k$, whose translation term is canceled out; ϵ_k denotes Gaussian noise if both (i, j) and (i', j') pairs are true inliers and ϵ_k has a large irregular error, otherwise.

Then, the rotation is estimated in translation invariant space as follows:

$$\hat{\mathbf{R}} = \underset{\mathbf{R}\in\mathrm{SO}(3)}{\operatorname{argmin}} \sum_{k=1}^{K} \min\left(w_k r(\boldsymbol{\beta}_k - \mathbf{R}\boldsymbol{\alpha}_k), \, \bar{c}^2\right) \quad (4)$$

where K denotes the cardinality of \mathcal{A} , i.e. $|\mathcal{A}|$; w_k is the weight to control the influence of each TIM pair; \bar{c} is a truncation parameter. Thus, (4) means that if either α_k or β_k is outlier, $r(\beta_k - \mathbf{R}\alpha_k)$ has large residual. Then, w_k is exploited to suppress the effect of the potential outlier. However, if the first summand, i.e. $w_k r(\beta_k - \mathbf{R}\alpha_k)$, is still large, this residual term is substituted with the predefined constant \bar{c}^2 to cut off the influence of the outliers on the estimation. Consequently, the undesirable effect of these potential outliers can be safely suppressed. The details of how to estimate w_k are presented in Section 4.7.

Then, the rotation estimation is followed by COTE (Yang et al. 2020b). As the term COTE itself indicates, 3D relative translation is estimated in an element-wise manner as follows:

$${}^{l}\hat{\mathbf{t}} = \operatorname*{argmin}_{{}^{l}\mathbf{t}\in\mathbb{R}} \sum_{(i,j)\in\mathcal{A}} \min\left(\frac{r({}^{l}\mathbf{t} - {}^{l}\mathbf{v}_{ij})}{\sigma_{ij}^{2}}, \, \bar{c}^{2}\right) \qquad (5)$$

where $\mathbf{v}_{ij} = \mathbf{q}_j - \mathbf{\hat{R}}\mathbf{p}_i$ denotes the translation discrepancy, σ_{ij} is the noise bound, and $l(\cdot)$ denotes the *l*-th element of a 3D vector. That is, l = 1, 2, 3 and each value corresponds to x, y, and z translation in the ascending order, respectively.

In summary, (3) is decomposed into (4) and (5). More details are explained in Sections 4.7 and 4.8.

4.6. Quasi-SO(3) in Urban Environments

Next, the concept of quasi-SO(3) is proposed based on the premise in Section 4.1. To be more robust against degeneracy, Quasi-SO(3) aims to reduce the DoF of rotation estimation from three to one, which is built upon a key observation that a pure yaw rotation is relatively dominant than roll and pitch rotations in a loop closing situation of urban environments. As shown in Fig. 5(a), once a revisit occurs, the roll and pitch differences between the source and target clouds are not large. Accordingly, two viewpoints of the source and target clouds are likely to be located in the coplanar space.

The experimental evidence also supports our rationale that the rotation affected by roll and pitch motions is negligible in loop closing situations. To be more concrete, the probability distribution function empirically shows that the effect of roll and pitch rotations is usually smaller than 8° (Fig. 5(b)). In general, the error of small angle assumption does not exceed 1.0 % if the angle is lower than 8° (Youn et al. 2021). In that



Fig. 5. (a) Visual description to show the geometrical characteristics of loop closing situations in urban canyons. (a) Revisit usually occurs in either forward, reverse, or perpendicular direction. Accordingly, the roll and pitch differences between the two viewpoints of source and target clouds are sufficiently small. (b) Qualitative analysis of the actual geodesic magnitude of relative roll ($\angle \mathbf{R}_x$), pitch ($\angle \mathbf{R}_y$), and the combined rotations ($\angle (\mathbf{R}_y \cdot \mathbf{R}_x))$) except yaw rotation, where $\angle \mathbf{R} = \cos^{-1} \frac{\text{Tr}(\mathbf{R}) - 1}{2}$. The loop pairs whose distances are between 0.5 to 10 m apart in the KITTI dataset were used. (c) Surrogate function $\rho_{\mu}(\cdot)$ when $\bar{c} = 0.15$. As the parameter μ gradually grows, the kernel has more non-linearity. It becomes truncated least squares once μ reaches ∞ (best viewed in color).

respect, it was demonstrated that these angles are sufficiently insignificant.

Therefore, let us express the relative rotation as $\mathbf{R} = \mathbf{R}_z \cdot \mathbf{R}_y \cdot \mathbf{R}_x$, where \mathbf{R}_z , \mathbf{R}_y , and \mathbf{R}_x denote the yaw, pitch, and roll rotations, respectively. Then, \mathbf{R} can be approximated as $\mathbf{R} \approx \mathbf{R}_z$ based on our assumption that the magnitudes of roll and pitch angles are insignificant in urban canyons (Chen et al. 2022), so the effect of \mathbf{R}_y and \mathbf{R}_x can be negligible as $\mathbf{R}_y \cdot \mathbf{R}_x \approx \mathbf{I}_3$, where \mathbf{I}_n denotes the $n \times n$ identity matrix. Consequently, the DoF of the relative rotation is reduced from three to one through this assumption. For brevity, the approximated \mathbf{R} is denoted by \mathbf{R}_+ .

One might argue that our assumption occasionally does not hold once the pose discrepancy is somewhat large and the loop closing is performed in non-flat areas. However, this problem can be easily resolved by utilizing the estimated roll and pitch angles from an inertial navigation system (INS). We propose to tackle this problem using the INS in Section 5.5.

4.7. Quasi-SO(3) Estimation using Graduated Non-Convexity to Avoid Degeneracy

Based on the assumption above, GNC-based truncated least square (Tzoumas et al. 2019) is explained to estimate \mathbf{R}_+ . To this end, (4) can be modified as $\operatorname{argmin}_{\mathbf{R}_+} \sum_{k=1}^{K} \rho_{\mu} \left(r \left(\beta_k - \mathbf{R}_+ \alpha_k \right) \right)$ by using a surrogate function, $\rho_{\mu}(\cdot)$, that controls the non-linearity of loss function governed by parameter μ (Fig. 5(c)).

Next, this equation is redefined based on Black-Rangarajan duality as follows (Zhou et al. 2016; Yang et al. 2020b):

$$\hat{\mathbf{R}}_{+} = \underset{\substack{\mathbf{R}_{+} \in \mathrm{SO}(3);\\w_{k} \in [0,1]}}{\operatorname{argmin}} \sum_{k=1}^{K} \left[w_{k} r \left(\boldsymbol{\beta}_{k} - \mathbf{R}_{+} \boldsymbol{\alpha}_{k} \right) + \Phi_{\rho_{\mu}} \left(w_{k} \right) \right]$$
(6)

where $\Phi_{\rho_{\mu}}(w_k) = \frac{\mu(1-w_k)}{\mu+w_k} \bar{c}^2$ is a penalty term (Yang et al. 2020a; Lim et al. 2022). However, minimizing the residuals while simultaneously assigning the corresponding weights does not guarantee optimality (Zhou et al. 2016).

To tackle this problem, the alternating optimization is exploited to solve (6). That is, the optimal rotation is estimated with the fixed weights updated in the previous iteration and then the weights are updated with the fixed optimized rotation. To be more concrete, the alternating optimization is expressed as follows:

$$\hat{\mathbf{R}}_{+}^{(t)} = \operatorname*{argmin}_{\mathbf{R}_{+} \in \mathrm{SO}(3)} \sum_{k=1}^{K} \hat{w}_{k}^{(t-1)} r\left(\boldsymbol{\beta}_{k} - \mathbf{R}_{+} \boldsymbol{\alpha}_{k}\right), \quad (7)$$

$$\hat{\mathbf{W}}^{(t)} = \underset{w_k \in [0,1]}{\operatorname{argmin}} \sum_{k=1}^{K} \left[w_k r \left(\boldsymbol{\beta}_k - \hat{\mathbf{R}}_+^{(t)} \boldsymbol{\alpha}_k \right) + \Phi_{\rho_{\mu}} \left(w_k \right) \right]$$
(8)

where the superscript (t) denotes the *t*-th iteration and $\hat{\mathbf{W}}$ is a matrix form of the weights, i.e. $\hat{\mathbf{W}} =$ diag $(\hat{w}_1, \hat{w}_2, \dots, \hat{w}_K)$. Each \hat{w}_k can be obtained by solving the following equation:

$$\frac{\partial}{\partial w_k} \left(w_k r \left(\boldsymbol{\beta}_k - \hat{\mathbf{R}}_+^{(t)} \boldsymbol{\alpha}_k \right) + \Phi_{\rho_\mu} \left(w_k \right) \right) = 0.$$
 (9)

Once we expand (9) in terms of \hat{w}_k , \hat{w}_k is expressed in a truncated manner as follows:

$$\hat{w}_{k}^{(t)} = \begin{cases} 0 & \text{if } \tilde{r}_{k} \in \left[\frac{\mu+1}{\mu}\bar{c}^{2}, +\infty\right) \\ \bar{c}\sqrt{\frac{\mu(\mu+1)}{\bar{r}_{k}}} - \mu & \text{if } \tilde{r}_{k} \in \left[\frac{\mu}{\mu+1}\bar{c}^{2}, \frac{\mu+1}{\mu}\bar{c}^{2}\right) \\ 1 & \text{otherwise} \end{cases}$$
(10)

where $\tilde{r}_k \geq 0$ is the simplified notation of $r(\boldsymbol{\beta}_k - \mathbf{R}^{(t)}_+ \boldsymbol{\alpha}_k)$.

Note that the value of μ grows for every iteration to gradually increase the non-linearity of the loss function. That is, μ is assigned as $\mu^{(t)} \leftarrow \kappa \cdot \mu^{(t-1)}$, where $\kappa > 1$ is a constant parameter. By doing so, the kernel becomes more non-linear, as presented in Fig. 5(c). For initialization, $\mu^{(0)}$ is set to be $\bar{c}^2/(\max(r(\beta_k - \alpha_k)) - \bar{c}^2)$. The optimization ends if either the differential of the loss function is sufficiently small or the iteration reaches the maximum number, N_{iter} .

In summary, the estimation of \mathbf{R}_+ has two major advantages over that of \mathbf{R} . First, estimation of the \mathbf{R}_+ helps to avoid the degeneracy situations in itself by reducing the DoF of rotation from three to one. Sometimes, the number of remaining inliers is less than three, because GNC does not guarantee to leave more than three inliers. Even though the degeneracy occurs, our proposed rotation estimation enables to tolerate the case where only one or two



Fig. 6. Changes of weights by GNC in quasi-SO(3) estimation in a degeneracy case: under the situation where the points having weights below 0.5 (dotted line) are filtered out as outliers, quasi-inliers help to estimate quasi-SO(3) by acting as inliers, thus our proposed method enables to tolerate the degeneracy.

correspondences remain because a single correspondence is enough to estimate \mathbf{R}_+ .

It may be thought that this phenomenon does not occur often; however, as shown in Figs. 4(d), 4(e), and 6, the degeneracy occurs once the distance between the two viewpoints becomes more than 9 meters away in outdoor scenes, empirically. Therefore, being robust against the degeneracy can be interpreted as being able to help the convergence of local registration better when performing loop closing in the distant cases (see Sections 5.4 and 7.6).

Second, \mathbf{R}_+ allows measurements that are originally outliers to be treated as inliers. To explain this reason, we introduce the concept of *quasi-inliers* each of whose error only exists along the perpendicular direction of the ground plane. That is, let us decompose the noise term in the translation invariant space into two components: a) a perpendicular term and b) a parallel term to the ground plane. Then, the quasi-inlier represents the measurement whose former error term has a value greater than zero, yet the latter one has a near-zero value.

Note that these quasi-inliers often occur owing to the ambiguity of the descriptors in urban environments. That is, the majority of urban structures are generally orthogonal to the ground, so the local geometrical properties along the perpendicular direction are highly likely to be similar, resulting in false matching.

The point is that the quasi-inliers behave like the actual inliers during the estimation of \mathbf{R}_+ . That is, the perpendicular error terms rarely affect the estimation of \mathbf{R}_+ because \mathbf{R}_+ primarily considers the effect of yaw rotation. In other words, the estimation of \mathbf{R}_+ is equivalent to the estimation of yaw rotation with the measurements projected into the *xy*-plane, so the effect of error along the perpendicular direction can be negligible.

In short, as these quasi-inliers play a role as the actual inliers, the quasi-inliers help to increase the number of inliers and prevent degeneracy, as shown in Fig. 6.

4.8. Component-Wise Translation Estimation

Finally, COTE is exploited to estimate the component-wise relative translation by the following four steps. For each *l*-th element, boundary interval set, which is a $2|\mathcal{A}|$ -tuples and denoted as ${}^{l}\mathcal{E}$, is firstly set which comprises the lower bound, ${}^{l}\mathbf{v}_{ij} - \sigma_{ij}\bar{c}$, and the upper bound, ${}^{l}\mathbf{v}_{ij} + \sigma_{ij}\bar{c}$ ([and] in Fig. 4(f), respectively). Note that it is assumed that all the elements of ${}^{l}\mathcal{E}$ are sorted in ascending order. That is, let the *g*-th element in ${}^{l}\mathcal{E}$ be ${}^{l}\mathcal{E}(g)$, whose value could be the

lower or upper bound, then it satisfies ${}^{l}\mathcal{E}(g) < {}^{l}\mathcal{E}(g+1)$ for $g = 1, 2, \dots, 2|\mathcal{A}| - 1$.

Second, the g-th consensus set is assigned as follows:

$${}^{l}\mathcal{I}_{g} = \{(i,j) | \frac{({}^{l}\phi_{g} - {}^{l}\mathbf{v}_{ij})^{2}}{\sigma_{ij}^{2}} \le \bar{c}^{2} \}$$
(11)

where ${}^{l}\phi_{g} \in \mathbb{R}$ denotes an arbitrary value which satisfies ${}^{l}\mathcal{E}(g) < {}^{l}\phi_{g} < {}^{l}\mathcal{E}(g+1)$.

Third, the optimal solution for each non-empty ${}^{l}\mathcal{I}_{g}$, which is denoted as ${}^{l}\hat{\mathbf{t}}_{g}$, is estimated as follows:

$${}^{l}\hat{\mathbf{t}}_{g} = \left(\sum_{(i,j)\in {}^{l}\mathcal{I}_{g}} \frac{1}{\sigma_{ij}^{2}}\right)^{-1} \sum_{(i,j)\in\mathcal{I}_{g}} \frac{1}{\sigma_{ij}^{2}} {}^{l}\mathbf{v}_{ij}.$$
 (12)

Finally, among the optimal solution candidates estimated by (12), the one that minimizes the truncated objective function, i.e. (5), is chosen as the global optimum (the red dashed rectangle in Fig. 4(i)) as follows:

$${}^{l}\hat{\mathbf{t}} = \operatorname*{argmin}_{{}^{l}\mathbf{t}_{g}\in\mathbf{H}} \sum_{(i,j)\in\mathcal{A}} \min\left(\frac{r({}^{l}\mathbf{t}_{g} - {}^{l}\mathbf{v}_{ij})}{\sigma_{ij}^{2}}, \bar{c}^{2}\right)$$
(13)

where **H** is a set of the optimal solution candidates whose size is up to $2|\mathcal{A}| - 1$. Note that COTE presumes that $\hat{\mathbf{R}}_+$ is sufficiently precise.

5. Quatro++ in Back-End of LiDAR SLAM

In this section, the details of how Quatro++ is integrated into the LiDAR SLAM frameworks. Originally, our proposed method is a front-end agnostic method, so it can be easily exploited in any LiDAR SLAM frameworks. In doing so, Quatro++ is utilized as an initial alignment method to help the estimate of local registration successfully converge into the global optimum.

5.1. Generic Pose Graph Optimization in LiDAR SLAM

As mentioned in Section 1, PGO is used to minimize global trajectory errors, which are caused by local drift of odometry. That is, these errors are accumulated as time goes by, resulting in a large pose discrepancy between the estimated poses and the actual poses.

To tackle this problem, loop detection and loop closing are utilized, which mainly consists of three steps. First, loop detection is performed to identify inter-related nodes (Kim and Kim 2018; Chen et al. 2022; Cattaneo et al. 2022). Second, a registration algorithm is exploited to estimate the actual pose difference. Finally, global pose errors are minimized via PGO by utilizing the pose differences estimated by the registration between the two nonconsecutive nodes.

In summary, the objective function of PGO is expressed as follows:

$$\hat{\mathcal{X}} = \underset{\mathbf{X}\in\mathcal{X}}{\operatorname{argmin}} \sum_{i} \left\| \mathbf{X}_{i}^{-1} \mathbf{X}_{i+1} \boxminus \mathbf{z}_{i}^{\operatorname{odom}} \right\|_{\Sigma_{i}^{\operatorname{odom}}}^{2} + \sum_{(j,k)\in\mathcal{L}} \left\| \rho_{jk} \left(\mathbf{X}_{j}^{-1} \mathbf{X}_{k} \boxminus \mathbf{z}_{jk}^{\operatorname{reg}} \right) \right\|_{\Sigma_{jk}^{\operatorname{loop}}}^{2}$$
(14)

9

where the first summand denotes the summation of Mahalanobis distances estimated by the odometry and the last summand denotes that by the loop closing. The definitions of the variables are as follows: \mathcal{X} and $\hat{\mathcal{X}}$ are sets of 3D poses and the optimized poses, respectively; i, j, and k are the indices of the node in a graph structure where the j-th and k-th nodes are not the successive nodes; \mathcal{L} denotes an indices set of loops and $\mathbf{X}_i \in SE(3)$ is the i-th pose that corresponds to the i-th node; \boxminus denotes box-minus operation (He et al. 2021) which is an equivalent operator of subtraction between two manifolds, i.e. in a 3D space, $\boxminus : SE(3) \times SE(3) \to \mathbb{R}^6$; $\mathbf{z}_i^{\text{odom}}$ denotes the measurement from the odometry and $\mathbf{z}_{jk}^{\text{reg}}$ from the registration; Σ_i^{odom} and $\mathbf{z}_{jk}^{\text{reg}}$, respectively.

5.2. Mean Squared Error-Based False Loop Rejection

In general, the loops, i.e. \mathcal{L} , can be redefined as $\mathcal{D} \setminus \mathcal{O}_{\mathcal{L}}$ where \mathcal{D} is all the loop candidates searched by a loop detection method and $\hat{\mathcal{O}}_{\mathcal{L}}$ denotes the rejected loop candidates by considering them as false loops.

Most LiDAR SLAM methods (Shan and Englot 2018; Shan et al. 2020; Li et al. 2021; Ramezani et al. 2022) heavily rely on mean squared error (MSE) of local registration methods to determine whether the loop candidate is considered as a false loop or not, i.e. the output of getFitnessScore() in point cloud library (Rusu and Cousins 2011; Aldoma et al. 2012). The MSE is the average of squared Euclidean distances from the transformed source cloud points by the estimated relative rotation and translation to the corresponding target cloud points (Rusu and Cousins 2011). Accordingly, the MSE indicates how the source cloud is tightly matched to the target cloud.

Once the estimate successfully converges into the actual global minimum (Fig. 7(a)), the MSE usually shows a small value because each distance between the target point and the warped source point is near zero. In contrast, the MSE shows a larger value in the divergence cases if the estimate goes to the local minimum and thus two point clouds are not aligned (Fig. 7(b)). Based on these observations, the loop is rejected if the MSE is somewhat large because a number of distances between the transformed source and target are not zero.

Thus, this MSE-based rejection effectively filters out most false loops, leaving some actual loops among the loop candidates.

5.3. Potential Limitations of the Mean Squared Error-Based False Loop Rejection

However, there are two problems in terms of a) false negative and b) false positive loops. First, the narrow convergence region of the local registration methods triggers too many rejections even though these loops are the actual loops. As mentioned earlier, the local registration methods estimate correspondences based on the strong assumption that the nearest points between the source and target clouds are valid. However, once the pose difference becomes large and only a small portion of two point clouds is overlapped, this assumption does not hold, so the estimate generally



Fig. 7. Registration results once (a)-(b) local registration is only applied and (c) a coarse-to-fine alignment is applied by utilizing our proposed method. The texts in the upper left corner represent the initial pose differences. The green and red boxes indicate whether the registration succeeds or not, respectively. The cyan and yellow colors denote the transformed source cloud and the target cloud, respectively (best viewed in color).

diverges. For this reason, many loop candidates whose pose discrepancies are large are highly likely to be rejected, causing lots of false negative loops.

Second, there are some failure cases having low MSEs, causing false positive loops. This phenomenon usually occurs in corridor-like scenes where the surroundings have few geometrical characteristics and high ambiguity. For this reason, it is hard to discern whether the registration succeeded or not solely based on the MSE. As a result, these false positive loops deteriorate the quality of loop constraints. That means, $\mathbf{z}_{jk}^{\text{reg}}$ in (14) becomes more imprecise.

5.4. Why Quatro++?: For More Loops and Accurate Measurements

To tackle these two problems, our proposed method can be utilized as a coarse alignment to provide an initial guess, helping the local registration methods to successfully estimate the relative pose, as presented in Fig. 7. As the performance of Quatro++ is less affected by the initial pose difference, our global registration can initially transform the source cloud to the region where the estimate of local registration can converge into the global minimum.

In summary, the coarse-to-fine alignment-based loop closing has two main advantages over solely using a local registration method. First, the loop closing with our proposed method significantly lessens wrong rejection of the loop candidates, i.e. false negatives, owing to the lowered MSEs. Second, our proposed coarse-to-fine-based loop closing in itself increases the accuracy of measurements. In particular, it was shown that our proposed method is robust against the corridor-like environments, so that the local minima cases can be reduced as well (see Section 7.1). Consequently, the coarse-to-fine loop closing increases the quality of z_{jk}^{reg} in (14).

Therefore, these two merits result in more accurate and precise PGO results (see Section 7.6).

5.5. Quatro++ With Roll-Pitch Compensation Using INS

As mentioned earlier, our proposed method sacrifices the estimation of relative roll and pitch angles to improve robustness. However, this limitation can be easily resolved by leveraging the INS system. That is, the relative roll and pitch angles between the two point clouds are easily obtained by the INS system. These angles are highly relevant to gravity direction, which can be fully observable in the INS system. In other words, the roll and pitch angles can be estimated with respect to the world frame, which indicates that these estimated angles are drift-free.

Based on these observations, the source point cloud is first transformed into the roll-pitch-compensated frame by using the estimated rotation, $\hat{\mathbf{R}}_{\text{INS}} = \hat{\mathbf{R}}_y \cdot \hat{\mathbf{R}}_x$, where $\hat{\mathbf{R}}_y$ and $\hat{\mathbf{R}}_x$ denote the elements of relative rotation affected by pitch and roll angles, respectively. The combination of using our proposed method with $\hat{\mathbf{R}}_{\text{INS}}$ is very simple, yet this compensation significantly reduces the relative pose errors while preserving the robustness of our proposed method (see Sections 7.1 and 7.5).

6. Experiments

6.1. Dataset

We conducted experiments with various datasets to evaluate the performance of Quatro++ (and Quatro). Accordingly, the following datasets were used: KITTI dataset (Geiger et al. 2012, 2013) to evaluate the success rate and robustness of global registration, NAVER LABS localization dataset (Lee et al. 2021) to check the applicability to more sparse LiDAR scans, MulRan dataset (Kim et al. 2020) to integrate our proposed method as a loop closing module, and Hilti-Oxford dataset (Zhang et al. 2022) to show feasibility of our Quatro++ on hand-held sensor configurations. These datasets were captured by Velodyne HDL-64E, Velodyne VLP-16, Ouster OS1-64, and HESAI XT32, respectively. That is, our experiments were conducted by using various mechanicallyspinning-type LiDAR sensors to check the generality and versatility.

We look closely at three points: a) performance comparison with other state-of-the-art (SOTA) global registration methods, b) the effect of ground segmentation on global registration, and c) application of our proposed method to LiDAR SLAM frameworks. In particular, we place more emphasis on the distant cases where the viewpoints of two point clouds are far apart. Accordingly, the experiments are further categorized into three parts: a) loop closing test, b) odometry test, and c) augmented rotation test.

That is, first, let us define all the loop pairs as \mathcal{S}_{All} as follows:

$$S_{\text{All}} = \left\{ (s,t) | r_{\min} \le \| \mathbf{t}_s - \mathbf{t}_t \|^2 \le r_{\max}, \\ |s-t| \ge m, \forall s, t \in \mathbb{N}_{\text{traj}} \right\},$$
(15)

where s and t are the indices of source and target point clouds, respectively; \mathbf{t}_s and \mathbf{t}_t are the translation vectors of ground truth poses that correspond to the source and target point clouds, respectively; r_{\min} and r_{\max} are the minimum and maximum distances between the source and target point clouds, respectively; m is the minimum number of frames between the source and target point clouds to prevent the loop pairs from being too close to each other; and \mathbb{N}_{traj} is the set of all the indices of the point cloud and ground truth pose pairs.

Then, loop pairs for the loop closing test, S_L , is sampled from S_{All} , i.e. $S_L \subset S_{All}$, where $|S_L| = N_{\text{sample}}$. In this



Fig. 8. (a)-(b) Examples of loop pairs of Seq. 00 and Seq. 08 in the KITTI dataset, each of whose distances is between 10 and 12 m, i.e. $r_{\min} = 10$ and $r_{\max} = 12$ m. The rainbow color indicates the poses over time and the trajectory is airborne for better understanding. The cyan-colored lines are the sampled loop pairs, S_L , whose size is 1,000, between two poses in loop closing situations (best viewed in color).

paper, we set $N_{\text{sample}} = 1,000$ and the symbol "A ~ B" denotes that $r_{\min} = A$ and $r_{\max} = B$. S_L is presented in Fig. 8 for better understanding. In this paper, the maximum distance we aim to achieve is set to 12 m, i.e. B= 12, which is double the criteria used in previous researches (Kim and Kim 2018; Cattaneo et al. 2022) when setting true loop pairs. Thus, our experiments present the robustness of registration against large pose discrepancy. In addition, when the pose difference exceeds 12 m, two point clouds overlap so partially that the output of local registration as fine alignment may not reach the global minimum, in turn leading to inaccurate loops.

In case of the odometry test, the set of pairs, S_O , is defined as follows:

$$S_O = \left\{ (\Delta n + 1, \Delta (n+1) + 1) | n \in \mathbb{Z}, \\ 0 \le n < \frac{\mathbb{N}_{\text{traj}} - 1}{\Delta} - 1 \right\},$$
(16)

where $\Delta > 0$ denotes the interval between two frames. Finally, the pairs for augmented rotation test, S_A , is similar to S_O , but we additionally rotate the target cloud along the yaw direction to check the robustness against large angle discrepancy.

6.2. Error Metrics

To evaluate algorithms quantitatively, both average pose errors (average translation error, t_{avg} , and average rotation error, r_{avg}) and the relative pose errors (relative translation error, t_{rel} , and relative rotation error, r_{rel}) were exploited; t_{avg} and r_{avg} are defined as follows:

•
$$t_{\text{avg}} = \sum_{n=1}^{N} (\mathbf{t}_{n,\text{GT}} - \hat{\mathbf{t}}_n)^2 / N,$$

• $r_{\text{avg}} = \frac{180}{\pi} \cdot \sum_{n=1}^{N} \cos^{-1}(\frac{\text{Tr}(\hat{\mathbf{R}}_n^{\mathsf{T}} \mathbf{R}_{n,\text{GT}}) - 1}{2}) / N$

where $\mathbf{t}_{n,\text{GT}}$ and $\mathbf{R}_{n,\text{GT}}$ are the *n*-th ground truth translation and rotation, respectively; *N* denotes the number of total samples; t_{rel} and r_{rel} are used for the odometry test, which are calculated by RPG evaluation tools (Zhang and Scaramuzza 2018).

Furthermore, we evaluate global registration methods in terms of success rate, which is a more important metric to directly evaluate the robustness of the global registration needed in loop closing situations of LiDAR SLAM. The criteria of the success rate is based on Kim et al. (2019): a pose error that is sufficiently within the narrow convergence region where the estimate of local registration successfully converges into the global optimum. Thus, the registration is counted as a success if both translation and rotation errors are lower than 2 m and 10°, respectively.

6.3. Parameters of Quatro++

Parameter settings of our proposed method in our experiments are presented in Table 1. It should be noted that parameters of FPFH are set to satisfy the following relations: $\nu < r_{\text{normal}} < r_{\text{FPFH}}$. Further implementation details can be found in our open-source code.

7. Experimental Results

In our experiments, other well-known global registration methods are employed as follows: RANSAC (Fischler and Bolles 1981) whose maximum iteration is set to 10,000, FGR (Zhou et al. 2016), and TEASER++ (Yang et al. 2020b). Here, the main differences between our proposed method and the baseline methods are explained. First, RANSAC and FGR have no outlier rejection modules like MCIS. FGR and TEASER++ are also GNC-based methods, so alternating optimization is performed to reject the outlier pairs. The biggest difference between our proposed method and each baseline method is as follows: FGR uses linearization in SE(3) manifold and the pose difference is estimated by gradient-based optimization; Similar to our proposed method, TEASER++ also utilizes the decoupling method. However, TEASER++ estimates SO(3) estimation, unlike our quasi-SO(3) estimation, and does not exploit ground segmentation. To avoid confusion, we stress that the meanings of "++" in Quatro++ and TEASER++ are different: TEASER++ is a speeded up version of TEASER at an optimization level and Quatro++ denotes the combination of Quatro with ground segmentation.

Furthermore, we compare our Quatro++ with deep learning-based methods: MDGAT-Matcher (Shi et al. 2021a)

Table 1. Parameters of each module in our global registration pipeline. Note that the parameters of FPFH should be set differently depending on the sensor configuration or the number of laser scans. The values before, in the middle of, and after the slash are the parameters for Velodyne VLP-16, HDL-64E, and Ouster OS1-64 sensors, respectively, to consider different sparsity characteristices, i.e. differences of vertical resolution between two adjacent rays.

| | | Param. | Description | Value |
|-------|--------|--|---|--|
| tro++ | FPFH | $ u \\ r_{ m normal} \\ r_{ m FPFH} $ | Voxel sampling size Radius for normal estimation Radius of FPFH descriptor | 0.1 / 0.3 / 0.6 m 0.3 / 0.5 / 1.5 m 0.45 / 0.65 / 2.25 m |
| Quat | Quatro | $egin{array}{c} ar{c} \\ N_{	ext{iter}} \\ \kappa \end{array}$ | Noise bound Maximum iteration of quasi-SO(3) estimation A factor to control the increase of GNC | 0.3 50 1.4 |



Fig. 9. (a) Source (red) and target clouds (green), where the left-bottom texts represent the pose discrepancy in each scene. (b)-(c) Registration results of our Quatro and Quatro++ between 1,939 to 4,531 frames in Seq. 00 and between 165 to 999 frames in Seq. 08 in the KITTI dataset, where the blue points denote the warped source cloud by the estimated pose. Quatro++ successfully registers the source and target clouds even though both cloud points are measured in distant and unobstructed cases, where few artificial structures are located. Black dashed boxes highlight the non-ground objects that have to be tightly aligned. The solid red and green boxes indicate the algorithms failed and succeeded, respectively (best viewed in color)

Table 2. Comparison of success rates on the KITTI dataset. The bold and the gray-highlight denote the best and the second-best performance, respectively (unit: %).

| | Sequence | | 00 | | | 02 | | | 05 | |
|-------|--------------------|-----------------------|--------------|------------------------|----------------------------------|-------------|------------------------|---------------------|-------------|------------------------|
| | Diff. of viewpoint | $2 \sim 6 \mathrm{m}$ | $6\sim 10m$ | $10 \sim 12\mathrm{m}$ | $2 \mathrm{m} \sim 6 \mathrm{m}$ | $6\sim 10m$ | $10 \sim 12\mathrm{m}$ | $2\sim 6\mathrm{m}$ | $6\sim 10m$ | $10 \sim 12\mathrm{m}$ |
| Conv. | RANSAC | 20.3 | 8.5 | 3.3 | 32.3 | 12.0 | 5.8 | 30.3 | 9.3 | 7.7 |
| | FGR | 99.2 | 79.3 | 45.5 | 71.2 | 41.8 | 21.2 | 99.5 | 83.7 | 48.1 |
| | TEASER++ | 100.0 | 94.7 | 84.6 | 95.8 | 78.7 | 71.8 | 99.9 | 99.1 | 93.9 |
| | Quatro (Ours) | 99.9 | 94.5 | 84.7 | 96.1 | 79.5 | 71.9 | 100.0 | 99.3 | 94.5 |
| | Quatro++ (Ours) | 99.9 | 94.8 | 87.1 | 98.7 | 97.1 | 90.8 | 99.9 | 99.7 | 98.5 |
| Deep | MDGAT-Matcher | 94.0 | 87.3 | 69.9 | 12.0 | 0.0 | 0.0 | 94.5 | 89.9 | 81.4 |
| | LCDNet (fast) | 78.5 | 12.8 | 0.8 | 70.5 | 15.8 | 2.9 | 87.7 | 21.2 | 0.9 |
| | LCDNet | 100.0 | 90.3 | 62.1 | 91.4 | 66.4 | 41.3 | 100.0 | 98.8 | 78.8 |
| | Sequence | | 06 | | | 07 | | | 08 | |
| | Diff. of viewpoint | $2\sim 6\mathrm{m}$ | $6\sim 10m$ | $10 \sim 12\mathrm{m}$ | $2 \mathrm{m} \sim 6 \mathrm{m}$ | $6\sim 10m$ | $10 \sim 12\mathrm{m}$ | $2\sim 6\mathrm{m}$ | $6\sim 10m$ | $10\sim 12\text{m}$ |
| Conv. | RANSAC | 19.1 | 8.7 | 3.3 | 17.0 | 5.6 | 5.3 | 18.6 | 6.4 | 2.7 |
| | FGR | 99.4 | 92.6 | 59.2 | 99.7 | 81.8 | 41.2 | 36.2 | 24.0 | 15.9 |
| | TEASER++ | 99.9 | 99.5 | 96.0 | 99.9 | 99.7 | 95.8 | 100.0 | 99.7 | 96.6 |
| | Quatro (Ours) | 100.0 | 99.7 | 97.7 | 100.0 | 99.7 | 95.8 | 100.0 | 99.8 | 98.0 |
| | Quatro++ (Ours) | 100.0 | 100.0 | 100.0 | 100.0 | 99.9 | 98.0 | 99.9 | 99.8 | 98.9 |
| Deep | MDGAT-Matcher | 99.1 | 99.9 | 91.7 | 92.1 | 94.1 | 68.3 | 17.9 | 55.6 | 7.7 |
| | LCDNet (fast) | 47.4 | 4.5 | 3.1 | 70.4 | 11.3 | 0.0 | 56.5 | 5.7 | 0.2 |
| | LCDNet | 100.0 | 99.8 | 85.2 | 100.0 | 98.6 | 80.3 | 100.0 | 95.9 | 75.7 |

and LCDNet (Cattaneo et al. 2022). We evaluate two versions of LCDNet: one using an unbalanced optimal transport algorithm, which estimates the transformation matrix in a learning-based manner, denoted as LCDNet (fast), and the other one uses RANSAC with 50,000 iterations, denoted as LCDNet.

In particular, we compare the generalization capabilities of our approach and the deep learning-based methods, so the networks trained on the KITTI dataset are tested on the MulRan dataset. Note that both datasets were acquired by 64channel LiDAR sensors, but the KITTI dataset was acquired by the Velodyne HDL-64E sensor, while the MulRan dataset was acquired by the OS1-64 sensor. Therefore, we can check the generalization capabilities in terms of both environmental changes and different sensor configuration.

Prepared using sagej.cls

When evaluating the coarse-to-fine alignment performance, we use the term Quatro++-c2f, which is comprised of the proposed Quatro++ as a coarse alignment and G-ICP (Segal et al. 2009) as a fine alignment.

7.1. Performance Comparison Between Quatro and Quatro++

First, the performance of Quatro and Quatro++ was analyzed. Note that the difference between Quatro++ and Quatro is with and without the ground segmentation module before the feature extraction and matching steps. As shown in Fig. 9, Quatro occasionally fails to perform registration once the pose discrepancy between the source and target clouds is large, so the correspondences from the central ground points are dominant (Fig. 9(b)). In contrast, Quatro++ showed

| | Sequence | | | DCC0 | 1 | | DCC02 | | | DCC03 | |
|-------|---|--------------------|--|---|--|---|---|---|--|--|--|
| | Diff. of vie | wpoin | t $ 2 \sim 6 \mathrm{m}$ | $6\sim 10\mathrm{m}$ | n $10 \sim 12 \mathrm{m}$ | $2\mathrm{m}\sim 6\mathrm{m}$ | $6\sim 10m$ | $10 \sim 12\mathrm{m}$ | $2\sim 6\mathrm{m}$ | $6\sim 10m$ | $10 \sim 12 \mathrm{m}$ |
| Conv. | RANSAC FGR TEASER+ Quatro (Ou Quatro++ (| + irs) Ours) | 74.2 63.5 95.2 95.3 95.4 | 46.8 43.9 92.1 92.6 92.6 | 32.2 41.0 86.4 86.5 86.9 | 78.0 63.3 99.0 98.8 99.0 | 48.0 50.9 93.8 93.7 95.0 | 34.6 44.1 88.4 90.5 90.5 | 76.2 65.4 99.4 99.3 99.0 | 50.3 45.7 97.8 97.9 97.7 | 30.0 35.7 91.8 91.3 91.8 |
| Deep | LCDNet (fa LCDNet | ast) | 16.6 93.8 | 0.4 76.1 | 0.1 48.0 | 28.3 98.6 | 0.0 85.2 | 0.0 50.0 | 19.2 98.6 | 0.0 87.0 | 0.0 61.7 |
| | Sequence | | 1 | Riversi | de01 | R | liverside | 02 | : | Riverside | e03 |
| | Diff. of vie | wpoin | t $ 2 \sim 6 m$ | $6\sim 10\mathrm{r}$ | n $10 \sim 12 \mathrm{m}$ | $2\mathrm{m}\sim 6\mathrm{m}$ | $6\sim 10m$ | $10\sim 12m$ | $2\sim 6m$ | $6\sim 10m$ | $10 \sim 12 \mathrm{m}$ |
| Conv. | RANSAC FGR TEASER+- Quatro (Ou Quatro++ (| + irs) Ours) | 60.2 85.9 85.5 85.7 91.2 | 40.7 77.3 78.1 79.9 87.0 | 29.1 62.6 73.4 78.9 84.1 | 65.2 96.9 92.3 93.1 98.2 | 42.6 91.0 83.9 86.3 93.1 | 36.0 76.9 73.5 76.2 85.8 | 46.7 80.9 80.2 80.2 91.1 | 18.9 45.5 48.2 49.0 70.7 | 25.0 49.5 65.4 68.3 75.8 |
| Deep | LCDNet (fa | ast) | 27.5 86.0 | 0.1 67.4 | 0.0 38.0 | 65.8 93.2 | 10.3 76.1 | 0.1 51.4 | 20.4 79.8 | 2.5 48.1 | 0.1 56.8 |
| | | | Sequence | | KAI | ST02 | | KAISI | 03 | | |
| | | | Diff. of view | point 2 | $\sim 6 \mathrm{m}$ $6 \sim 1$ | 10 m $10 \sim 1$ | 2 m $2 \text{ m} \sim$ | 6 m $6 \sim 10$ | 0m 10~ | - 12 m | |
| | | Deep Conv. | RANSAC FGR TEASER++ Quatro (Our Quatro++ (C LCDNet (fas LCDNet | s) Durs) | 74.8 60 73.8 72 94.7 93 94.0 92 93.5 91 16.3 0. 94.4 86 | 8 37.8 3 68.8 0 88.5 5 87.8 3 88.5 7 0.4 2 71.9 | 8 87. 8 99. 5 99. 5 99. 5 100 20. 9 99. | 3 69.5 9 98.8 9 98.5 9 98.5 9 98.5 1 2.2 9 93.6 | $\begin{array}{c} 0 & 44 \\ 3 & 92 \\ 0 & 93 \\ 0 & 92 \\ \hline 0 & 92 \\ \hline 5 & 98 \\ \hline 0 \\ 0 \\ 5 & 73 \end{array}$ | 4.3 2.2 3.3 4.3 3.0 0.0 3.7 | |

Table 3. Comparison of success rates on the MulRan dataset. The bold and the gray-highlight denote the best and the second-best performance, respectively (unit: %).

more robust performance under these situations, successfully warping the source cloud into the target cloud (Fig. 9(c)).

Consequently, as shown in Tables 2 and 3, Quatro++ presented a higher success rate, which is the most important evaluation metric to check the suitability of global registration methods as an initial alignment. In particular, we demonstrate that Quatro++ becomes more robust in distant cases, whose difference in viewpoints between the source and target clouds ranges between 10 to 12 m. Furthermore, Quatro++ showed a remarkable improvement when our proposed method is used in Seq. 02, which is a rural environment that contains few non-ground objects and the ground points are more dominant.

7.2. Effect of Ground Segmentation on Global Registration

In addition, we conducted three in-depth analyses of the effect of the ground segmentation on the global registration as follows: impact of ground segmentation a) when two point clouds with large discrepancy of viewpoints are given and b) when more spurious correspondences are given, and c) influence of the performance of ground segmentation on global registration.

First, as shown in Fig. 10 and Table 4, ground segmentation substantially increases the performance of all global registration methods when the viewpoint discrepancy between source and target is large. The main reason for the poor performance without the ground segmentation is because of the too many correspondences from the central ground points of the source to those of the target. That is, both ground points of the source and target close to the origin are sufficiently dense regardless of the large pose

discrepancy. Accordingly, even though these central points are not measured from the same viewpoint, the high density makes descriptors from the central ground points similar, leading to gross outlier correspondences. Consequently, these false correspondences finally lead to wrong pose estimation (empirically, global registration methods estimate the relative translation as a near-zero vector in this case as shown in Fig. 9(b)). In summary, the high density of the central points is misunderstood as a geometrical characteristic itself.

While the ground segmentation module occasionally decreases the performance slightly, Quatro with ground segmentation showed better performance when the viewpoint discrepancy between source and target is large (note that the combination of Quatro and ground segmentation is equal to Quatro++). Because Quatro estimates the relative rotation by using the quasi-SO(3) estimation, its reduced DoF allows us to perform robust global registration even though few correspondences are given. That means Quatro is more robust to the reduced number of correspondences due to the rejection of many ground points than TEASER++. In contrast, TEASER++, whose DoF for estimation of rotation is three, sometimes showed worse performance because rejection of ground points substantially reduces the number of correspondences, resulting in degeneracy. Finally, the degeneracy potentially leads to large rotation errors of TEASER++. Therefore, these results demonstrate the potential of synergy between our Quatro and ground segmentation.

Second, we validate that ground segmentation especially improves the performance of global registration even when more spurious correspondences are given, as shown in Table 4. We deliberately set the parameters of FPFH as $\nu = 0.3 \,\mathrm{m}$, $r_{\mathrm{normal}} = 0.5 \,\mathrm{m}$, and $r_{\mathrm{FPFH}} = 0.65 \,\mathrm{m}$ in the MulRan dataset (note that the appropriate parameters for MulRan dataset are $\nu = 0.6$ m, $r_{\text{normal}} = 1.5$ m, and $r_{\text{FPFH}} =$ 2.25 m, as presented in Table 1). Once more spurious correspondences are given, the success rates of global registration methods without ground segmentation are more dramatically decreased, as shown in Table 4. However, after the application of ground segmentation, global registration approaches showed better success rates, which supports that ground segmentation helps to increase the coverage of the global registration methods. In particular, ground segmentation brought much greater performance increase when the spurious correspondences were given than when accurate ones were given, which also corroborates that ground segmentation effectively prevents the catastrophic failure caused by imprecise feature matching.

Finally, we analyzed the results of the global registration according to the performance of the ground segmentation. To this end, we employed GPF (Zermas et al. 2017), LineFit (Himmelsbach et al. 2010), CascadedSeg (Narksri et al. 2018), R-GPF (Lim et al. 2021a), and Patchwork (Lim et al. 2021b). As reported in Lim et al. (2021b) and Seo et al. (2022), while GPF, LineFit, and CascadedSeg showed high precision with low recall, R-GPF showed high recall with low precision. In contrast, Patchwork showed both high precision and recall with little perturbation of ground segmentation performance, showing ths highest F_1 -score.

As shown in Fig. 11, the performance of ground segmentation and global registration exhibits a logarithmic relationship. This is because higher F₁-score implies that ground points are precisely and repeatably rejected. By successfully rejecting the ground points from point clouds observed at other viewpoints, ground segmentation with high F₁-score allows global registration to take filtered points whose most parts are consistently overlapped as input. In contrast, if we employ ground segmentation with low F1score, ground points are inconsistently rejected. By doing so, even though the ground points are rejected from the source or target cloud, the ground may not be clearly rejected from the other point cloud. This phenomenon rather increases the number of outliers. For this reason, our approach with other ground segmentation approaches showed lower performance than one with Patchwork.

Therefore, we conclude that exploiting ground segmentation enables global registration to robustly estimate the relative pose between the two point clouds whose pose discrepancy is large. In particular, the experimental evidence supports that ground segmentation with high F_1 -score effectively prevents failure cases by rejecting the gross outliers from the ground points in advance.

7.3. Performance Comparison With State-of-the-Art Methods

In Loop Closing Situations Next, the performance of global registration methods is compared in loop closing situations. As shown in Tables 2 and 3, all the baseline methods showed high success rates once the relative pose between two viewpoints of source and target was sufficiently close ($2 \sim 6 \text{ m case}$). However, as the relative pose between two viewpoints of source and target became distant, our



(c) TEASER++

Fig. 10. Average success rate changes of other global registration methods depending on the absence (w/\circ) and presence (w/) of ground segmentation (GS) module in the loop closing situations of the KITTI dataset where the position discrepancy is between $10\sim12$ m.

Quatro++ showed higher success rate than other global registration methods ($10 \sim 12 \text{ m}$ case). Again, we place more emphasis on the success rate that can directly check whether the global registration methods can successfully estimate the relative pose as an initial alignment even though distant loop pairs are given. Thus, the higher the success rate is, the more suitable method is for performing loop closing.

In addition, it is also noticeable that our preliminary version, Quatro, was on par with TEASER++ (Tables 2 and 3) or showed higher success rate when more imprecise correspondences are given (Table 4). This is because quasi-SO(3) estimation, which reduces the minimum of DoF from three to one, enables to be robust against degeneracy. By doing so, our proposed method prevents catastrophic failure when estimating rotation, unlike TEASER++. For these reasons, Quatro showed higher success rate and smaller errors in the distant cases ($8 \sim 10 \text{ m case}$, Table 5).

Furthermore, some cases where our Quatro++ only succeeded are presented in Fig. 12, demonstrating the robustness of our proposed method against corridor-like environments (the 1st~3rd rows, Fig. 12), three forked roads (the 4th~5th rows, Fig. 12), and intersection or opposing viewpoint scenarios that have large angular

| | Method | Ground | | DCC | | | Riverside | | | KAIST | |
|---------------------------------|---------------|--------------|------------------------------------|------------------------------|-----------------------------|------------------------------|------------------------------|--|------------------------------|------------------------------|------------------------------|
| | | seg. | $2 \sim 6 \mathrm{m}$ | $6\sim 10\text{m}$ | $10\sim 12\text{m}$ | $2m\sim 6m$ | $6\sim 10m$ | $10\sim 12\mathrm{m}$ | $2\sim 6m$ | $6\sim 10m$ | $10\sim 12\mathrm{m}$ |
| 65 m s given) | RANSAC | \checkmark | 22.3 44.5 (22.2 ↑) | 10.2 27.8 (17.6 ↑) | 7.6 24.9 (17.3 ↑) | 15.6 39.3 (23.7 ↑) | 4.8 24.0 (19.2 ↑) | 4.0 17.2 (13.2 ↑) | 25.9 53.7 (27.8 ↑) | 10.9 36.9 (26.0 ↑) | 7.7 23.2 (15.5 ↑) |
| '0.5/0. mprecise ces are | FGR | \checkmark | 50.8 56.5 (5.7 ↑) | 16.4 37.5 (21.1 ↑) | 4.7 24.9 (20.2 ↑) | 56.9 88.6 (31.7 ↑) | 9.1 56.7 (47.6↑) | 1.6 28.7 (27.1 ↑) | 67.1 81.2 (14.1 ↑) | 33.0 76.6 (43.6 ↑) | 12.6 57.1 (44.5 ↑) |
| ns: 0.3/ (more ii sponden | TEASER++ | \checkmark | 95.6 95.4 (0.2 ↓) | 72.8 78.0 (5.2 ↑) | 44.0 52.9 (8.9 ↑) | 82.5 89.8 (7.3 ↑) | 46.5 66.4 (19.9 ↑) | $\begin{array}{c} 22.4 \\ \textbf{41.3} (18.9 \uparrow) \end{array}$ | 89.8 96.8 (7.0↑) | 63.4 86.8 (23.4 ↑) | 39.4 63.0 (23.6↑) |
| Paran corres | Quatro (Ours) | ✓ | 95.9 96.4 (0.5 ↑) | 75.0 83.2 (8.2↑) | 48.0 61.1 (13.1↑) | 82.9 91.6 (8.7 ↑) | 48.1 74.5 (26.4 ↑) | 24.2 49.8 (25.6 ↑) | 90.1 97.0 (6.9 ↑) | 64.4 90.8 (26.4 ↑) | 41.7 71.3 (29.6↑) |
| 25 m iven) | RANSAC | ~ | 76.1 83.6 (7.5 ↑) | 48.4 55.4 (7.0↑) | 32.3 35.3 (3.0↑) | 67.2 81.8 (14.6 ↑) | 44.5 61.9 (17.4 ↑) | 31.7 45.1 (13.4 ↑) | 81.1 89.7 (8.6 ↑) | 65.4 74.5 (9.1 ↑) | 41.1 56.0 (14.9 ↑) |
| 1.5/2.5 precise ses are g | FGR | \checkmark | 64.1 68.8 (4.7 ↑) | 46.8 51.1 (4.3 ↑) | 40.3 44.9 (4.6↑) | 87.9 94.4 (6.5 ↑) | 71.3 84.2 (12.9 ↑) | 63.0 81.0 (18.0 ↑) | 86.9 87.8 (0.9 ↑) | 85.6 86.0 (0.4 ↑) | 80.5 85.7 (5.2↑) |
| ts: 0.6/ (less in pondenc | TEASER++ | \checkmark | 97.9 97.7 (0.2 ↓) | 94.6 94.4 (0.2↓) | 88.3 86.4 (1.9↓) | 92.4 95.1 (2.7 ↑) | 86.6 87.1 (0.5 ↑) | 79.6 80.2 (0.6↑) | 97.3 97.2 (0.1 ↓) | 96.0 94.2 (1.8↓) | 90.9 91.0 (0.1 ↑) |
| Param | Quatro (Ours) | \checkmark | 97.8 97.8 (0.0 -) | 94.7 95.1 (0.4 ↑) | 89.4 89.7 (0.3 ↑) | 92.7 96.1 (3.4↑) | 88.0 92.6 (4.6↑) | 82.1 87.2 (5.1 ↑) | 97.0 96.8 (0.2↓) | 95.7 95.5 (0.3 ↓) | 91.1 93.3 (2.2 ↑) |

Table 4. Comparison of success rate changes before and after the application of ground segmentation (Lim et al. 2021b) on MulRan dataset when more spurious correspondences are given by setting the parameters of FPFH as $\nu = 0.3$ m, $r_{\text{normal}} = 0.5$ m, and $r_{\text{FPFH}} = 0.65$ m (unit: %).



Fig. 11. Influence of the performance of ground segmentation on the success rate of our Quatro. Note that Quatro with Patchwork is equal to Quatro++ (best viewed in color).

Table 5. Comparison of average translation and rotation errors, i.e. t_{avg} and r_{avg} , with the state-of-the-art methods in loop closing situations of Seq. 06 in the KITTI dataset. The ranges in the first row represent position discrepancy boundaries of loop pairs. The bold and the gray-highlight denote the best and the second-best performance, respectively (units for t_{avg} : m, r_{avg} : deg).

| | 0~ | 2 m | 4~ | 6 m | $8\sim 10{ m m}$ | | |
|--|----------------|--------------------|-----------------------|----------------|-----------------------|-----------------------|--|
| Method | $t_{ m avg}$ | $r_{ m avg}$ | $t_{ m avg}$ | $r_{\rm avg}$ | $t_{ m avg}$ | r_{avg} | |
| RANSAC | 2.369 | 14.22 | 5.010 | 27.35 | 8.341 | 33.34 | |
| FGR | 0.057 | 0.222 | 0.103 | 0.301 | 1.821 | 1.828 | |
| TEASER++ | 0.070 | 0.285 | 0.131 | 0.481 | 0.498 | 1.469 | |
| Quatro (Ours) | 0.067 | 0.324 | 0.120 | 0.465 | 0.471 | 0.724 | |
| Quatro++ (Ours) | 0.078 | 0.338 | 0.110 | 0.452 | 0.163 | 0.561 | |
| Quatro w/ INS (Ours) Quatro++ w/ INS (Ours) | 0.059 0.067 | 0.207 0.208 | 0.101 0.089 | 0.230 0.222 | 0.429 0.117 | 0.346 0.243 | |

discrepancies between the viewpoints of source and target clouds (the 6th \sim 8th rows, Fig. 12).

The reasons why other state-of-the-art methods failed are presumed as follows. As mentioned earlier, as the pose discrepancy between two viewpoints becomes farther, the number of true inliers decreases and the ratio of outliers increases simultaneously. Under this circumstance, RANSAC is highly likely to fail to estimate correct relative pose because RANSAC is relatively weak to the outliers. In fact, it was reported that RANSAC often fails once the ratio of outliers is over 60% (Tzoumas et al. 2019).

FGR showed better performance compared with RANSAC because FGR was based on GNC. Unfortunately, FGR linearizes SE(3) during the optimization, yet this linearization is sometimes not valid once the pose discrepancy becomes larger, resulting in performance degradation ($8 \sim 10$ m case, Table 5). In addition, FGR has no graph-based pruning module, i.e. MCIS, so that FGR cannot tolerate gross outliers.

On the other hand, TEASER++ showed a promising performance even in distant cases because TEASER++ also utilizes the graph-based outlier pruning method. However, the point is that MCIS also occasionally rejects too many correspondences by considering them as outliers, potentially leading to degeneracy in the GNC-based SO(3) rotation estimation step. As a result, the rotation estimation output flipped (the 3rd, 5th, and 8th rows, Fig. 12) or tilted results (the other rows, Fig. 12) owing to the degeneracy. Note that TEASER++ is based on the decoupling of rotation and translation, so the failure of rotation estimation inevitably brings failure of translation.

Deep learning-based methods showed a substantial and inconsistent degradation in performance as the viewpoint distance between the source and target increased. This is because these learning-based approaches were only trained to predict the relative pose in cases where the viewpoint difference between the source and target is sufficiently small and rarely trained in distant loop cases. Furthermore, as presented in Table 3, the performance of deep learning-based methods was more significantly degraded than our approach when the surroundings or sensor configuration was different from the training data. These results support our claim that our approach is more suitable for loop closing and applicable to various environments because our approach does not need any training procedure.



Fig. 12. Qualitative comparison in distant cases, where only our approach succeeded in performing registration of two partially-overlapped point clouds in the KITTI dataset. Even though distant source (red) and target (green) clouds are given, our proposed method showed tightly aligned results, overcoming the sparsity and degeneracy issues. The closer the warped source cloud (blue) and target cloud are, the better. The solid black boxes are zoomed views that highlight the misalignment of other state-of-the-art methods. The left-top texts represent the initial pose discrepancy in each scene. Note that the RANSAC returns the identity matrix, i.e. I_4 , in the failure cases, so the estimates of RASNAC are identical to the source clouds (best viewed in color).

In Odometry Test Next, the registration methods were tested in the odometry test. Our proposed method is also compared with the conventional odometry methods and deep learning-based methods. Note that global registration basically does not aim to be used as an odometry, but we

can check which method shows consistent performance by analyzing trajectory errors. In other words, better global registration shows better odometry performance once the frame interval Δ is larger than one, i.e. $\Delta = 3$ and $\Delta = 5$ in Table 6.

Table 6. Comparison of odometry test with the state-of-the-art methods on Seq. 00 of the KITTI dataset. Δ denotes the frame interval; for instance, $\Delta = \delta$ means that the $(i + \delta)$ -th and *i*-th point clouds are taken as the source and target clouds, respectively. The suffix c2f means that the global registration is employed as an initial alignment followed by the local registration as a fine alignment. The results of deep learning-based methods are from Li and Wang (2020) and Fischer et al. (2021) (units for $t_{\rm rel}$: %, $r_{\rm rel}$: deg/100m).

| _ | | Δ : | = 1 | $\Delta =$ | = 3 | $\Delta = 5$ | |
|--------|---|--|--------------------------------------|----------------------------------|----------------------------------|---------------------------|----------------------------------|
| | Method | $t_{\rm rel}$ | $r_{\rm rel}$ | $t_{\rm rel}$ | $r_{\rm rel}$ | $t_{\rm rel}$ | $r_{\rm rel}$ |
| Local | ICP | 6.88 | 2.99 | 21.92 | 8.70 | 21.14 | 8.51 |
| | G-ICP | 1.26 | 0.45 | 5.50 | 1.45 | 14.20 | 3.32 |
| | VGICP | 1.03 | 0.30 | 11.83 | 1.65 | 19.11 | 6.32 |
| Global | FGR | 2.73 | 0.69 | 7.17 | 1.58 | 14.66 | 4.12 |
| | TEASER++ | 2.11 | 0.91 | 2.64 | 1.11 | 3.19 | 0.91 |
| | Quatro (Ours) | 1.45 | 0.41 | 1.38 | 0.24 | 1.94 | 0.46 |
| | Quatro++ (Ours) | 1.90 | 0.53 | 1.45 | 0.32 | 0.99 | 0.28 |
| Deep | $\label{eq:LO-Net} \begin{array}{c} LO-Net\\ LO-Net+M\\ DMLO^{\dagger}\\ DMLO+M^{\dagger}\\ A-LOAM+StickyPillars^{\dagger} \end{array}$ | 1.47 0.78 0.83 0.73 0.65 | 0.72 0.42 0.44 0.44 0.26 | N/A N/A N/A N/A 0.79 | N/A N/A N/A N/A 0.31 | N/A N/A N/A 1.29 | N/A N/A N/A N/A 0.48 |
| Conv. | SuMa | 0.68 | 0.23 | 1.69 | 0.61 | 2.36 | 0.51 |
| | A-LOAM | 0.70 | 0.27 | 0.97 | 0.38 | 31.16 | 12.10 |
| | Quatro-c2f (Ours) | 0.65 | 0.21 | 0.67 | 0.21 | 0.67 | 0.21 |
| | Quatro++-c2f (Ours) | 0.68 | 0.23 | 0.61 | 0.21 | 0.50 | 0.20 |

†: Seq. 00 is used for training of the network

As baseline methods, we employ three local registration methods: ICP (Besl and McKay 1992), G-ICP (Segal et al. 2009), and VGICP (Koide et al. 2021). Furthermore, we compare our approach with learning-based odometry methods: LO-Net (Li et al. 2019), DMLO (Li and Wang 2020), and A-LOAM with StickyPillars (Fischer et al. 2021). Note that the codes of deep learning-based odometry approaches are not available, so we employ the results from the original papers. We also employ conventional LiDAR odometry methods: SuMa (Behley and Stachniss 2018) and A-LOAM (Zhang and Singh 2014).

As shown in Table 6, local registration methods showed promising performance only if the frame interval is not large, i.e. $\Delta = 1$. These results make sense because the pose discrepancy of two viewpoints is sufficiently within the coverage of the local registration, so their estimate can successfully converge into the global optimum. However, as the interval of frames gets larger, the performance of local registration is dramatically degraded because their strong assumption that the nearest point pairs between the source and target clouds are correlated, does not hold. As a result, VGICP (Koide et al. 2021), which is one of the state-of-theart local registration methods, showed catastrophic failure of odometry estimation, as shown in Fig. 13(b).

In terms of global registration, we stress two experimental observations. First, our Quatro and Quatro++ especially showed less decreased performance than the other global registration methods, which supports that forgoing the estimation of roll and pitch angles makes the algorithms more robust and consistent. That is, the trajectories of other global registration methods were likely to be twisted owing to undesirable pitch and roll errors, whereas our proposed method showed more consistent trajectories, as seen from the xz-viewpoint in Fig. 13. Put differently, the trajectories from the other global registration methods showed the large

discrepancy along the z-axis, whereas the result from our Quatro++ was closer to the ground truth.

Second, using ground segmentation as a preprocessing step helps the global registration avoid local minima because Quatro++ showed a promising performance in the case where $\Delta = 5$ in Table 6. Interestingly, when the intervals were small, Quatro showed lower errors than Quatro++. It can be interpreted that once the pose discrepancy between two viewpoints is not large, ground points rather help to estimate the relative pose because the ground points of source and target are mostly overlapped, leading to more true inliers.

However, these errors do not directly affect the performance of fine alignment and thus are negligible for the local registration in $\Delta = 1$ and $\Delta = 3$ cases. That is, there is an insignificant performance difference when Quatro or Quatro++ is utilized as an initial alignment when $\Delta = 1$ and $\Delta = 3$, showing similar level of errors in Table 6. Moreover, Quatro++ still robustly outputs the estimate to transform the viewpoint of the source cloud into the boundary within the narrow convergence region of the local registration. Consequently, Quatro++-c2f showed a promising performance when $\Delta = 5$. Furthermore, Quatro++-c2f showed better performance compared with the state-of-the-art methods, even including conventional and deep learning-based approaches. It was remarkable that though some deep learning-based methods were trained using whole KITTI sequences including Seq. 00, our proposed method showed a competing performance without any prior knowledge.

Therefore, the suitability of our Quatro++ was demonstrated from the perspective of the coarse-to-fine alignment, which is our ultimate goal to achieve successful loop closing, helping local registration algorithms perform the fine alignment.

In Augmented Rotation Situations Next, the performance was tested in augmented rotation situations in NAVER LABS localization dataset. The augmented rotation situations mean that the source cloud is additionally rotated along the yaw direction. By doing so, the feasibility of the algorithms on the opposing viewpoint or intersection cases could be checked.

As shown in Fig. 14, our proposed method showed lower errors and maintained the level of errors even though the augmented angle reached 180°. FGR showed better performance if the augmented angle was not large; however, as the augmented angle became larger, the performance of FGR was degraded because the linearization assumption did not hold.

TEASER++ showed successful and robust registration results, yet it yielded undesirable rotation errors, showing the offset between our proposed method and TEASER++ in Fig. 14. In other words, the estimate by TEASER++ includes unintended roll and pitch errors. In contrast, our Quatro++ showed rather precise performance because our assumption that $\angle(\mathbf{R}_y \cdot \mathbf{R}_x)$ can be approximated as the identity matrix is met in most indoor environments. In doing so, our quasi-SO(3) estimation prevents these roll and pitch errors, as shown in Fig. 15.

Therefore, it was proven that our proposed method is also advantageous in indoor situations where the ground is mostly



Fig. 13. (L-R): Qualitative comparison of trajectories by the state-of-the-art registration methods in Seq. 00 and Seq. 05 of the KITTI dataset. Δ denotes the frame interval; for instance, $\Delta = \delta$ means that the $(i + \delta)$ -th and *i*-th point clouds are taken as the source and target clouds, respectively. Our proposed methods showed more consistent results although the distant and partially overlapped two point clouds are given. In contrast, the trajectories of other global registration methods were likely to be twisted owing to the undesirable pitch and roll errors and some failure situations.



Fig. 14. Performance changes with respect to the augmented yaw rotation when the frame interval is set to 5 (i.e. $\Delta = 5$) in HD_B1 of the NAVER LABS localization dataset (best viewed in color).



Fig. 15. Example of registration when large rotation discrepancy exists in HD_B1 of the NAVER LABS localization dataset. (c) TEASER++ showed successful and robust registration, yet it yielded the undesirable angular error, highlighted as the red double arrow in the side view. The solid red and green boxes indicate the algorithms failed and succeeded, respectively (best viewed in color).

flat and the non-ground objects are mostly orthogonal to the ground.

7.4. Computational Cost of Modules in Quatro++

Here, the computational costs with respect to the feature extraction/matching and optimization of our global registration were analyzed. As shown in Fig. 16, the application of ground segmentation effectively reduced the time taken to set correspondences. In addition, as shown in Fig. 17, the optimization part in our proposed method showed fastest speed compared with other methods.

Therefore, the total time of Quatro++, which is the summation of the time taken for preprocessing, correspondence estimation, and Quatro, takes less than one second, so we confirm that our proposed method can be used in the loop closing module of SLAM frameworks, which requires over 1 Hz.

Fig. 16. (L-R): The average time for each module in the preprocessing and correspondence estimation on Intel(R) Core(TM) i7-7700K and Intel(R) Core(TM) i9-13900. Note that the application of ground segmentation effectively reduces the time taken to set correspondences.



Fig. 17. Average optimization time in the KITTI dataset and the NAVER LABS localization dataset on Intel(R) Core(TM) i9-9900KF (Avg. of our Quatro: 5.0 and 6.4 ms, respectively).



Fig. 18. Registration results in a hand-held system in Oxford-Hilti dataset. Once the yaw and pitch angles are compensated by an INS system, our proposed method is also effective on a hand-held system. The blue and green points denote the warped source and target clouds, respectively (best viewed in color).

7.5. Application I: Leveraging an INS in Non-Flat Regions

Our Quatro++ (or Quatro) sacrifices roll and pitch angles estimation to be more robust. However, Quatro++ and Quatro can achieve more precise initial alignment results by utilizing the INS measurements to compensate roll and pitch angles. This compensation means that the approximated $\mathbf{R}_{y} \cdot \mathbf{R}_{x}$ is replaced by the rotation from the roll and pitch angles measured by INS, which is explained in Section 5.5. As presented in Table 5, even though the raw measurements by the INS system were used, the pitch and roll compensation using these measurements significantly reduced the rotation errors, which is followed by the improvement of translation estimation. This is because our proposed method is based on the decoupling method, which estimates the relative rotation followed by translation estimation; thus, the quality of rotation estimation directly affects the quality of translation estimation.

Not only for the improvement in performance, but we also stress that our proposed method can easily utilize these INS measurements without any additional fusion technologies, such as weighted average; thus, our proposed method is more INS-friendly than other methods. One may argue that other methods can also improve performance by using INS measurements. However, other methods require an additional manner to appropriately average the estimated rotation and



Fig. 19. (a)-(b) Mapping results of LeGO-LOAM without and with the application of our proposed method as a loop closing in Seq. 05 of the KITTI dataset. Note that all the parameters of LeGO-LOAM are set to be the same.

the rotation measured by the INS. Even though the weighted average is performed, this sensor fusion technique does not let the failed estimates successfully go into the narrow convergence region, still showing large pose errors. In contrast, as mentioned in Section 5.5, the INS measurements can be easily exploitable for our proposed method by substituting the approximated roll and pitch rotations with those from INS measurements. By doing so, our proposed method showed lower errors while preserving higher success rates.

In addition, we conducted a feasibility study on a dataset made by a hand-held system. As shown in Fig. 18, even though the axes were rotated, estimation of $\hat{\mathbf{R}}_{\text{INS}}$, which was explained in Section 5.5, followed by our proposed method showed successful registration results. Therefore, it was demonstrated that our proposed method is also effective on a hand-held system once roll and pitch angles are compensated as an initial guess of our global registration.

Therefore, these results demonstrated the feasibility of our proposed method in INS systems where the attitude of xy-plane of the sensor frame is not likely fixed*.

7.6. Application II: Quatro++ in LiDAR SLAM Frameworks

So far, we investigated performance of a single registration given a pair of point clouds. Here, the impact of Quatro++ to the LiDAR SLAM was analyzed. Note that our Quatro++ can be employed in keyframe-based pose-graph SLAM systems.

First, we compared the performance before and after the application of our global registration into LeGO-LOAM (Shan and Englot 2018). LeGO-LOAM utilizes the radius search to find loop candidates by using the current position as a query, which is followed by ICP. However, as mentioned in Section 5.3, ICP usually fails to estimate the correct relative pose in distant cases, rejecting many loop candidates. That is, even though the radius search finds the actual loop, the loop candidate whose pose discrepancy is large is rejected owing to the narrow convergence region of the local registration methods, resulting in a high MSE. Consequently, these loops are not included in the graph structure, leaving trajectory errors (Fig. 19(a)).

In contrast, LeGO-LOAM with our proposed method showed a more precise mapping result (Fig. 19(b)). It was observed that the number of false positive rejections can be significantly reduced and the quality of loop constraints is improved when our Quatro++ provides an initial alignment to the local registration. As a result, even if a large drift is accumulated through the large loop, the global trajectory errors can be successfully minimized by leveraging more accurate and abundant loop constraints.

The positive impact on the quality of mapping becomes more distinguishable and significant in large-scale environments. In this experiment, three types of LeGO-LOAM were compared: the original LeGO-LOAM, SC-LeGO-LOAM that exploits ScanContext (Kim and Kim 2018) as a loop detection, TSC-LeGO-LOAM that employs both TEASER++ and ScanContext, and QSC-LeGO-LOAM that utilizes both our proposed method (Quatro++) and ScanContext.

As shown in Fig. 20, LeGO-LOAM without any loop detection and closing approaches showed imprecise mapping results. In large-scale environments, the radius search could not find potential loop candidates owing to the large pose drift. For this reason, loop closing was performed only locally, so its global trajectory errors were not reduced by PGO (Fig. 20(a)).

The application of the loop detection module, i.e. SC-LeGO-LOAM, showed a more precise map by finding additional loop candidates that could not be searched solely by the radius search. However, some trajectory errors still remained due to the undesirable false positive loops. This is because the loop detection sometimes wrongly estimates two geometrically similar scenes as a loop candidate. For instance, the scenes where the buildings are densely

https://hilti-challenge.com/leader-board-2023.html

^{*}Consequently, we achieved the 4th and 1st place by utilizing our proposed method as a loop closing module in the 2022 and 2023 Hilti SLAM challenges, respectively. Please refer to the following site:



Fig. 20. (T-B): Point cloud mapping results and the corresponding loop constraints in DCC01 of the MulRan dataset. (a) The original LeGO-LOAM. (b) LeGO-LOAM with ScanContext as a loop detection. (c) LeGO-LOAM with both TEASER++ and ScanContext, named as TSC-LeGO-LOAM. (d) LeGO-LOAM with both Quatro++ and ScanContext, named as QSC-LeGO-LOAM, showing a tightly aligned mapping result and fewer false loops. The green lines denote the true positive loop constraints and red lines denote the false positive loop constraints (best viewed in color).



Fig. 21. Visualized loop constraints of (a) The original LeGO-LOAM, (b) LeGO-LOAM with ScanContext as a loop detection and (c) TSC-LeGO-LOAM, and (d) QSC-LeGO-LOAM in Riverside01 of the MulRan dataset. Note that QSC-LeGO-LOAM reduced false positive loops while increasing true positive loops simultaneously. The green lines denote the true positive loop constraints and red lines denote the false positive loop constraints (best viewed in color).

located on the lateral side of the robot have ambiguous geometrical characteristics. Accordingly, sometimes distant



Fig. 22. Box plot of absolute pose errors after the application of ScanContext (SC) as a loop detection, both TEASER++ and ScanContext (TSC), and both our Quatro++ and ScanContext (QSC). The **** annotations indicate measurements with *p*-value $< 10^{-4}$ after a paired *t*-Test.

 Table 7.
 Absolute pose errors of SLAM results in the MulRan dataset.

 For all the metrics, the lower, the better.
 Image: Comparison of the sector of

| | Algorithm | Mean | Median | RMSE | Stdev. |
|-------------|----------------------|--------------|--------------|--------------|--------------|
| DCC01 | LeGO-LOAM | 28.89 | 29.21 | 32.37 | 14.60 |
| | SC-LeGO-LOAM | 6.13 | 5.59 | 6.46 | 1.89 |
| | TSC-LeGO-LOAM | 6.10 | 5.49 | 5.56 | 2.58 |
| | QSC-LeGO-LOAM (Ours) | 5.65 | 5.33 | 5.97 | 1.92 |
| KAIST02 | LeGO-LOAM | 24.21 | 18.48 | 30.25 | 18.14 |
| | SC-LeGO-LOAM | 4.28 | 3.72 | 5.33 | 3.17 |
| | TSC-LeGO-LOAM | 4.27 | 3.85 | 5.12 | 2.82 |
| | QSC-LeGO-LOAM (Ours) | 3.88 | 2.71 | 4.86 | 2.93 |
| Riverside01 | LeGO-LOAM | 96.60 | 62.57 | 124.94 | 79.27 |
| | SC-LeGO-LOAM | 26.87 | 28.99 | 30.59 | 26.87 |
| | TSC-LeGO-LOAM | 21.23 | 21.23 | 32.89 | 19.50 |
| | QSC-LeGO-LOAM (Ours) | 20.19 | 19.57 | 23.01 | 19.57 |

two point clouds measured in the corridor-like environments or opposing viewpoint cases are considered as the same place. The problem is that the MSEs of these scenes are likely to have small values, so some false loops are not rejected via the MSE-based thresholding. Therefore, these wrong loops are formed on the graph structure, resulting in the misaligned map (Fig. 20(b)).

In contrast, the application of our proposed method in the loop closing showed a tightly aligned mapping result and fewer false loops, as shown in Figs. 20(d) and 21(d). As explained earlier, our proposed method showed more robust performance against these ambiguous scenes and reverse cases by leveraging ground segmentation. Accordingly, even though distant or reversed loop pairs are given by the loop detection module, our proposed method can successfully estimate the relative pose in a coarse manner. Consequently, the local registration method, which follows the global registration, can estimate the correct relative pose as a fine alignment. In particular, we also highlight that our proposed method increases the number of true positive loops in widely opened environments, showing more dense green lines in Fig. 21(d) compared with Fig. 21(b).

Certainly, TSC-LeGO-LOAM also increases the number of loop pairs (Fig. 21(c)), yet TEASER++ triggers some false positive loops. This is because TEASER++ showed the lower success rates in the riverside scenes, as presented in Table 3, so TEASER++ occasionally let the fine alignment by local registration converge to the local minimum by providing a wrong initial guess as a coarse alignment. Compared with TSC-LeGO-LOAM, our QSC-LeGO-LOAM only led to few false positive loop pairs due to the robustness of Quatro++ against the outliers and sparsity issues, showing better SLAM performance (Fig. 22 and Table 7). In addition, by leveraging its robustness, Quatro++ can perform registration as a loop closing even when moving objects, which may lead to erroneous feature matching, are included (Fig. 23).

Through these experiments, we demonstrate that our proposed method has positive effect on the LiDAR SLAM: our proposed method increases the quality of loop constraints by reducing both false positive and false negative loops, reducing the pose errors, as presented in Fig. 22 and Table 7.

8. Conclusion

In this study, a robust global registration method, *Quatro++*, has been proposed. Our proposed method was proven to be more robust against the sparsity and degeneracy issues than the state-of-the-art global registration methods by utilizing ground segmentation. Furthermore, we applied our proposed method to the loop closing module in LiDAR SLAM and confirmed that the quality of the loop constraints was improved, showing more precise mapping results.

Although we demonstrated that our proposed method outperformed the state-of-the-art methods in terms of global registration and loop closing in LiDAR SLAM, there are some limitations. That is, we placed more emphasis on the back-end part of the correspondence-based global registration, so the improvement of feature extraction and matching was beyond our scope. In future works, we plan to improve the front-end of the correspondence-based global registration. Accordingly, the way to increase the quality of the estimated correspondences should be addressed to tackle the sparsity and degeneracy issues in terms of the front-end part as well.

9. Funding

This work was financially supported in part by Korea Evaluation Institute of Industrial Technology (KEIT) funded by the Korea Government (MOTIE) under Grant No.20018216, Development of mobile intelligence SW for autonomous navigation of legged robots in dynamic and atypical environments for real application, and the Unmanned Swarm CPS Research Laboratory Program of Defense Acquisition Program Administration and Agency for Defense Development (UD220005VD). The researchers are supported by BK21 FOUR.

References

- Aldoma A, Marton ZC, Tombari F, Wohlkinger W, Potthast C, Zeisl B, Rusu RB, Gedikli S and Vincze M (2012) Tutorial: Point cloud library: Three-dimensional object recognition and 6 DoF pose estimation. *IEEE Robotics & Automation Magazine* 19(3): 80–91.
- Ao S, Hu Q, Yang B, Markham A and Guo Y (2021) SpinNet: Learning a general surface descriptor for 3D point cloud



Fig. 23. Registration results of Quatro++ in distant cases (T-B): Registration results between 467 and 4,175; 691 and 4,887 frames in DCC01 of the MulRan dataset. (a) Even though the partially overlapped source (red) and target (green) clouds, which contain moving objects, are given, (b) our proposed method showed tightly aligned results. The closer the warped source cloud (blue) and target cloud are, the better. (c) The solid black circles highlight the existence of moving objects and rainbow colors indicate the closest distances between the points of source and target clouds. The top-left texts represent the initial pose discrepancy in each scene (best viewed in color).

registration. In: *IEEE/CVF Conference on Computer Vision* and Pattern Recognition. pp. 11753–11762.

- Arun KS, Huang TS and Blostein SD (1987) Least-squares fitting of two 3D point sets. *IEEE Transactions on Pattern Analysis* and Machine Intelligence (5): 698–700.
- Ashbrook A, Fisher RB, Robertson C and Werghi N (1998) Finding surface correspondence for object recognition and registration using pairwise geometric histograms. In: *European Conference* on Computer Vision. pp. 674–686.
- Bai X, Luo Z, Zhou L, Chen H, Li L, Hu Z, Fu H and Tai CL (2021) PointDSC: Robust point cloud registration using deep spatial consistency. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 15859–15869.
- Behley J and Stachniss C (2018) Efficient surfel-based SLAM using 3D laser range data in urban environments. In: *Robotics: Science and Systems*.
- Belongie S, Malik J and Puzicha J (2002) Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24(4): 509–522.
- Bernreiter L, Ott L, Nieto J, Siegwart R and Cadena C (2021) PHASER: A robust and correspondence-free global point cloud registration. *IEEE Robotics and Automation Letters* 6(2): 855– 862.
- Besl PJ and McKay ND (1992) Method for registration of 3D shapes. In: Sensor Fusion IV: Control Paradigms and Data Structures. pp. 586–606.
- Black MJ and Rangarajan A (1996) On the unification of line processes, outlier rejection, and robust statistics with applications in early vision. *International Journal of Computer Vision* 19(1): 57–91.
- Briales J and Gonzalez-Jimenez J (2017) Convex global 3D registration with Lagrangian duality. In: *IEEE/CVF*

Conference on Computer Vision and Pattern Recognition. pp. 4960–4969.

- Brown M, Windridge D and Guillemaut JY (2019) A family of globally optimal branch-and-bound algorithms for 2D-3D correspondence-free registration. *Pattern Recognition* 93: 36– 54.
- Carlone L, Calafiore GC, Tommolillo C and Dellaert F (2016) Planar pose graph optimization: Duality, optimal solutions, and verification. *IEEE Transaction on Robotics* 32(3): 545–565.
- Cattaneo D, Vaghi M and Valada A (2022) LCDNet: Deep loop closure detection and point cloud registration for LiDAR SLAM. *IEEE Transactions on Robotics* 38(4): 2074–2093.
- Chen X, Läbe T, Milioto A, Röhling T, Behley J and Stachniss C (2022) OverlapNet: A siamese network for computing LiDAR scan similarity with applications to loop closing and localization. *Autonomous Robots* 46(1): 61–81.
- Chetverikov D, Svirko D, Stepanov D and Krsek P (2002) The trimmed iterative closest point algorithm. In: Object Recognition Supported by User Interaction for Service Robots. pp. 545–548.
- Cho YK and Gai M (2014) Projection-recognition-projection method for automatic object recognition and registration for dynamic heavy equipment operations. *Journal of Computing in Civil Engineering* 28(5): A4014002.
- Choi S, Kim T and Yu W (1997) Performance evaluation of RANSAC family. *Journal of Computer Vision* 24(3): 271–300.
- Choy C, Dong W and Koltun V (2020) Deep global registration. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2514–2523.
- Chua CS and Jarvis R (1996) 3D free-form surface registration and object recognition. *International Journal of Computer Vision* 17(1): 77–99.

- Chum O, Matas J and Kittler J (2003) Locally optimized RANSAC. In: *Joint Pattern Recognition Symposium*. pp. 236–243.
- Curless B and Levoy M (1996) A volumetric method for building complex models from range images. In: *Conference on Computer Graphics and Interactive Techniques*. pp. 303–312.
- Dellenbach P, Deschaud JE, Jacquet B and Goulette F (2022) CR-ICP: Real-time elastic LiDAR odometry with loop closure. In: *IEEE International Conference on Robotics and Automation*. pp. 5580–5586.
- Deng H, Birdal T and Ilic S (2018) PPF-FoldNet: Unsupervised learning of rotation invariant 3D local descriptors. In: *European Conference on Computer Vision*. pp. 602–618.
- Dong Z, Yang B, Liu Y, Liang F, Li B and Zang Y (2017) A novel binary shape context for 3D local surface description. *ISPRS Journal of Photogrammetry and Remote Sensing* 130: 431– 452.
- Du J, Wang R and Cremers D (2020) DH3D: Deep hierarchical 3D descriptors for robust large-scale 6DoF relocalization. In: *European Conference on Computer Vision*. pp. 744–762.
- Fischer K, Simon M, Olsner F, Milz S, Gross HM and Mader P (2021) StickyPillars: Robust and efficient feature matching on point clouds using graph neural networks. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 313–323.
- Fischler MA and Bolles RC (1981) Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* 24(6): 381–395.
- Geiger A, Lenz P, Stiller C and Urtasun R (2013) Vision meets robotics: The KITTI dataset. *International Journal of Robotics Research* 32(11): 1231–1237.
- Geiger A, Lenz P and Urtasun R (2012) Are we ready for autonomous driving? The KITTI vision benchmark suite. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3354–3361.
- Glent Buch A, Yang Y, Kruger N and Gordon Petersen H (2014) In search of inliers: 3D correspondence by local and global voting. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2067–2074.
- Gojcic Z, Zhou C, Wegner JD, Guibas LJ and Birdal T (2020) Learning multiview 3D point cloud registration. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1759–1769.
- Gojcic Z, Zhou C, Wegner JD and Wieser A (2019) The perfect match: 3D point cloud matching with smoothed densities.In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5545–5554.
- Greenspan M and Godin G (2001) A nearest neighbor method for efficient ICP. In: *IEEE International Conference on 3-D Digital Imaging and Modeling*. pp. 161–168.
- Grisetti G, Kümmerle R, Stachniss C and Burgard W (2010) A tutorial on graph-based SLAM. *IEEE Intelligent Transportation Systems Magazine* 2(4): 31–43.
- Hartley RI and Kahl F (2009) Global optimization through rotation space search. *International Journal of Computer Vision* 82(1): 64–79.
- He D, Xu W and Zhang F (2021) Kalman filters on differentiable manifolds. *arXiv preprint arXiv:2102.03804*.

- Himmelsbach M, Hundelshausen FV and Wuensche HJ (2010) Fast segmentation of 3D point clouds for ground vehicles. In: *IEEE Intelligent Vehicles Symposium (IV)*. pp. 560–565.
- Horn BK (1987) Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America A* 4(4): 629–642.
- Kim G and Kim A (2018) Scan Context: Egocentric spatial descriptor for place recognition within 3D point cloud map. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 4802–4809.
- Kim G, Park YS, Cho Y, Jeong J and Kim A (2020) MulRan: Multimodal range dataset for urban place recognition. In: *IEEE International Conference on Robotics and Automation*. pp. 6246–6253.
- Kim H, Song S and Myung H (2019) GP-ICP: Ground plane ICP for mobile robots. *IEEE Access* 7: 76599–76610.
- Koide K, Yokozuka M, Oishi S and Banno A (2021) Voxelized GICP for fast and accurate 3D point cloud registration. In: *IEEE International Conference on Robotics and Automation*. pp. 11054–11059.
- Komorowski J (2021) MinkLoc3D: Point cloud based large-scale place recognition. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 1790–1799.
- Lawler EL and Wood DE (1966) Branch-and-bound methods: A survey. *Operations Research* 14(4): 699–719.
- Lee D, Ryu S, Yeon S, Lee Y, Kim D, Han C, Cabon Y, Weinzaepfel P, Guérin N, Csurka G et al. (2021) Large-scale localization datasets in crowded indoor spaces. In: *IEEE/CVF Conference* on Computer Vision and Pattern Recognition. pp. 3227–3236.
- Lei H, Jiang G and Quan L (2017) Fast descriptors and correspondence propagation for robust global point cloud registration. *IEEE Transactions on Image Processing* 26(8): 3614–3623.
- Leordeanu M and Hebert M (2005) A spectral technique for correspondence problems using pairwise constraints. PhD Thesis, Carnegie Mellon University, USA.
- Li H (2009) Consensus set maximization with guaranteed global optimality for robust geometry estimation. In: *IEEE/CVF International Conference on Computer Vision*. pp. 1074–1080.
- Li K, Li M and Hanebeck UD (2021) Towards high-performance solid-state-LiDAR-inertial odometry and mapping. *IEEE Robotics and Automation Letters* 6(3): 5167–5174.
- Li Q, Chen S, Wang C, Li X, Wen C, Cheng M and Li J (2019) Lo-Net: Deep real-time LiDAR odometry. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8473–8482.
- Li Z and Wang N (2020) DMLO: Deep matching LiDAR odometry. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems.* pp. 6010–6017.
- Lim H, Hwang S and Myung H (2021a) ERASOR: Egocentric ratio of pseudo occupancy-based dynamic object removal for static 3D point cloud map building. *IEEE Robotics and Automation Letters* 6(2): 2272–2279.
- Lim H, Hwang S, Shin S and Myung H (2020) Normal Distributions Transform is enough: Real-time 3D scan matching for pose correction of mobile robot under large odometry uncertainties.
 In: *International Conference on Control, Automation, and Systems.* pp. 1155–1161.

- Lim H, Oh M and Myung H (2021b) Patchwork: Concentric zonebased region-wise ground segmentation with ground likelihood estimation using a 3D LiDAR sensor. *IEEE Robotics and Automation Letters* 6(4): 6458–6465.
- Lim H, Yeon S, Ryu S, Lee Y, Kim Y, Yun J, Jung E, Lee D and Myung H (2022) A single correspondence is enough: Robust global registration to avoid degeneracy in urban environments. In: *IEEE International Conference on Robotics* and Automation. pp. 8010–8017.
- Lin YK, Lin WC and Wang CC (2022) K-closest points and maximum clique pruning for efficient and effective 3-D Laser scan matching. *IEEE Robotics and Automation Letters* 7(2): 1471–1477.
- Maron H, Dym N, Kezurer I, Kovalsky S and Lipman Y (2016) Point registration via efficient convex relaxation. *ACM Transactions on Graphics* 35(4): 1–12.
- Narksri P, Takeuchi E, Ninomiya Y, Morales Y, Akai N and Kawaguchi N (2018) A slope-robust cascaded ground segmentation in 3D point cloud for autonomous vehicles. In: *IEEE International Conference on Intelligent Transportation Systems.* pp. 497–504.
- Olsson C, Kahl F and Oskarsson M (2008) Branch-andbound methods for Euclidean registration problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31(5): 783–794.
- Pais GD, Ramalingam S, Govindu VM, Nascimento JC, Chellappa R and Miraldo P (2020) 3DRegNet: A deep neural network for 3D point registration. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition.*
- Pan J, Min Z, Zhang A, Ma H and Meng MQH (2019) Multiview global 2D-3D registration based on branch and bound algorithm. In: *IEEE International Conference on Robotics and Biomimetics*. pp. 3082–3087.
- Papazov C, Haddadin S, Parusel S, Krieger K and Burschka D (2012) Rigid 3D geometry matching for grasping of known objects in cluttered scenes. *International Journal of Robotics Research* 31(4): 538–553.
- Pomerleau F, Colas F, Siegwart R and Magnenat S (2013) Comparing ICP variants on real-world data sets. *Autonomous Robots* 34(3): 133–148.
- Qin T, Li P and Shen S (2018) VINS-Mono: A robust and versatile monocular visual-inertial state estimator. *IEEE Transaction on Robotics* 34(4): 1004–1020.
- Ramezani M, Khosoussi K, Catt G, Moghadam P, Williams J, Borges P, Pauling F and Kottege N (2022) Wildcat: Online continuous-time 3D LiDAR-inertial SLAM. arXiv preprint arXiv:2205.12595.
- Rodolà E, Albarelli A, Bergamasco F and Torsello A (2013) A scale independent selection process for 3D object recognition in cluttered scenes. *International Journal of Computer Vision* 102(1): 129–145.
- Rossi RA, Gleich DF, Gebremedhin AH, Patwary MMA and Ali M (2015) Parallel maximum clique algorithms with applications to network analysis and storage. *SIAM Journal on Scientific Computing* 37(5): 589–616.
- Rouhani M and Sappa AD (2011) Correspondence free registration through a point-to-model distance minimization. In: *IEEE International Conference on Computer Vision*. pp. 2150–2157.
- Rusinkiewicz S and Levoy M (2001) Efficient variants of the ICP algorithm. In: *IEEE International Conference on 3-D Digital*

Imaging and Modeling. pp. 145-152.

- Rusu RB, Blodow N and Beetz M (2009) Fast point feature histograms (FPFH) for 3D registration. In: *IEEE International Conference on Robotics and Automation*. pp. 3212–3217.
- Rusu RB and Cousins S (2011) 3D is here: Point cloud library (PCL). In: *IEEE International Conference on Robotics* and Automation. pp. 1–4.
- Scaramuzza D (2011) 1-Point-RANSAC structure from motion for vehicle-mounted cameras by exploiting non-holonomic constraints. *International Journal of Computer Vision* 95(1): 74–85.
- Schnabel R, Wahl R and Klein R (2007) Efficient RANSAC for point-cloud shape detection. In: *Computer Graphics Forum*, volume 26. pp. 214–226.
- Segal A, Haehnel D and Thrun S (2009) Generalized-ICP. In: *Robotics: Science and Systems*, 4. pp. 435–442.
- Seo DU, Lim H, Lee S and Myung H (2022) PaGO-LOAM: Robust ground-optimized LiDAR odometry. In: International Conference on Ubiquitous Robots. pp. 1–7.
- Shan T and Englot B (2018) LeGO-LOAM: Lightweight and ground-optimized LiDAR odometry and mapping on variable terrain. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 4758–4765.
- Shan T, Englot B, Meyers D, Wang W, Ratti C and Rus D (2020) LIO-SAM: Tightly-coupled LiDAR inertial odometry via smoothing and mapping. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 5135–5142.
- Shi C, Chen X, Huang K, Xiao J, Lu H and Stachniss C (2021a) Keypoint matching for point cloud registration using multiplex dynamic graph attention networks. *IEEE Robotics* and Automation Letters 6(4): 8221–8228.
- Shi J, Yang H and Carlone L (2021b) ROBIN: a graphtheoretic approach to reject outliers in robust estimation using invariants. In: *IEEE International Conference on Robotics and Automation.* pp. 13820–13827.
- Song S, Lim H, Lee AJ and Myung H (2022) DynaVINS: A visualinertial SLAM for dynamic environments. *IEEE Robotics and Automation Letters* 7(4): 11523–11530.
- Sun L (2021a) IRON: Invariant-based highly robust point cloud registration. *arXiv preprint arXiv:2103.04357*.
- Sun L (2021b) RANSIC: Fast and highly robust estimation for rotation search and point cloud registration using invariant compatibility. arXiv preprint arXiv:2104.09133.
- Thrun S and Montemerlo M (2006) The graph SLAM algorithm with applications to large-scale mapping of urban structures. *The International Journal of Robotics Research* 25(5-6): 403– 429.
- Tzoumas V, Antonante P and Carlone L (2019) Outlier-robust spatial perception: Hardness, general-purpose algorithms, and guarantees. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 5383–5390.
- Vandenberghe L and Boyd S (1996) Semidefinite programming. SIAM Review 38(1): 49–95.
- Wang Y and Solomon JM (2019) Deep closest point: Learning representations for point cloud registration. In: IEEE/CVF International Conference on Computer Vision. pp. 3523–3532.
- Yang H, Antonante P, Tzoumas V and Carlone L (2020a) Graduated non-convexity for robust spatial perception: From non-minimal solvers to global outlier rejection. *IEEE Robotics and*

Automation Letters 5(2): 1127-1134.

- Yang H, Shi J and Carlone L (2020b) TEASER: Fast and certifiable point cloud registration. *IEEE Transaction on Robotics* 37(2): 314–333.
- Yang J, Li H, Campbell D and Jia Y (2015) Go-ICP: A globally optimal solution to 3D ICP point-set registration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 38(11): 2241–2254.
- Yang J, Xiao Y, Cao Z and Yang W (2019) Ranking 3D feature correspondences via consistency voting. *Pattern Recognition Letters* 117: 1–8.
- Yew ZJ and Lee GH (2018) 3DFeatNet: Weakly supervised local 3D features for point cloud registration. In: *European Conference on Computer Vision*. pp. 607–623.
- Yin P, Yuan S, Cao H, Ji X, Zhang S and Xie L (2023) Segregator: Global point cloud registration with semantic and geometric cues. arXiv preprint arXiv:2301.07425.
- Youn W, Lim H, Choi HS, Rhudy MB, Ryu H, Kim S and Myung H (2021) State estimation for HALE UAVs with deep-learningaided virtual AOA/SSA sensors for analytical redundancy. *IEEE Robotics and Automation Letters* 6(3): 5276–5283.
- Yue R, Xu H, Wu J, Sun R and Yuan C (2019) Data registration with ground points for roadside LiDAR sensors. *Remote Sensing* 11(11): 1354–1169.
- Zeng A, Song S, Nießner M, Fisher M, Xiao J and Funkhouser T (2017) 3DMatch: Learning local geometric descriptors from RGB-D reconstructions. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1802–1811.
- Zermas D, Izzat I and Papanikolopoulos N (2017) Fast segmentation of 3D point clouds: A paradigm on LiDAR data for autonomous vehicle applications. In: *IEEE International Conference on Robotics and Automation*. pp. 5067–5073.
- Zhang J and Singh S (2014) LOAM: LiDAR odometry and mapping in real-time. In: *Robotics: Science and Systems*, volume 2.
- Zhang L, Helmberger M, Fu LFT, Wisth D, Camurri M, Scaramuzza D and Fallon M (2022) Hilti-Oxford Dataset: A millimetre-accurate benchmark for simultaneous localization and mapping. *arXiv preprint arXiv:2208.09825*.
- Zhang Z and Scaramuzza D (2018) A tutorial on quantitative trajectory evaluation for visual(-inertial) odometry. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- Zhou QY, Park J and Koltun V (2016) Fast global registration. In: *European Conference on Computer Vision*. pp. 766–782.