

G-Loc: Tightly-coupled Graph Localization with Prior Topo-metric Information

Lorenzo Montano-Oliván, Julio A. Placed, Luis Montano, and María T. Lázaro, *Member, IEEE*

Abstract—Localization in already mapped environments is a critical component in many robotics and automotive applications, where previously acquired information can be exploited along with sensor fusion to provide robust and accurate localization estimates. In this work, we offer a new perspective on map-based localization by reusing prior topological and metric information. Thus, we reformulate this long-studied problem to go beyond the mere use of metric maps. Our framework seamlessly integrates LiDAR, inertial and GNSS measurements, and cloud-to-map registrations in a sliding window graph fashion, which allows to accommodate the uncertainty of each observation. The modularity of our framework allows it to work with different sensor configurations (*e.g.*, LiDAR resolutions, GNSS denial) and environmental conditions (*e.g.*, mapless regions, large environments). We have conducted several validation experiments, including the deployment in a real-world automotive application, demonstrating the accuracy, efficiency, and versatility of our system in online localization.

Index Terms—Robot Localization, Graph Optimization, Topo-metric Map.

I. INTRODUCTION

ROBUST and precise (global) localization is essential for autonomous robots and self-driving vehicles. Simultaneous Localization and Mapping (SLAM) solutions [1], [2] have emerged as a response to 3D reconstruction and online localization, growing exponentially in recent years also in the autonomous driving realm [3]. However, there are two situations where these methods suffer from limitations. On the one hand, they assume that the robot is teleoperated, or that the motion actions are known *a priori*. Active SLAM tackles this by integrating online path planning into the problem [4]. On the other hand, it is desirable for many robotic applications to navigate repeatedly in previously visited (and mapped) areas without the need of further updating the built model of the environment. For example, autonomous vehicles often navigate familiar routes, and robots typically operate in known environments when not engaged in exploratory tasks. Using a previously built environment model in such situations can scale down the problem of SLAM to map-based localization. This work specifically focuses on the latter, aiming to leverage existing *prior* information to address localization.

Manuscript received: May 6, 2024; Revised: June 27, 2024; Accepted: August 25, 2024.

This paper was recommended for publication by Editor Sven Behnke upon evaluation of the Associate Editor and Reviewers' comments. This work was partially supported by DGA_FSE T73_23R and PTAS-20211016 DIGIZITY.

Lorenzo Montano-Oliván, Julio A. Placed, and María T. Lázaro are with the Instituto Tecnológico de Aragón (ITA), María de Luna 7, 50018 Zaragoza, Spain {lmontano, jplaced, mtlazaro}@ita.es

Luis Montano is with the Instituto de Investigación en Ingeniería de Aragón (I3A), Universidad de Zaragoza, María de Luna 1, 50018 Zaragoza, Spain montano@unizar.es

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




Fig. 1: Visualization of G-Loc in a sequence of the *Newer College Extension* dataset [5]. The prior topological (orange pose-graph) and metric (colored pointcloud) model is exploited for online robust localization. An *active graph* (blue) is optimized, containing certain vertices from the prior graph and the most recent robot states.

At first glance, without the need of updating an environment model, the use of Global Navigation Satellite System (GNSS) in (semi-)urban areas may seem convenient for localization, although its use alone has long been proven to be neither robust nor accurate enough for the required applications [6]. Combining measurements from multiple sensors and finding correspondences with a reference map model is key to robust, reliable and accurate localization in automotive applications. Most of the existing work tackles this task by using filters for measurement integration, and by exploiting dense high-resolution point cloud maps previously generated by SLAM algorithms [7]–[10]. However, modern SLAM systems, based on probability and graph theories [1], contain much more information beyond the resulting reconstruction: a rich underlying graph that captures the localization estimates and accommodates the observation uncertainties. This valuable topological information has been overlooked in the literature.

In this work, we present G-Loc (see Figure 1), a novel tightly-coupled graph localization approach that exploits prior topo-metric knowledge of the environment. We propose to go beyond purely geometric maps and use all the information provided by modern SLAM systems. In addition, we efficiently fuse this information with LiDAR (Light Detection And Ranging), IMU (Inertial Measurement Unit), and GNSS observations to create a robust and accurate localization system.

The rest of the paper is organized as follows. Section II discusses related work and the contributions of this paper, Section III formally defines the problem, Section IV details

the components of our system, and Section V contains the experimental results. This letter is concluded in Section VI.

II. RELATED WORK

Levinson *et al.* [6] are among the first to exploit a previously-built map for localization. They correlate LiDAR intensity measurements with the prior map (a grid representation containing a Gaussian distribution of intensity in each cell) using a particle filter. They also incorporate GNSS measurements to limit particle dispersion. However, this and other related work [11], reduce the estimation problem to 3 degrees of freedom, assuming that the rest are estimated, *e.g.*, from an IMU. [12] estimates a global 3D pose instead after incorporating altitude into the distribution of the cells.

A different approach is followed in [13], where the reference map consists of road primitives. An Unscented Kalman Filter (UKF) is used to combine GNSS, odometry and the matching between the prior map and the online primitives extracted from images via convolutional neural networks. The challenge of maintaining and matching a dense point cloud map is also addressed in [14], where such representations are transformed into a sparse multi-layer vector map, and [15], where features are extracted from lateral galleries in tunnel-like environments.

More recently, advances in point cloud registration have brought these methods to the forefront of map-based localization, whether feature-based or not. Yoneda and Mita [7] first matched LiDAR measurements to a georeferenced point cloud map using Iterative Closest Point (ICP) [16]. The limitations of this single sensor approach and the sensitivity of ICP to the initial guess are addressed in subsequent works, all of them being filter-based [8]–[10], [17], [18]. Nagy and Benedek [9] fuse point cloud and semantic matching with GNSS measurements to localize a robot in a high resolution 3D semantic cloud. In [19], instead of matching the online LiDAR data to the reference point cloud map, the latter is rendered in 2D; this allows to match monocular images. In [20] LiDAR Odometry (LO) and IMU data are combined in a UKF to improve the matching of the current scan with a portion of the prior point cloud map; this portion is given by cropping the map around the GNSS pose estimation. Autoware [8], a widely used open-source library constantly under development, features global localization in a prior map (either in the form of a point cloud or a high-definition vector map). This approach relies on aligning the input clouds with the reference map using the Normal Distribution Transform (NDT) [21] technique. The optimized transformation is integrated with IMU and GNSS measurements within an extended Kalman filter module. Xia *et al.* [22] combine GNSS, IMU and NDT-based map matching on a Kalman filter. However, they use a light map instead of selecting regions of interest from a dense map. On the downside, this requires a special offline process. *PoseMap* [23] incorporates the trajectory into the prior model. Each robot pose is assigned a submap of surfels and the online cloud is aligned only with the k -closest neighbors, reducing the computational cost and allowing to use the full LiDAR range.

In contrast to the filter-based techniques mentioned above, a handful of methods have recently started to use a graph

formulation, borrowing from the advances in graph SLAM and due to its improved robustness and accuracy [24]. Dubé *et al.* [25] estimate the robot’s pose in a prior global map by extracting segments from a 3D point cloud and matching neighbors in the feature space. The resulting transformations are later fed into a pose-graph SLAM framework along with pre-integrated IMU measurements. [26] also employs a pose-graph formulation for localization, combining wheel odometry, GNSS, and semantic matching with the reference map. In parallel to our work, Koide *et al.* [27] recently proposed a map-based localization approach that tightly couples range and IMU data within a sliding window factor graph optimization.

The localization method proposed in this letter falls into the graph-based category. G-Loc is a modular and optimized framework that adapts to a wide variety of sensor and environment configurations, capable of handling large-scale environments, GNSS denial, mapless regions, etc. Unlike previous work [25]–[27] we go beyond geometric representations and fully exploit the knowledge gathered during a previous mapping task, somewhat following [23] but leveraging the pose-graph topology instead of just the trajectory poses. More specifically, our main contributions are as follows:

- A novel approach to map-based localization that exploits the topology of the prior model both for an efficient search of regions of interest, and for accommodating their uncertainty in a graph optimization process.
- A complete localization system that combines precise LiDAR-inertial odometry, GNSS and map matching constraints. It can accurately locate a robot under various sensor configurations and environmental challenges.
- Thorough experiments demonstrating the accuracy and versatility of our system compared to state-of-the-art localization and SLAM methods, as well as real world deployment in an autonomous driving scenario.

III. PROBLEM FORMULATION

Map-based localization aims to determine the pose of a robot or sensor within an existing map of the environment. This can intuitively be seen as a coordinate transformation process, which consists of establishing correspondences between the map information and the current perception [28].

First of all, let us denote the map and robot base frames by m and r , respectively. To simplify the notation, all the sensors will be assumed to be in r . In its most basic form, map-based 3D localization can be formulated as the problem of estimating the robot poses that best match the sensor observations with each other and with the map, *i.e.*, finding the *maximum a posteriori* estimates

$$\mathcal{T}^* \triangleq \underset{\mathcal{T}}{\operatorname{argmax}} p(\mathcal{T} | \mathcal{M}, \mathcal{Z}), \quad (1)$$

where $p(\cdot)$ denotes probability, \mathcal{M} is the prior map, \mathcal{Z} is the set of sensor measurements, and $\mathcal{T} \triangleq \{ {}^m\mathbf{T}_{r_i} \in SE(3) | i \in \mathcal{K}_t \}$ is the set of relative transformations between the map and the latest $N_{\mathcal{K}}+1$ robot frames. To maintain tractability, the set of frame indices $\mathcal{K}_t \triangleq \{ \max(0, t-N_{\mathcal{K}}), \dots, t \}$ has been limited in a sliding window fashion.

On the one hand, we assume that \mathcal{M} has been previously built by a graph-SLAM algorithm, and thus contains metric and topological information (in the form of point clouds and a pose-graph, respectively). Furthermore, the metric part is associated with the nodes of the graph in such a way that each vertex is assigned a local map (or *submap*, \mathcal{S}') of its surroundings (see upper part of Figure 4). Then:

$$\mathcal{M} \triangleq \{\mathcal{G}, \mathcal{S}_\ell \mid \ell \in \mathcal{R}\}, \quad (2a)$$

$$\mathcal{G} \triangleq \{\mathcal{V}, \mathcal{E}\}, \quad (2b)$$

where $\mathcal{R} \triangleq \{0, \dots, \dim(\mathcal{V})\}$ is the reference frame index set, and $(\mathcal{V}, \mathcal{E})$ are the sets of vertices and edges in the pose-graph. The ℓ -th submap will be formed by concatenating the point clouds between the ℓ and $\ell - N_s$ LiDAR frames, as described in [29], or just a region around the reference vertex. The latter can be easily obtained by splitting the resulting map of a typical SLAM process [30] into grids. Notice that the submaps of nearby vertices might overlap and that elements of the pose-graph will live in $SE(3)$. Additionally, the graph may contain GNSS information, making also the map georeferenced and aligned to the East, North, Up (ENU) global coordinate system. On the other hand, the set of measurements collected up to time t consists of LiDAR point clouds, IMU observations between consecutive frames, and GNSS data: $\mathcal{Z}_t \triangleq \{\mathcal{Z}_i^{\text{LiDAR}}, \mathcal{Z}_{i,j}^{\text{IMU}}, \mathcal{Z}_i^{\text{GNSS}}\}_{i,j \in \mathcal{K}_t}$.

Assuming (i) that the robot poses are normally distributed (i.e., $\mathbf{T} \triangleq \bar{\mathbf{T}} \exp(\delta) \in SE(3)$, with $\bar{\mathbf{T}}$ a large mean transformation and $\delta \in \mathbb{R}^6$ a random vector normally distributed around zero) and pairwise independent, (ii) the measurements are conditionally independent w.r.t. the poses, and (iii) LiDAR and IMU observations are combined to provide odometry estimates, (1) can be rewritten as follows:

$$\mathcal{T}^* = \underset{\mathcal{T}}{\operatorname{argmin}} F(\mathcal{T}) \quad (3)$$

$$\text{s.t. } F(\mathcal{T}) = \sum_{i,j \in \mathcal{K}_t} \|e_{i,j}^{\text{LIO}}\|^2 + \sum_{i \in \mathcal{K}_t} \|e_i^{\mathcal{M}}\|^2 + \sum_{i \in \mathcal{K}_t} \|e_i^{\text{GNSS}}\|^2,$$

where $\|e\|^2 \triangleq (e^T \Sigma^{-1} e) \in \mathbb{R}$ is the quadratic error for a measurement with covariance Σ . $e_{i,j}^{\text{LIO}}$ are the residual errors associated with the LiDAR-inertial odometry (LIO) estimates, $e_i^{\mathcal{M}}$ are those associated with the correspondences between LiDAR observations and the prior map, and e_i^{GNSS} are the translation errors associated with the GNSS measurements.

IV. METHODS

G-Loc is divided into three main threads: LIO, Localization and Dynamic Loading, to which the following subsections are dedicated. Figure 2 provides an overview of the system, showing the different modules and their connections. Thanks to its parallelized architecture, the entire system runs in real time within a ROS2 Humble [31] framework. It is able to use GPU-accelerated cloud registration and is robust to a wide range of sensor configurations with minimal parameter tuning.

LIO combines range and inertial measurements within a graph optimization framework. Online LiDAR data is aligned with most recent point clouds (which are contained in a submap), and high-frequency IMU data is pre-integrated for an efficient optimization. The output of LIO is then fed into

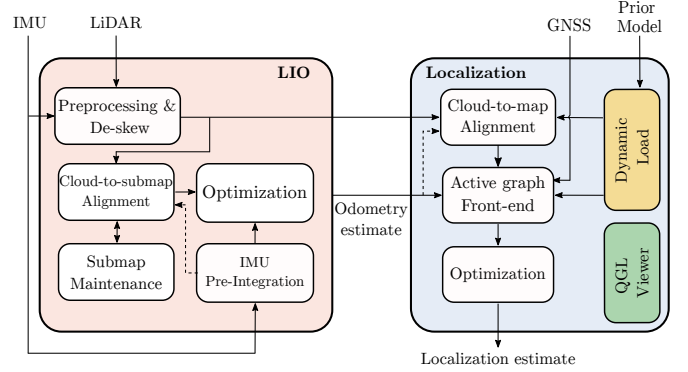


Fig. 2: Block diagram of G-Loc. The dashed lines represent the initial guess fed to the registration methods.

the localization module, along with GNSS measurements and the results from cloud-to-map registration. A second pose-graph is optimized taking into account the previous constraints. Splitting the optimization problem into two different graphs improves the stability and convergence of the system and facilitates the noise configuration [30]. The dynamic loader allows to work with large prior maps by limiting the number of point clouds loaded in memory. Finally, visualization is based on OpenGL and handled in a separate thread.

A. LiDAR-inertial Odometry (LIO)

Relative measurements between consecutive robot poses are computed in a graph-based LIO framework. The incremental inertial state can be defined as

$$\mathcal{X} \triangleq \{\mathbf{x}_i \mid i \in \mathcal{I}_t\}, \quad (4a)$$

$$\mathbf{x}_i \triangleq \{\mathbf{T}_i, \mathbf{v}_i, \mathbf{b}_i\}, \quad (4b)$$

where $\mathbf{v} \in \mathbb{R}^3$ is the linear velocity and $\mathbf{b} = (\mathbf{b}^g, \mathbf{b}^a) \in \mathbb{R}^6$, with $\mathbf{b}^g, \mathbf{b}^a$ in \mathbb{R}^3 the gyroscope and accelerometer bias, respectively. The dimension of \mathcal{X} is constrained by $\mathcal{I}_t \triangleq \{\max(0, t - N_{\mathcal{I}}), \dots, t\}$, similarly to \mathcal{K}_t . The map frame superindex notation has been dropped for readability.

First, a preprocessing step is performed. The input cloud is downsampled to a desired voxel resolution, and points too close/distant are removed. Then, a *de-skew* operation is performed to mitigate the noise in the input scan \mathcal{L}_t . In this way, the angular distortion caused by the robot's ego-motion while capturing the data is corrected. Since each point has an associated acquisition timestamp, angular velocity measurements from the IMU can be integrated and then interpolated to compute the relative orientation w.r.t. the start of the scan.

The de-skewed cloud is then aligned with the previous submap \mathcal{S}_{t-1} created by concatenating the N_s previous point clouds (see Figure 3). *Scan-to-submap alignment* relies either on the efficient and GPU-accelerated implementation of voxelized generalized ICP (FastGICP) [32], or the parallelized version of NDT (NDTOmp) [33]. This process outputs LO estimates between frames (i.e., the transformations ${}^{i-1}\tilde{\mathbf{T}}_i^{\text{LIO}}$), which can be introduced into the graph as prior edges to constrain the pose of the inertial state after computing

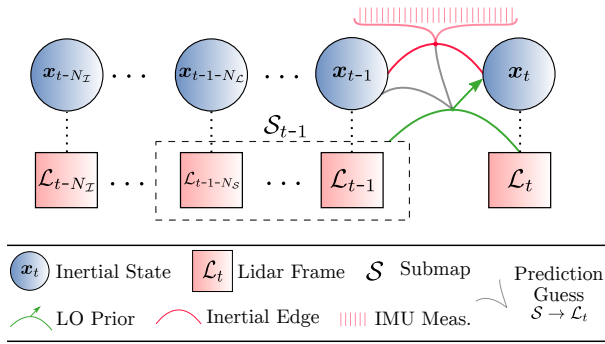


Fig. 3: Scheme of how the different constraints are integrated into the graph-based LIO optimization framework.

${}^i\tilde{\mathbf{T}}^{\text{LO}} = \left(\mathbf{T}_{i-1} {}^{i-1}\tilde{\mathbf{T}}_i^{\text{LO}} \right)^{-1}$. The residuals for the i -th frame will be given by:

$$\mathbf{e}_i^{\text{LO}} \triangleq \log \left(\mathbf{T}_i {}^i\tilde{\mathbf{T}}^{\text{LO}} \right)^\vee \in \mathbb{R}^6, \quad (5)$$

where the logarithm and vee operators are, respectively, $\log(\cdot) : SE(3) \mapsto \mathfrak{se}(3)$ and $(\cdot)^\vee : \mathfrak{se}(3) \mapsto \mathbb{R}^6$, being $\mathfrak{se}(3)$ the Lie algebra of $SE(3)$. The covariance of this measurement is inherited from the registration method [32]. Note that only unaccented variables will be optimization variables.

Simultaneously with the above, IMU measurements between LiDAR frames are pre-integrated following [34]. *Pre-integration* provides an estimate of the robot's motion between successive frames, also allowing the scan-to-submap matching to be fed with a decent initial guess (see Figure 3). High-frequency odometry (at IMU rate) can be obtained by composing the LO with the relative motion predictions from the preintegrated measurement model. IMU data is efficiently incorporated into the graph optimization framework by creating an inertial edge [2], [30] between frames; this edge is actually composed of several different edges that correlate the poses and velocities between frames, and also model the bias variation over time (assumed to follow a Brownian motion pattern). We use the method proposed in [34] to define the relative increments between frames and the preintegrated measurement model, *i.e.*, $\Delta\tilde{\mathbf{R}}_{i,j}$, $\Delta\tilde{\mathbf{v}}_{i,j}$ and $\Delta\tilde{\mathbf{p}}_{i,j}$ (see [34] for details). This allows us to define their residual error as $\mathbf{e}^{\text{IMU}} \triangleq (e_{\Delta R}, e_{\Delta v}, e_{\Delta p}, e_{\Delta b})^T \in \mathbb{R}^{15}$, where:

$$\mathbf{e}_{\Delta R_{i,j}} \triangleq \log(\Delta\tilde{\mathbf{R}}_{i,j}^T \mathbf{R}_i^T \mathbf{R}_j)^{\vee}, \quad (6a)$$

$$\mathbf{e}_{\Delta v_{i,j}} \triangleq \mathbf{R}_i^T (\mathbf{v}_j - \mathbf{v}_i - \mathbf{g} \Delta t_{i,j}) - \Delta\tilde{\mathbf{v}}_{i,j}, \quad (6b)$$

$$\mathbf{e}_{\Delta p_{i,j}} \triangleq \mathbf{R}_i^T (\mathbf{p}_j - \mathbf{p}_i - \mathbf{v}_i \Delta t_{i,j} - \frac{1}{2} \mathbf{g} \Delta t_{i,j}^2) - \Delta\tilde{\mathbf{p}}_{i,j}, \quad (6c)$$

all of which live in \mathbb{R}^3 , and $e_{\Delta b_{i,j}} \triangleq \mathbf{b}_j - \mathbf{b}_i \in \mathbb{R}^6$. Note that the vee and logarithmic operators are now over $SO(3)$.

Finally, the graph containing the above constraints (see Figure 3) is optimized using g2o [35], solving the optimization:

$$\mathcal{X}^* = \underset{\mathcal{X}}{\operatorname{argmin}} \left(\sum_{i \in \mathcal{I}_t} \|\mathbf{e}_i^{\text{LO}}\|^2 + \sum_{i,j \in \mathcal{I}_t} \|\mathbf{e}_{i,j}^{\text{IMU}}\|^2 \right) \quad (7)$$

The success of registration algorithms depend on a good initial guess and sufficient overlap between the input clouds. The first issue has already been addressed, and the second

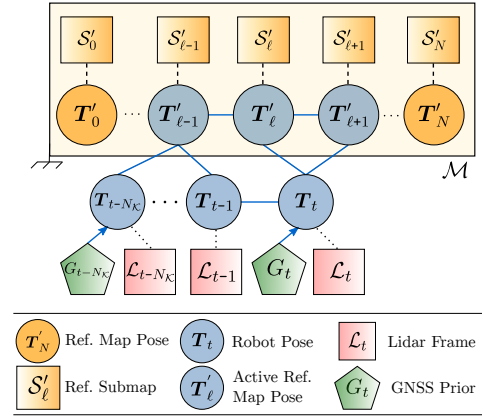


Fig. 4: Overview of the proposed graph-based localization method using prior topological and metric information. Orange elements form the *reference graph*, and blue elements form the *active graph*.

is mitigated by the submaps, which in turn require some attention. We use the scan-to-submap alignment to add new point clouds to the submap. Then, the latter is downsampled (*e.g.*, via uniform sampling) to prune redundant points. To restrain complexity, the submap size is limited to include the last N_S point clouds. A side effect of this approach is high adaptability to different range sensors: we can handle densities from 16 to 128 beams and different patterns by changing N_S . By effectively configuring the submaps (*e.g.*, $N_S = 1$ for 128-plane LiDARs), we keep the computational cost manageable while achieving the precision required for each application.

B. 3D Graph-based Localization

The global localization process is based on optimizing the *active graph*, which composed of a *sliding graph* with the latest robot states (\mathcal{T}), and regions of interest of the *reference graph* (*i.e.*, certain reference vertices, \mathcal{T}'). See Figure 4.

The localization process starts with estimating the robot's initial pose within the prior map using GNSS measurements¹. It then queries the prior topo-metric model to find the k -closest reference vertices (for the exemplary graph shown in Figure 4, at time t , they will be \mathbf{T}'_ℓ and $\mathbf{T}'_{\ell+1}$). Next, the current scan is aligned with the submaps associated with these vertices ($\mathcal{S}'_\ell, \mathcal{S}'_{\ell+1}$). The *active graph* combines the constraints enforced by the results of these alignments with the relative motion data provided by the LIO and GNSS measurements (if available). Therefore, this new graph will contain vertices from both the prior reference graph and the online sliding graph. The latter encodes current and past robot states in a sliding window approach (*i.e.*, \mathcal{T}), being older nodes marginalized as they leave the window. Figure 1 displays the active (blue) and reference (orange) graphs on top of the prior point cloud map for a localization task. The online cloud is shown in red. Note how multiple reference vertices can be associated with a single vertex in the sliding graph (as is also shown in Figure 4).

The optimization of the active graph is governed by (3), *s.t.*:

$$\mathbf{e}_{i,j}^{\text{LIO}} \triangleq \log \left(\mathbf{T}_i {}^i\tilde{\mathbf{T}}_j^{\text{LIO}} (\mathbf{T}_j)^{-1} \right)^\vee \in \mathbb{R}^6, \quad (8a)$$

¹In the absence of GNSS, this initial estimate can be given manually by the user, *e.g.*, using a ROS2 service.




Fig. 5: Visualization of dynamic loading. The prior graph is shown in orange, while the active graph is depicted in blue. For efficiency, only a subset of the geometric model is loaded, specifically the submaps corresponding to vertices within the robot’s neighborhood.

$$\mathbf{e}_i^{\mathcal{M}} \triangleq \sum_{k \in \mathcal{R}_i} \log \left(\mathbf{T}_i {}^i \tilde{\mathbf{T}}_k^{\mathcal{M}} (\mathbf{T}'_k)^{-1} \right)^{\vee} \in \mathbb{R}^6, \quad (8b)$$

$$\mathbf{e}_i^{\text{GNSS}} \triangleq \mathbf{p}_i - \tilde{\mathbf{p}}_i \in \mathbb{R}^3, \quad (8c)$$

where ${}^i \tilde{\mathbf{T}}_j^{\text{LIO}}$ is the output from the previous subsection, ${}^i \tilde{\mathbf{T}}_k^{\mathcal{M}}$ are the registration results between the current scan and the prior map, and $\mathbf{T}'_k \in \mathcal{G}$. Note that correspondences with the prior map are only evaluated for a subset of the reference vertices, *i.e.*, $k \in \mathcal{R}_i \subset \mathcal{R}$. Finally, $\tilde{\mathbf{p}}_i = {}^{\text{UTM}} \tilde{\mathbf{p}}_i - {}^{\text{UTM}} \tilde{\mathbf{p}}_m$, with ${}^{\text{UTM}} \tilde{\mathbf{p}}_i$ the actual GNSS measurement and ${}^{\text{UTM}} \tilde{\mathbf{p}}_m$ the origin of the prior model (this origin will be associated with a Universal Transverse Mercator coordinate). Relativizing GNSS allows to merge it with cloud-to-map matching and LIO constraints.

On the one hand, LIO is used to connect sequential vertices of the sliding graph with the covariance depending on the alignment result (defined as the mean error between all matched points, weighted by the match percentage). In addition, it provides an initial guess for the cloud-to-map registration. On the other hand, registrations between the current scan and the reference submaps are introduced as binary edges between vertices in the reference and sliding graphs, with information also depending on the alignment result. Also, the transformation with the smallest covariance is used to compute the initial estimate for the optimization process (handled by g2o). The previous constraints connect the two graphs and form the active graph (blue elements in Figure 4). Should the sliding and reference graphs become disconnected, either due to alignment failure or the lack of nearby reference vertices, the algorithm will rely on LIO and GNSS until re-localization occurs. Finally, GNSS measurements are used as prior unary edges when they are synchronized with the online cloud and if their uncertainty is below a certain threshold.

C. Dynamic Loading

This module allows loading only nearby parts of the reference model, \mathcal{M} , as the robot moves; this is especially useful for large maps where maintaining (and visualizing) all submaps at high resolution is intractable. Since both the reference and sliding graphs are georeferenced, it is possible to load the submaps near the current localization estimate in a sliding window fashion. To do this, as a new vertex is

added to the sliding graph, we query the closest vertex in the reference graph and load its associated submap and those of its k -closest neighbors. While the number of reference submaps in memory is below a user-specified limit, the process is repeated after blacklisting the loaded submaps to avoid duplicates. In addition, the submaps that leave this window are removed from memory. This method allows us to control the total load by specifying the maximum number of submaps (in practice, we load 15 to 30 submaps). Figure 5 illustrates the dynamic load of a portion of the environment. Thus, we maintain a constant RAM memory usage during operation (*e.g.*, ~ 50 MB for 30 submaps). In contrast, for some of the experiments in large environments that we present in the next section, loading the entire dense cloud would require up to 2 GB of memory.

V. EXPERIMENTAL RESULTS

In this section, we present a thorough evaluation of G-Loc in different applications and under different sensor configurations, demonstrating the versatility of our approach. We use three real urban datasets: *European Union Long-Term (EULT)* [36], *Newer College Extension (NCE)* [5], and our own dataset collected on the *Campus of the University of Zaragoza (CUZ)*. In addition, we present the results of using G-Loc in a real-world automotive application, in the context of an autonomous driving bus. The dataset experiments were performed on a laptop equipped with an Intel® Core™ i7-1165G7 CPU and an Nvidia GeForce® RTX® 3080 GPU, while the bus was equipped with an Intel® Core™ i9-12900TE and an Nvidia RTX® A2000 (note that many other autonomous driving processes were running concurrently on this hardware). Unless otherwise noted, [32] was used for registration. Table I contains a short summary of the experiments and the sensors installed in the mobile platforms. The error metrics presented in this section were obtained using `evo` [37].

A. Dataset Evaluation

The *EULT* dataset contains several sequences recorded with a car in a downtown. They all follow a similar trajectory under different conditions (*e.g.*, traffic, weather); this is particularly interesting for testing map-based localization, since we need different but overlapping reference and evaluation trajectories. Sequence and sensor details are shown in Table I.

To evaluate the performance of G-Loc, we first generated a georeferenced model of the environment using one of the available sequences, which we will refer to as the *reference sequence*². To obtain the required prior model with topological (a pose-graph) and metric (point cloud submaps) information, we used [29], a SLAM system with a similar architecture to G-Loc but replacing the localization module with a mapping module for robust loop closing and global optimization. To evaluate the performance in mapless regions, we removed an entire road from the prior model (~ 0.4 km); the SLAM-estimated trajectory is shown in orange in Figure 6a. The built model was fed into G-Loc while running two different *evaluation sequences*². Figure 6a also depicts the raw GNSS data (green)

²Ref. seq. ID: 2018-07-19. Eval. seq. ID: 2018-07-16 and 2018-07-17.

Dataset	Length (per sequence)	Δ Elevation	LiDAR	IMU	GNSS
EULT (ref. & eval.)	~ 5 km	~ 50 m	2x Velodyne HDL-32E	XSens MTi-28A53G25	Magellan ProFlex 500
NCE (ref. & eval.)	~ 0.3 km	~ 0 m	Ouster OS0-128		—
CUZ (ref.)	~ 2.7 km	~ 1 m	Ouster OS1-128		Ardusimple RTK2B
CUZ (eval.)	~ 1.5 km	~ 1 m	Livox HAP		
Bus	~ 8 km	~ 70 m	2x RS Helios 32-5515	XSens MTi-680-G	

TABLE I: Summary of the sensor configurations and environmental characteristics in the experiments performed.




Fig. 6: Evaluation results in EULT (a)-(b) and CUZ (c)-(d) datasets. (a), (c): georeferenced SLAM trajectories in the reference sequences (orange) and raw GNSS data in the evaluation sequences (green). (b), (d): Trajectories estimated by G-Loc in the evaluation sequences; the color map indicates 2D ATE (in meters, darker is smaller).

for one of them to facilitate comparison with the reference trajectory. The middle right section contains the cropped region, where there is GNSS signal but no prior information. When the robot traverses a mapped region, the prior model is exploited to improve localization and correct for potential drift in LIO. Conversely, in unmapped regions, the sliding graph disconnects from the reference graph, accumulating a larger error that is eventually corrected after re-localization. Global measurements in this large mapless region kept the absolute error bounded, enabling the system to re-localize with the map as it traversed previously mapped areas. Figure 6b contains the 2D Absolute Trajectory Error (ATE) for translation along the trajectory for one of the sequences³, showing that the highest errors occur on the unmapped road and that even under these conditions they are within reasonable limits (see colormap). Table II contains the numerical results (mean and standard deviation) for the cropped sequences, namely Root Mean Squared Error (RMSE) of 3D ATE for translation and rotation, and lateral and longitudinal error. Furthermore, we report the results of using a complete prior map (*i.e.*, uncropped); in this case, the errors were similar but exhibited a smaller deviation. We also evaluated the use of NDTomp, which is of particular interest for resource-constrained platforms. In this case, the overhead slightly increased the localization errors.

Finally, to benchmark G-Loc, we conducted the same experiments using Autoware [8], a popular localization system for autonomous vehicles. For a fair comparison, we fed velocity measurements from our LIO module and used NDTomp. Autoware failed to localize in the mapless region, leading to an unrecoverable state. Using the complete prior model, the translation errors were comparable to G-Loc but the rotation error was nearly doubled (see Table II). In addition, Autoware

required the use of an Intel® Core™ i9 CPU to work properly.

The CUZ dataset was captured using a sensor-equipped car. We recorded two sequences, which partially overlap, but with two different LiDAR sensors. See Table I. To evaluate this dataset, we followed a similar procedure as in EULT. First, we built the prior model with a reference sequence using a high-performance sensor with 360° Field of View (FoV) and 128 planes. Then, we evaluated G-Loc with an evaluation sequence, exploiting the prior topo-metric knowledge and using a LiDAR with 120° FoV and a completely different sensing pattern. Figure 6c displays the trajectory of the sequence (orange) and the raw GNSS data for the evaluation one (green). Both trajectories overlap for most of the route, and when they do not, they are close enough so that the prior map can be used (unlike in EULT). ATE error along the trajectory is shown in Figure 6d and numerical results are reported in the bottom of Table II. As in EULT, the use of FastGICP gave the best results, except for the lateral error (however, these lower mean values are accompanied by higher variance). This experiment demonstrates the adaptability of our method to the use of different sensors for building the prior map and localization. The submaps played a fundamental role here, mitigating the effects of the different sensing patterns and feeding the registration algorithms with sufficient overlapping regions. In contrast, Autoware accumulated much higher errors and had difficulty finding correspondences in many regions (which explains the large deviations), eventually leading to an unrecoverable state. The results reported in Table II for this experiment (marked with †) are only before this divergence.

The NCE dataset contains several sequences recorded using a handheld device in Oxford. This dataset allows us to test G-Loc under different conditions than the previous ones. It includes aggressive motion, where IMU integration is crucial for undistorting the point cloud and predicting the motion between successive LiDAR frames. We have used *Quad*, *Math*, and *Underground* collections of NCE, each of which contains three sequences of increasing difficulty and slightly different

³Due to the lack of ground-truth in this dataset and the noise levels of GNSS in certain regions, errors were calculated w.r.t. SLAM. This algorithm fuses GNSS with LIO to mitigate the effects of noisy measurements and GNSS-denied areas (see zoomed region of Figure 6a). The error between GNSS and SLAM is negligible when the former is reliable.

Dataset	Method	ATE (trans., m)	Long. Error (m)	Lat. Error (m)	ATE (rot., deg)	Time per frame (ms)
EULT (cropped)	G-Loc (FastGICP)	0.209 (0.395)	0.087 (0.280)	0.048 (0.228)	0.423 (0.445)	25.13 (7.82)
	G-Loc (NDTomp)	0.240 (0.300)	0.086 (0.304)	0.021 (0.208)	1.008 (1.160)	38.01 (14.86)
	Autoware	×	×	×	×	×
EULT (complete)	G-Loc (FastGICP)	0.194 (0.201)	0.101 (0.240)	0.014 (0.197)	0.440 (0.576)	26.47 (4.31)
	G-Loc (NDTomp)	0.218 (0.196)	0.119 (0.245)	0.009 (0.203)	0.953 (0.836)	40.13 (12.99)
	Autoware	0.212 (0.194)	0.123 (0.232)	0.006 (0.203)	0.784 (1.380)	[‡] 37.31 (14.52)
CUZ	G-Loc (FastGICP)	0.154 (0.102)	0.071 (0.132)	0.019 (0.126)	0.473 (0.537)	34.62 (9.14)
	G-Loc (NDTomp)	0.214 (0.151)	0.131 (0.179)	0.014 (0.158)	0.650 (0.878)	72.15 (31.54)
	Autoware	[†] 0.308 (1.251)	[†] 0.400 (0.655)	[†] 0.086 (1.353)	[†] 2.146 (0.850)	[‡] 59.74 (71.47)

TABLE II: Localization errors (mean and std. dev.) and time consumed per frame for G-Loc (FastGICP and NDTomp) and Autoware in EULT and CUZ datasets. Best results are bold, and × indicates an unrecoverable failure in localization.

Method / Seq.	Quad-m	Quad-h	Math-m	Math-h	Und-m	Und-h	Avg.
LIO-SAM [30]	0.067	0.13	0.13	0.085	0.065	0.42	0.15
NV-LIO [38]	<u>0.076</u>	0.18	<u>0.095</u>	0.094	<u>0.072</u>	<u>0.078</u>	<u>0.10</u>
LiDAR-Loc [39]	0.11	×	0.10	×	0.11	×	0.11
HDL-Loc [40]	0.10	×	0.11	×	0.11	0.12	0.11
G-Loc (ours)	0.067	0.087	0.090	<u>0.093</u>	0.074	0.075	0.081

TABLE III: ATE RMSE (m) in NCE dataset. Best results are bold and second best are underlined. × indicates an unrecoverable failure.

trajectories. First, we used [29] to build a prior model with the *easy* sequences and then ran our system in the *medium* and *hard* ones. Figure 1 shows the resulting prior model for the *Math* collection. We compared the performance of G-Loc with two SLAM (LIO-SAM [30], NV-LIO [38]) and two map-based localization methods (LiDAR-Loc [39], HDL-Loc [40]). For a fair comparison, all localization methods used the same prior map. Table III contains ATE RMSE, with the values for the SLAM methods extracted from [38]. The results showcase the robustness and accuracy of our system, which outperformed the localization methods and succeeded in every sequence. The failures of [39], [40] were caused by the lack of prior maps in certain regions (*Quad-h*) and aggressive motion (*Math-h*). G-Loc maintained accurate localization in these regions thanks to LIO, allowing for re-localization. G-Loc also outperformed the SLAM methods in most sequences, achieving an average error of 8 cm (note that the localization errors ultimately depend on the error of the prior model, see [29] for detailed values).

B. On the Time Consumption

While the above section aimed to demonstrate the robustness and precision of our approach, we now seek to demonstrate its real-time performance. The processing times per frame consumed by each algorithm in the conducted experiments are listed in Table II. G-Loc required, on average, 25 ms in EULT and 34 ms in CUZ —this is much faster than the typical LiDAR frame rate (10 Hz) and therefore meets real-time requirements. In both cases, LIO accounts for 60% of that time, and the localization module for the rest (visualization and dynamic load are excluded as they operate in the background). G-Loc with NDTomp took considerably longer, but the frame rate was generally not exceeded. We also report the time consumed by Autoware (marked with [‡], as it used a better CPU), which is similar to G-Loc using NDTomp.

Finally, we compared the use of G-Loc to running a full SLAM system to reveal the benefits of exploiting the known prior topo-metric map instead of building a new map. We used the EULT dataset because it is the largest and therefore where a new map would be more costly to generate and maintain.

The time required to process a frame is similar to G-Loc (the LIO update), but the graph optimization and loop closing took 190.4 ± 101.4 ms. These demanding processes slowed down the entire pipeline and caused interruptions in the localization.

C. Real-world Deployment

G-Loc has been deployed on an autonomous bus (see Figure 7a) on an urban bus line in the city of Zaragoza (Aragón, Spain) as part of the pioneering Spanish project on urban driving, DIGIZITY. The sensors installed on the vehicle and the route specifics appear in Table I. Our system requires coordination with other navigation systems to provide a robust and reliable localization, ensuring safe navigation even in areas with poor GNSS coverage. The prior model of the entire route is shown in Figure 7b. In addition, localization starts as soon as the bus is turned on. This means that the localization has to work without a prior model from the departure point until the route is reached (~ 2 km). The system localizes in the unmapped region using LIO and GNSS in a sliding-graph optimization fashion until re-localization with the prior model occurs. Figure 7c depicts the trajectory estimated by G-Loc for a single trip, and Figure 7d shows the improvement of these estimates (blue) over the raw GNSS data (green) in one of the many low coverage areas. During the deployment, the bus operated seamlessly for about 100 hours over a span of 20 days, covering about 800 km. Notably, there were no localization failures and the errors remained within bounds comparable to previous results (*i.e.*, at the cm level), further showcasing the robustness and accuracy of G-Loc.

VI. CONCLUSION

In this work we have presented G-Loc, a localization system that seamlessly integrates GNSS data, LIO and cloud-to-map registration into a graph optimization framework. We have proposed a novel method for map-based localization that leverages the results of a previous graph-SLAM process, utilizing both the geometric map and the underlying pose-graph representation. We have tested G-Loc in several datasets and in a real-world automotive application, demonstrating its accuracy, robustness, and adaptability to different sensor configurations and environmental conditions (*e.g.*, the lack of prior knowledge in some parts of the environment). Long-term autonomy requires the sporadic execution of mapping to update the robot’s knowledge of the environment. This work sets the stage for developing a long-term approach where the prior model is updated and refined with online information.

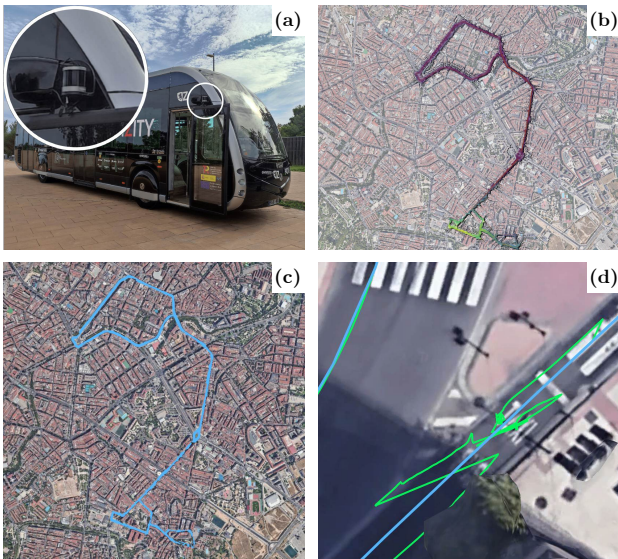


Fig. 7: Results of the real-world experiment. (a) Platform and sensor setup. (b) Prior topo-metric model (colored by height, darker is lower). (c) Complete trajectory estimated by G-Loc. (d) Comparison between the raw GNSS (green) and the localization estimates (blue).

REFERENCES

- [1] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Trans. on Robotics*, vol. 32, no. 6, pp. 1309–1332, 2016.
- [2] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. Montiel, and J. D. Tardós, "ORB-SLAM3: An accurate open-source library for visual, visual-inertial, and multimap SLAM," *IEEE Trans. on Robotics*, vol. 37, no. 6, pp. 1874–1890, 2021.
- [3] G. Bresson, Z. Alsayed, L. Yu, and S. Glaser, "Simultaneous localization and mapping: A survey of current trends in autonomous driving," *IEEE Trans. on Intelligent Vehicles*, vol. 2, no. 3, pp. 194–220, 2017.
- [4] J. A. Placed, J. Strader, H. Carrillo, N. Atanasov, V. Indelman, L. Carlone, and J. A. Castellanos, "A survey on active simultaneous localization and mapping: State of the art and new frontiers," *IEEE Trans. on Robotics*, vol. 39, no. 3, pp. 1686–1705, 2023.
- [5] L. Zhang, M. Camurri, D. Wisth, and M. Fallon, "Multi-camera LiDAR inertial extension to the Newer College dataset," *arXiv preprint arXiv:2112.08854*, 2021.
- [6] J. Levinson, M. Montemerlo, and S. Thrun, "Map-based precision vehicle localization in urban environments," in *Proc. of Robotics: Science and Systems*, vol. 4, 2007, p. 1.
- [7] K. Yoneda, H. Tehrani, T. Ogawa, N. Hukuyama, and S. Mita, "LiDAR scan feature for localization with highly precise 3D map," in *IEEE Intelligent Vehicles Symp.*, 2014, pp. 1345–1350.
- [8] S. Kato, E. Takeuchi, Y. Ishiguro, Y. Ninomiya, K. Takeda, and T. Hamada, "An open approach to autonomous vehicles," *IEEE Micro*, vol. 35, no. 6, pp. 60–68, 2015.
- [9] B. Nagy and C. Benedek, "Real-time point cloud alignment for vehicle localization in a high resolution 3D map," in *Proc. of the European Conf. on Computer Vision Workshops*, 2019, pp. 226–239.
- [10] X. Shan, A. Cabani, and H. Chafouk, "A survey of vehicle localization: Performance analysis and challenges," *IEEE Access*, vol. 11, pp. 107 085–107 107, 2023.
- [11] R. W. Wolcott and R. M. Eustice, "Robust LiDAR localization using multiresolution Gaussian mixture maps for autonomous driving," *Int. Journal of Robotics Research*, vol. 36, no. 3, pp. 292–319, 2017.
- [12] G. Wan, X. Yang, R. Cai, H. Li, Y. Zhou, H. Wang, and S. Song, "Robust and precise vehicle localization based on multi-sensor fusion in diverse city scenes," in *IEEE ICRA*, 2018, pp. 4670–4677.
- [13] F. Poggenhans, N. O. Salscheider, and C. Stiller, "Precise localization in high-definition road maps for urban regions," in *IEEE/RSJ IROS*, 2018, pp. 2167–2174.
- [14] E. Javanmardi, Y. Gu, M. Javanmardi, and S. Kamijo, "Autonomous vehicle self-localization based on abstract map and multi-channel LiDAR in urban area," *IATSS Research*, vol. 43, no. 1, pp. 1–13, 2019.
- [15] T. Seco, M. T. Lázaro, J. Espelósín, L. Montano, and J. L. Villarroel, "Robot localization in tunnels: Combining discrete features in a pose graph framework," *Sensors*, vol. 22, no. 4, 2022.
- [16] P. J. Besl and N. D. McKay, "Method for registration of 3-D shapes," in *Sensor fusion IV: control paradigms and data structures*, vol. 1611, 1992, pp. 586–606.
- [17] W.-C. Ma, I. Tartavull, I. A. Bârsan, S. Wang, M. Bai, G. Mattyus, N. Homayounfar, S. K. Lakshmikanth, A. Pokrovsky, and R. Urtasun, "Exploiting sparse semantic HD maps for self-driving vehicle localization," in *IEEE/RSJ IROS*, 2019, pp. 5304–5311.
- [18] J. Jeong, Y. Cho, and A. Kim, "HDML-LOC: Exploiting high definition map image for precise localization via bitwise particle filter," *IEEE Robotics and Automation L.*, vol. 5, no. 4, pp. 6310–6317, 2020.
- [19] R. W. Wolcott and R. M. Eustice, "Visual localization within LiDAR maps for automated urban driving," in *IEEE/RSJ IROS*, 2014, pp. 176–183.
- [20] Y. Wu, Y. Li, W. Li, H. Li, and R. Lu, "Robust LiDAR-based localization scheme for unmanned ground vehicle via multisensor fusion," *IEEE Trans. on Neural Networks and Learning Systems*, vol. 32, no. 12, pp. 5633–5643, 2021.
- [21] P. Biber and W. Straßer, "The normal distributions transform: A new approach to laser scan matching," in *IEEE/RSJ IROS*, vol. 3, 2003, pp. 2743–2748.
- [22] X. Xia, N. P. Bhatt, A. Khajepour, and E. Hashemi, "Integrated inertial-LiDAR-based map matching localization for varying environments," *IEEE Trans. on Intelligent Vehicles*, 2023.
- [23] P. Egger, P. V. Borges, G. Catt, A. Pfrunder, R. Siegwart, and R. Dubé, "Posemap: Lifelong, multi-environment 3D lidar localization," in *IEEE/RSJ IROS*, 2018, pp. 3430–3437.
- [24] G. Grisetti, R. Kümmerle, C. Stachniss, and W. Burgard, "A tutorial on graph-based SLAM," *IEEE Intelligent Transportation Systems Mag.*, vol. 2, no. 4, pp. 31–43, 2010.
- [25] R. Dube, A. Cramariuc, D. Dugas, H. Sommer, M. Dymczyk, J. Nieto, R. Siegwart, and C. Cadena, "SegMap: Segment-based mapping and localization using data-driven descriptors," *Int. Journal of Robotics Research*, vol. 39, no. 2-3, pp. 339–355, 2020.
- [26] C. Guo, M. Lin, H. Guo, P. Liang, and E. Cheng, "Coarse-to-fine semantic localization with HD map for autonomous driving in structural scenes," in *IEEE/RSJ IROS*, 2021, pp. 1146–1153.
- [27] K. Koide, S. Oishi, M. Yokozuka, and A. Banno, "Tightly coupled range inertial localization on a 3D prior map based on sliding window factor graph optimization," in *IEEE ICRA*, 2024.
- [28] S. Thrun, "Probabilistic robotics," *Communications of the ACM*, vol. 45, no. 3, pp. 52–57, 2002.
- [29] L. Montano-Oliván, J. A. Placed, L. Montano, and M. T. Lázaro, "From underground mines to offices: A versatile and robust framework for range-inertial SLAM," *arXiv preprint arXiv:2407.14797*, 2024.
- [30] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus, "LIO-SAM: Tightly-coupled LiDAR inertial odometry via smoothing and mapping," in *IEEE/RSJ IROS*, 2020, pp. 5135–5142.
- [31] S. Macenski, T. Foote, B. Gerkey, C. Lalancette, and W. Woodall, "Robot Operating System 2: Design, architecture, and uses in the wild," *Science Robotics*, vol. 7, no. 66, 2022.
- [32] K. Koide, M. Yokozuka, S. Oishi, and A. Banno, "Voxelized GICP for fast and accurate 3D point cloud registration," in *IEEE ICRA*, 2021, pp. 11 054–11 059.
- [33] K. Koide, "NDT OMP: Multi-threaded and SSE-friendly NDT algorithm," https://github.com/koide3/ndt_omp, 2017, accessed: 2024.
- [34] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza, "On-manifold preintegration for real-time visual-inertial odometry," *IEEE Trans. on Robotics*, vol. 33, no. 1, pp. 1–21, 2016.
- [35] G. Grisetti, R. Kümmerle, H. Strasdat, and K. Konolige, "g2o: A general framework for (hyper) graph optimization," in *IEEE ICRA*, 2011, pp. 9–13.
- [36] Z. Yan, L. Sun, T. Krajník, and Y. Ruichek, "EU long-term dataset with multiple sensors for autonomous driving," in *IEEE/RSJ IROS*, 2020, pp. 10 697–10 704.
- [37] M. Grupp, "evo: Python package for the evaluation of odometry and SLAM," <https://github.com/MichaelGrupp/evo>, 2017, accessed: 2024.
- [38] D. Chung and J. Kim, "NV-LIO: LiDAR-inertial odometry using normal vectors towards robust SLAM in multifloor environments," *arXiv preprint arXiv:2405.12563*, 2024.
- [39] R. Sasaki, "lidar_localization_ros2," https://github.com/rsasaki0109/lidar_localization_ros2, 2023, accessed: 2024.
- [40] K. Koide, "Real-time 3D localization using a (velodyne) 3D LiDAR," https://github.com/koide3/hdl_localization, 2018, accessed: 2024.