



# LiODOM: Adaptive local mapping for robust LiDAR-only odometry

Emilio Garcia-Fidalgo<sup>\*</sup>, Joan P. Company-Corcoles, Francisco Bonnin-Pascual, Alberto Ortiz

Department of Mathematics and Computer Science, University of the Balearic Islands, Ctra. Valldemossa km 7.5, Palma de Mallorca, 07122, Spain  
 Institut d'Investigació Sanitària Illes Balears, Hospital Universitari Son Espases, Ctra. Valldemossa 79, Palma de Mallorca, 07120, Spain

## ARTICLE INFO

### Article history:

Received 4 February 2022  
 Received in revised form 20 May 2022  
 Accepted 22 July 2022  
 Available online 28 July 2022

### Keywords:

LiDAR odometry  
 Mapping  
 Localization

## ABSTRACT

In the last decades, Light Detection And Ranging (LiDAR) technology has been extensively explored as a robust alternative for self-localization and mapping. These approaches typically state ego-motion estimation as a non-linear optimization problem dependent on the correspondences established between the current point cloud and a map, whatever its scope, local or global. This paper proposes LiODOM, a novel LiDAR-only ODometry and Mapping approach for pose estimation and map-building, based on minimizing a loss function derived from a set of weighted point-to-line correspondences with a local map abstracted from the set of available point clouds. Furthermore, this work places a particular emphasis on map representation given its relevance for quick data association. To efficiently represent the environment, we propose a data structure that combined with a hashing scheme allows for fast access to any section of the map. LiODOM is validated by means of a set of experiments on public datasets, for which it compares favourably against other solutions. Its performance on-board an aerial platform is also reported.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

## 1. Introduction

Self-localization and mapping, either performed simultaneously or in a sequential fashion, are crucial abilities for a mobile robot to be useful in relevant applications, irrespective of whether the robot operates fully autonomously or in a semi-autonomous way. As stated many years ago, odometry estimation is a fundamental piece within this framework. A plethora of sensing devices have been adopted throughout the years, comprising tachometers/wheel encoders, inertial and heading sensors, time of flight sensors, and motion estimation devices, to name but a few. Among all of them, laser scanners and, for a few years now, cameras have turned out to be the sensors of choice. The latter have been extensively used [1,2] due to the rich perception of the surrounding world encoded in images. Vision-based estimation is however sensitive to lighting conditions, have a limited horizontal field of view and require additional calculations to acquire depth and shape perception. In contrast, 3D laser scanners provide a 360-degree overview of the platform surroundings, supply reliable range estimations and, especially motivated by

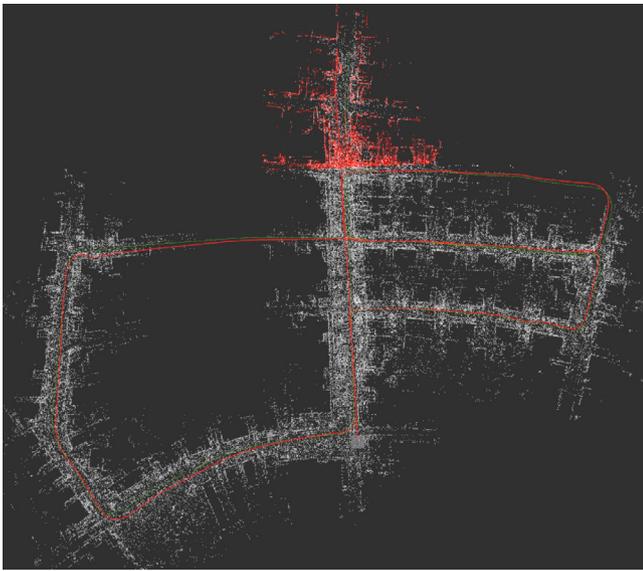
the development of self-driving cars, recently have become an affordable choice for pose estimation and mapping.

LiDAR odometry is typically stated as an optimization problem that is solved using the Iterative Closed Point (ICP) algorithm [3] or any of its variants. For this to happen in a satisfactory, fast and accurate way, a set of reliable correspondences between the current point cloud and a map must be found. A KD-tree is a popular choice to represent the whole map [4], although the resulting performance degrades as the number of points to be managed increases, what makes necessary a filtering step to screen most relevant points. An alternative is to build a local map using a sliding window [5,6], although this might discard useful associations that could be found if the search was performed over a global map.

This paper proposes LiODOM, a novel LiDAR odometry and mapping approach that is able to estimate the pose without additional sensors, e.g. IMU and/or GPS, unlike other recent approaches [5,6]. Our approach is formulated as a non-linear optimization problem based on a set of point-to-line constraints, weighted according to the distance from each point to the sensor centre. Furthermore, we propose an efficient data structure, based on a hashing scheme, to represent the map. As a result, a local map can be retrieved according to the pose of the robot in an effective way, and point cloud correspondences against the local map can be efficiently established. This *adaptive* solution naturally allows us to find correspondences between the current point cloud and revisited places (contrary to just using a sliding

<sup>\*</sup> Corresponding author at: Department of Mathematics and Computer Science, University of the Balearic Islands, Ctra. Valldemossa km 7.5, Palma de Mallorca, 07122, Spain

E-mail addresses: [emilio.garcia@uib.es](mailto:emilio.garcia@uib.es) (E. Garcia-Fidalgo), [joan.pep.company@gmail.com](mailto:joan.pep.company@gmail.com) (J.P. Company-Corcoles), [xisco.bonnin@uib.es](mailto:xisco.bonnin@uib.es) (F. Bonnin-Pascual), [alberto.ortiz@uib.es](mailto:alberto.ortiz@uib.es) (A. Ortiz).



**Fig. 1.** Example of map produced by LiODOM (KITTI 05 sequence), comprising an unoptimized global map generated during navigation (in white) and a local map (in red) that is retrieved according to the position of the vehicle, to be used for next pose estimation. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

window). Fig. 1 illustrates the performance of LiODOM. In brief, the main contributions of this work are:

- A LiDAR-only odometry framework that is based on an optimization problem supported by weighted point-to-line factors computed from the correspondences with a local map.
- A fast and efficient mapping approach based on a hash-based data structure that speeds up searches and permits to gracefully update large-scale maps.
- An extensive evaluation of the proposed approach on public datasets, including a comparison with other state-of-the-art methods.
- An additional evaluation of LiODOM on-board an aerial platform where it provides, among others, speed estimations that are used inside velocity control loops, what shows LiODOM's capability to operate in real time.
- A final contribution is the availability for the community of the source code<sup>1</sup> of our approach.

The rest of the paper is organized as follows: Section 2 overviews most relevant works in the field; the proposed framework is introduced in Sections 3, 4 and 5; Section 6 reports on the results obtained; to finish, Section 7 concludes the paper and suggests future research lines.

## 2. Related work

Laser scanners and vision cameras are the two sensor modalities that have attracted the largest number of researchers in the last years because of their respective success stories and well-known advantages regarding motion estimation – briefly speaking, reliable and scale-aware depth measurements in the first case, and low weight, low power and low cost in the second case. Although LiODOM belongs to the first class, in the following, we not only revise most recent laser-based works close to our solution, but also others relying on imaging sensors in order to gain

perspective on the respective methodologies and achievements, without being intent to be exhaustive. Finally, we also consider map indexing.

### 2.1. LiDAR-based solutions

Most recent approaches carry out LiDAR-based odometry in combination with an IMU for higher accuracy. These solutions are typically regarded as loosely- and tightly-coupled methods [6, 7]. Loosely-coupled methods estimate the state from each sensor separately. Arguably the most well-known method that falls into this category is LiDAR Odometry and Mapping (LOAM) [4], where *edges* and *surfaces* are detected and registered to a map through point-to-line and point-to-plane constraints within an optimization framework. In LOAM, an IMU can be optionally used to de-skew the input point cloud and provide a prior motion estimate. LOAM extensions can be found to be used specifically on ground vehicles [8] or with solid-state LiDARs [9,10]. More recently, a lightweight LOAM version named FLOAM [11] has been proposed. This is probably the closest work to our solution. In this respect, LiODOM introduces a simpler but more efficient pose optimization scheme that is based on LOAM edges, resulting into robust estimations, as it is shown in Section 6. Besides that, FLOAM uses downsampled global feature maps that are stored in 3D KD-trees. Our novel hashing-based mapping approach allows for a faster interaction with the map, especially regarding updates. Rozenberszki et al. [12] also describes a LiDAR-based approach. However, in contrast to our solution, this method assumes the existence of a predefined map and relies on a place recognition method. Finally, within this class of solutions, [13,14] propose a data fusion scheme based on an Extended Kalman Filter (EKF).

Tightly-coupled methods fuse sensor data jointly, either through optimization [5,6] or filtering [7,15]. In this regard, Ye et al. [5] introduces LIOM, a tightly-coupled odometry and mapping approach which jointly minimizes LiDAR and IMU observations in a sliding window. Despite its good performance, it is computationally expensive, making difficult its use in practical situations. In a more recent work [7], the same authors opt for an iterated Error-State Kalman Filter (ESKF), resulting into a faster solution. A recent work [6] introduces LIO-SAM as a new tightly-coupled method. In LIO-SAM, LiDAR-inertial odometry is stated as a factor graph, allowing an easy incorporation of any type of observation as a constraint, such as loop closures, GPS samples or IMU measurements. In a very recent work [16], the authors employ pre-integrated IMU readings to de-skew the input point cloud. Unlike the approaches surveyed so far, we tackle the problem of pose estimation using solely a LiDAR.

### 2.2. Vision-based solutions

In the last decades, cameras have been extensively used for mapping and localization, specially motivated by the low cost of cameras, the constant increase in computing power and the richness of the sensor data provided. These solutions are typically classified into feature-based (indirect) and direct methods, according to the strategy employed for data association across images. On the one hand, feature-based solutions utilize a set of feature points to reconstruct the world around the sensor. In this regard, Klein and Murray introduced PTAM [17], where tracking and mapping procedures run in parallel in two different threads. Following these ideas, Strasdat proposed a number of solutions [18–20] that served as a basis for ORB-SLAM [21–23], which can be considered as one of the most popular solutions in Visual SLAM. In LiODOM, we follow the computational approach adopted for these solutions, i.e. two threads run in parallel for

<sup>1</sup> <http://github.com/emiliofidalgo/liodom>.

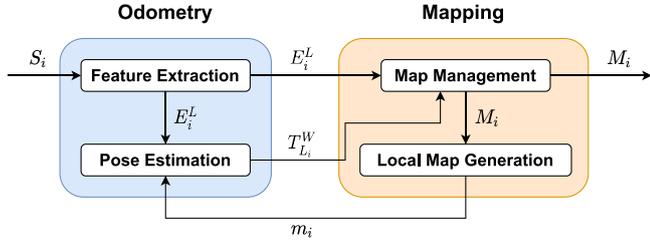


Fig. 2. Overview of LiODOM.

feature extraction and pose estimation; additionally, a set of edge points are extracted from the point cloud.

On the other hand, direct methods enable 3D reconstructions of the world of higher accuracy, since they use a larger amount of image pixels and minimize photometric errors for pose estimation. In this category, several solutions have been proposed, e.g. DTAM [24], Kinect Fusion [25], DVO [26], Elastic Fusion [27] and Kintinuous [28]. However, they typically rely on a GPU to solve the pose estimation problem, which is out of the scope of this work. In comparison with vision-based approaches, LiDAR-based solutions are less affected by illumination changes, they can provide a full overview of the surroundings and can supply reliable range estimations.

A final class of solutions combine LiDAR and cameras [29–31]. The main objective of these methods is to obtain the advantages of both types of sensors, although at the expense of additional calibration and synchronization processes that can significantly complicate their use together.

### 2.3. Map indexing

As mentioned above, establishing a set of correspondences between the input scan and a map is of prime importance for efficient pose estimation. Some authors have opted for indexing the points of a global map using a tree-based approach [4], although usually these solutions do not scale well. In our work, the global map is devised as a disjoint partition of the 3D space, and, inspired by other approaches [32,33], it is indexed using a hashing scheme. An alternative for fast data association is building a local map from a sliding window [5,6], instead of matching directly to a global map, but this option tends to discard useful correspondences. In this matter, LiODOM also introduces an adaptive local map mechanism, which can be seen as an alternative to the classical local mapping paradigm.

## 3. System overview

For a start, we define a *sweep* as a set of 360-degree 2D scans. A sweep received at time  $i$  is denoted as  $S_i$ . We also define two coordinate systems: (1)  $L$ , the LiDAR coordinate system, which is a frame attached to the geometric centre of the sensor; and (2)  $W$ , the world coordinate system, which coincides with  $L$  at the beginning. We denote by  $T_B^A \in SE(3)$  the transformation that maps a point  $p^B \in \mathbb{R}^3$  expressed in  $B$  to a point  $p^A \in \mathbb{R}^3$  expressed in  $A$ . The rotation matrix and the translation vector of  $T_B^A$  are respectively denoted by  $R_B^A \in SO(3)$  and  $t_B^A \in \mathbb{R}^3$ .

Fig. 2 illustrates LiODOM. As in other works [4,5], our solution consists of two main components, odometry and mapping, which run concurrently: while the odometry module (Section 4) computes a set of features  $E_i^L$  (LOAM edges) from  $S_i$  and estimates the current pose of the LiDAR  $T_{L_i}^W$ , the mapping module (Section 5) registers the resulting edges to a global map  $M_i$  and generates an adaptive local map  $m_i$  to be employed in the subsequent pose estimation step.

## Algorithm 1 LiDAR Odometry

**Input:**  $S_i, T_{L_{i-1}}^W, T_{L_{i-2}}^W, m_{i-1}$

**Output:**  $E_i^L, T_{L_i}^W$

- 1:  $E_i^L \leftarrow$  set of edges from  $S_i$
- 2:  $T_{L_i}^W \leftarrow \widehat{T}_{L_i}^W \equiv$  initial transformation estimate [Eq. (7)]
- 3: **for**  $n$  iterations **do**
- 4:   **for**  $p_j^{L_i} \in E_i^L$  **do**
- 5:      $p_j^W \leftarrow T_{L_i}^W p_j^{L_i}$  [Eq. (2)]
- 6:      $N(p_j^W) \leftarrow$  5 NN of  $p_j^W$  in  $m_{i-1}$
- 7:     **if**  $N(p_j^W)$  is a line **then**
- 8:       Compute  $d_e(p_j^W, l(p_j^W))$  [Eq. (3)]
- 9:       Compute residual  $\varrho_e(p_j^W, l(p_j^W))$  [Eq. (4)]
- 10:       Add residual  $\varrho_e(p_j^W, l(p_j^W))$
- 11:       to the optimization problem
- 12:     Optimize pose  $T_{L_i}^W$  [Eq. (6)]

## 4. LiDAR odometry

The LiDAR odometry module is organized into two synchronized execution threads that, hence, decouples feature extraction from pose estimation. Both are described next.

### 4.1. Feature extraction

For a start, each sweep  $S_i$  is divided into its different scans, discarding at the same time those points whose range do not fall within a certain interval  $[r_{\min}, r_{\max}]$ , which has to be configured accordingly to the sensor operating and noise characteristics. Each pre-processed scan is next considered, selecting a number of key points to reduce the computational requirements. In this work, we make use of LOAM edge features, given the utility they have shown in other works and their simpler computation [4,8,11,34]. Besides, to select the best features, in LiODOM we also calculate a local curvature measure  $c_j$  for each point  $p_j^{L_i}$  [4] as follows:

$$c_j = \frac{1}{|\Omega| \cdot \|p_j^{L_i}\|} \sum_{k \in \Omega, k \neq j} \|p_j^{L_i} - p_k^{L_i}\|, \quad (1)$$

being  $\Omega$  a set of consecutive (i.e. co-planar) points in the vicinity of  $p_j^{L_i}$  belonging to the same scan of the sweep  $S_i$ . In this way, we take advantage of the typical larger intra-beam resolution against the inter-beams resolution, to get more distinctive features [11]. Moreover, to distribute edges throughout the environment, a scan is further divided into equally-sized sectors, and a maximum number of edges is set for every sector. Unlike [4], we split each scan into 8 sectors and choose a maximum of 10 edges per sector after sorting them in decreasing order of curvature  $c$ . Furthermore, the selection applies non-maxima suppression, i.e. a point is chosen as an edge if none of its neighbours has been already selected. The result of this procedure is a set of edges  $E_i^L$  chosen from sweep  $S_i$ .

### 4.2. Pose optimization

Let us consider the transformation  $T_{L_i}^W$  from the LiDAR at time  $i$  to the world. Then, every point  $p_j^{L_i} \in E_i^L$  projects into the world frame  $W$  as:

$$p_j^W = T_{L_i}^W p_j^{L_i} = R_{L_i}^W p_j^{L_i} + t_{L_i}^W, \quad (2)$$

being  $R_{L_i}^W$  and  $t_{L_i}^W$  the respective rotation matrix and translation vector of  $T_{L_i}^W$ . We denote the set of transformed edges as  $E_i^W$ . Subsequently, a set of point-to-line correspondences between  $E_i^W$  and a local map are computed for pose estimation. We have opted for this solution rather than using, for instance, a global map, because it turns out to be more computationally stable as more frames are processed. In LiODOM, that local map is not built after pose estimation in a sequential way as in [5,6], but it is built concurrently with pose estimation by the mapping module (see Section 5).

Let us assume now the existence of a local map  $m_{i-1}$  at time  $i-1$ , which is a subset of the global map  $M_{i-1}$ . This map  $m_{i-1}$  contains the points in  $M_{i-1}$  closest to the LiDAR according to the latest pose estimate  $T_{L_{i-1}}^W$ . For each point  $p_j^W \in E_i^W$ , we obtain the  $k$  nearest points in  $m_{i-1}$ , where  $k=5$  in this work. We now denote this set as  $(N(p_j^W))$  and the  $n$ th nearest neighbour of  $p_j^W$  as  $(N_n(p_j^W))$ . Next, we assess whether points in  $(N(p_j^W))$  are aligned by analysing their scatter matrix [9]. If the largest eigenvalue of this matrix is, at least, three times the second largest eigenvalue, we consider that a valid point-to-line correspondence can be established between  $p_j^W$  and the line  $l(p_j^W)$  resulting from  $(N_1(p_j^W))$  and  $(N_2(p_j^W))$ . We then calculate the point-to-line distance  $d_e$  as

$$d_e(p_j^W, l(p_j^W)) = \frac{\|(p_j^W - N_1(p_j^W)) \times N_{12}\|}{\|N_{12}\|}, \quad (3)$$

with  $N_{12} = N_1(p_j^W) - N_2(p_j^W)$ . At this point, it is worth noting that increasing  $k$  leads to higher computational times for calculating the scatter matrix, without necessarily ensuring additional insight on whether the nearby points lie on a line in a local sense.

LiODOM, as an odometer, optimizes only the current pose of the LiDAR  $T_{L_i}^W$ . Within the optimization framework, each correspondence provides a constraint between  $T_{L_i}^W$  and the local map  $m_{i-1}$ , whose residual  $q_e$  is computed as:

$$q_e(p_j^W, l(p_j^W)) = \omega_j d_e(p_j^W, l(p_j^W)), \quad (4)$$

where  $\omega_j$  is a weighting term defined as:

$$\omega_j = 1 - \frac{r_j - r_{\min}}{r_{\max} - r_{\min}}, \quad (5)$$

being  $r_j$  the range returned by the LiDAR for edge  $p_j^L$ . The rationale behind this factor is that LiDARs tend to decrease their accuracy at longer distances and, thus, we give more importance to correspondences established at closer distances. We then compute the optimal transformation  $T_{L_i}^W$  as the minimizer of the loss function  $J(\tilde{T}_{L_i}^W, \mathcal{Y})$ :

$$J(\tilde{T}_{L_i}^W, \mathcal{Y}) = \frac{1}{2} \sum_{j \in \mathcal{Y}} \rho \left( \left\| q_e \left( \tilde{T}_{L_i}^W p_j^L, l \left( \tilde{T}_{L_i}^W p_j^L \right) \right) \right\|^2 \right) \\ T_{L_i}^W = \min_{\tilde{T}_{L_i}^W} J(\tilde{T}_{L_i}^W, \mathcal{Y}) \quad (6)$$

where  $\mathcal{Y}$  is the set of correspondences established between  $E_i^W$  and the local map  $m_{i-1}$ , and  $\rho$  is a Huber loss function to reduce the influence of outliers. The system of non-linear equations is solved by means of the Levenberg–Marquardt algorithm using the Ceres Solver library [35], using the transformation  $\hat{T}_{L_i}^W$  as initial guess:

$$\hat{T}_{L_i}^W = T_{L_{i-1}}^W \hat{T}_{L_i}^{L_{i-1}} \\ = T_{L_{i-1}}^W T_{L_{i-1}}^{L_{i-2}} = T_{L_{i-1}}^W \left( T_{L_{i-2}}^W \right)^{-1} T_{L_{i-1}}^W, \quad (7)$$

i.e. we assume the same motion as for the previously estimated pose. Although LiODOM deals only with LiDAR data, it is clear

that any additional motion estimate, e.g. from an IMU, can be incorporated at this point.

The full LiDAR odometry procedure is stated algorithmically in Alg. 1. In our experiments, 1 or 2 refining iterations are enough, i.e.  $n=2$  or 3 at line 3 of Alg. 1.

## 5. LiDAR mapping

The registration of the extracted edges  $E_i^L$  on the global map  $M_i$  is performed by the mapping module using the last optimized pose  $T_{L_i}^W$ . This module also generates the corresponding local map  $m_i$  as described next.

### 5.1. Map representation

Given the high frequency at which the map must be accessed, the type of data structure chosen to represent 3D space becomes crucial for fast operation. A single KD-tree has been typically used to this end [4]. However, this option presents several drawbacks: on the one hand, the full tree tends to change as points are added or deleted to/from the tree, and, on the other hand, the KD-tree performance decreases as more points need to be managed [4]. To overcome these issues, in LiODOM we introduce an efficient hashing data structure for representing the map, taking inspiration from other recent works [32,33]. To be more specific, the 3D space is partitioned into a set of disjoint cuboids of a fixed size that we name *cells*. A cell  $C_j$  is represented by its geometric centre, denoted by  $(c_{jx}, c_{jy}, c_{jz})$ , and includes all 3D points whose coordinates fall into its limits. We define a map at time  $i$  as  $M_i = \{\mathbb{H}_i, \mathbb{C}_i\}$ , where  $\mathbb{H}_i$  is a hash table and  $\mathbb{C}_i$  is the set of existing cells up to time  $i$ . The table  $\mathbb{H}_i$  allows us to rapidly get access to a specific cell  $C_j$  using a hash function of its coordinates, defined by:

$$H(C_j) = (c_{jx} \oplus (c_{jy} \ll 1)) \oplus (c_{jz} \ll 2), \quad (8)$$

where  $\oplus$  and  $\ll$  are, respectively, the bitwise XOR and the left shift operators. This function has been selected in order to minimize, as much as possible, hash collisions. That is to say, if bits of a binary word have roughly 50% chance of being 0 or 1, i.e. as randomly distributed as possible, the bitwise XOR between such binary words results into another word also following a random distribution. Furthermore, since the bitwise XOR is a symmetric operation, the order of the elements in the hash code is lost. To break this symmetry, we use the shift operator, at a limited computational cost.

### 5.2. Map updates

In LiODOM, map updates are performed once per sweep, being the set of edges  $E_i^L$ , extracted from  $S_i$ , and the last optimized transformation  $T_{L_i}^W$  the input data. Unlike other approaches [4], where the raw point cloud is used for mapping, in our approach, the map is built using directly the edges to speed up the mapping procedure, resulting into more sparse maps. Initially, every point  $p_j^L \in E_i^L$  is transformed to world coordinates using  $T_{L_i}^W$  and (2). Next, for each point  $p_j^W = (x, y, z)$ , we compute the geometric centre of the cell  $C_q$  in which the point should be stored as:

$$\begin{bmatrix} c_{qx} \\ c_{qy} \\ c_{qz} \end{bmatrix} = \begin{bmatrix} \lfloor x/s_{xy} \rfloor s_{xy} + \frac{1}{2} s_{xy} \\ \lfloor y/s_{xy} \rfloor s_{xy} + \frac{1}{2} s_{xy} \\ \lfloor z/s_z \rfloor s_z + \frac{1}{2} s_z \end{bmatrix}, \quad (9)$$

where  $s_{xy}$  and  $s_z$  are the metric cell sizes for the corresponding dimension. We next check if the cell  $C_q$  is already in the map

**Algorithm 2** LiDAR Mapping

---

**Input:**  $E_i^L, T_{L_i}^W, M_{i-1}$   
**Output:**  $M_i, m_i$

- 1:  $M_i \leftarrow M_{i-1}$
- 2: **for**  $p_j^{L_i} \in E_i^L$  **do**
- 3:    $p_j^W \leftarrow T_{L_i}^W p_j^{L_i}$  [Eq. (2)]
- 4:    $C_q \leftarrow$  cell where  $p_j^W$  should be [Eq. (9)]
- 5:   **if**  $H(C_q) \notin \mathbb{H}$  **then**
- 6:     Create new cell  $C_n$  using  $C_q$  coordinates
- 7:     Add  $p_j^W$  to  $C_n$
- 8:     Update  $M_i$  adding  $C_n$  to  $\mathbb{C}$
- 9:     Update  $M_i$  adding  $H(C_q)$  to  $\mathbb{H}$
- 10:   **else**
- 11:     Retrieve cell  $C_q$  using  $H(C_q)$
- 12:     Update  $M_i$  adding  $p_j^W$  to  $C_q$
- 13:     **if**  $C_q$  has more points than  $\tau$  **then**
- 14:       Update  $M_i$  filtering  $C_q$  using a 3D voxel grid
- 15:  $C_{L_i} \leftarrow$  cell where the LiDAR should be [Eq. (9)]
- 16:  $m_i \leftarrow \emptyset$
- 17: **for**  $C_i \in$  Neighbours of  $C_{L_i}$  in  $M_i$  **do**
- 18:    $m_i \leftarrow m_i \cup C_i$
- 19: **for**  $j \in \{1, 2, 3\}$  **do**
- 20:    $m_i \leftarrow m_i \cup E_{i-j}^W$

---

by querying the hash table  $\mathbb{H}$  using the key  $H(C_q)$ . If this is the case, the point is added to the existing cell. Otherwise, a new cell  $C_n$  is created with point  $p_j^W$  as seed, to be added next to  $\mathbb{C}$  and indexed on  $\mathbb{H}$  by  $H(C_n)$ . Finally, modified cells exceeding a certain number of points are filtered using a 3D voxel grid. Note that our data structure allows us to rapidly update just the required areas of the environment, avoiding the update of the whole map at each iteration. This fact contributes to speed up the mapping procedure, as will be shown in the experiments.

### 5.3. Adaptive local map computation

Lastly, the mapping module generates a local map  $m_i$ , which contains the points of  $M_i$  within a certain range from the current LiDAR pose. Assuming a moderate motion between two consecutive sweeps, these points are enough to find correspondences for the next pose estimation step. To build the local map, we first retrieve the cell  $C_{L_i}$  where the LiDAR is located at that moment using its current position  $T_{L_i}^W$  and (9). Next, assuming a 3D grid arranged over  $M_i$ , neighbouring cells of  $C_{L_i}$  up to a certain distance are further retrieved from  $M_i$ , and their corresponding points are merged to form the local map  $m_i$ . This operation results to be very fast due to the proposed hashing structure. Points on  $m_i$  are finally organized into a KD-tree to speed up nearest neighbour search. Note that this tree is very simple, as it just contains a small subset of the total map points, in contrast to managing the whole global map [4].

On the other side, we refer to this local map as *adaptive* since it always covers a specific area of the environment, contrary to a local map built by aggregation of a sliding window [5,6]. Besides, it provides us with correspondences with revisited areas of the environment in a natural way. Additionally, the availability of  $m_i$  avoids us to search for correspondences against the whole map, as done by other solutions [4]. Finally, to avoid reduced amounts of points from unexplored areas, we always add the last three sweeps to  $m_i$ . The complete mapping procedure is outlined in Alg. 2.

**Table 1**

Parameters values and section where they are defined.

$r_{\min}$ (Section 4.1)	3.0	$k$ (Section 4.2)	5
$r_{\max}$ (Section 4.1)	75.0	$S_{xy}$ (Section 5.2)	25
Scan sectors (Section 4.1)	8	$S_z$ (Section 5.2)	20
Edges per sector (Section 4.1)	10	Sweeps in $m_i$ (Section 5.3)	3

## 6. Experimental results

In this section, we report on the results of several experiments conducted to evaluate LiODOM, including a comparison with other solutions. A laptop featuring an Intel Core i7-10750H @2.6 GHz, 16 GB RAM has been used in all cases.

### 6.1. Methodology

We validate our approach using the KITTI odometry benchmark [36], as usual in the field [11,37,38]. This dataset consists of 22 sequences collected using a Velodyne HDL-64E sensor. Eleven of these sequences include GPS poses that can be used as ground truth. The average translational (%) and rotational (deg/100 m) errors are adopted in the following as main performance metrics. We additionally consider the Absolute Trajectory Error (ATE), although it rather focuses on the global consistency of the whole trajectory and thus it is more appropriate for SLAM systems.

To further validate LiODOM, we compare it with other pure LiDAR-based odometry and also with SLAM solutions, namely FLOAM [11], ISC-LOAM [34] and LeGo-LOAM [8]. We are aware of the existence of recent fusion-based [5,6] or even semantic-aided [37] solutions. They are not considered in this evaluation since, in contrast to our method, they imply additional complexities, such as synchronization and calibration procedures or increasing computational resources.

### 6.2. Algorithm configuration

LiODOM has been run several times against the K05 sequence, fine-tuning them towards increasing the localization accuracy, in order to find a suitable set of parameters. The values obtained for the most relevant parameters can be found in Table 1. They have been kept constant in all the remaining experiments, what, given the performance achieved, shows that this configuration works reasonably well and is able to tolerate different operating conditions: two different LiDAR devices, different processors and different datasets taken from different environments. This set-up represents, thus, a good trade-off between performance and accuracy. Nonetheless, those parameters with an impact on the computational complexity could need specific modifications depending on the available computational resources in order to keep response times at a reasonable value.

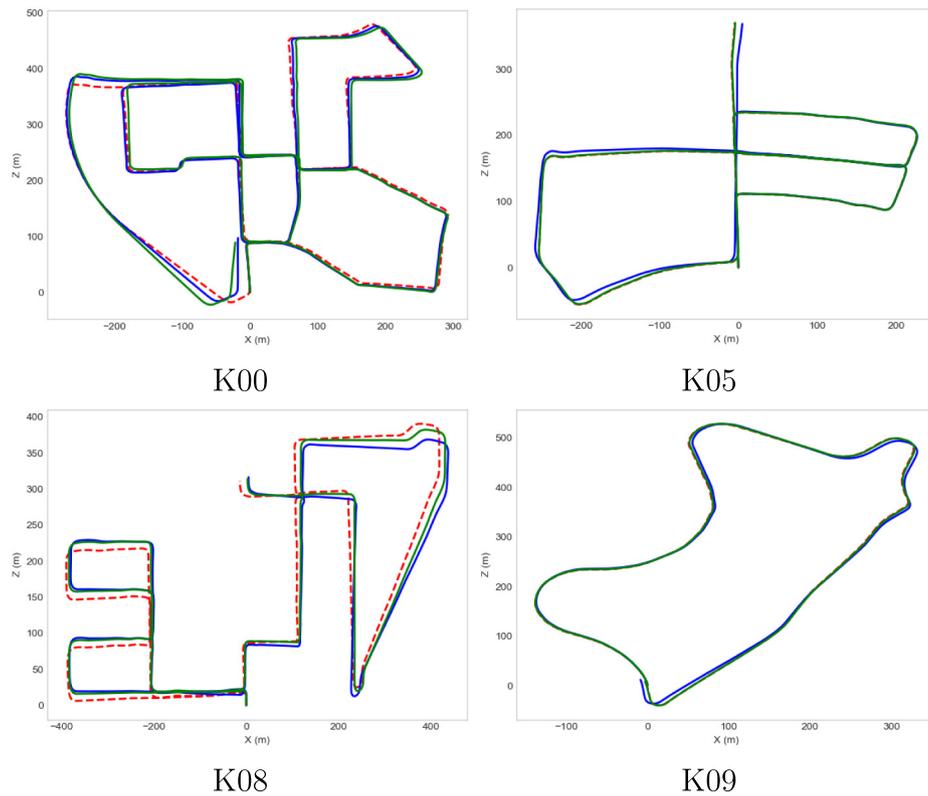
### 6.3. Odometry performance

Table 2 summarizes the results obtained in terms of translational and rotational errors. Results for FLOAM were obtained by ourselves using its open source implementation, while results for ISC-LOAM and LeGo-LOAM are directly reported from, respectively, [37,38]. As can be observed, LiODOM achieves competitive results in all sequences in terms of translation error. This can be observed even in sequences comprising loop closures, such as K05, K06 and K07, where our approach achieves the second best results, sometimes very close to complete SLAM solutions like ISC-LOAM. We obtain, on average, 1.038% drift in translation, outperforming the other solutions in this matter. Regarding rotation error, again our solution leads to the lowest errors in most of the

**Table 2**

Average translational and rotational errors for the KITTI odometry benchmark. Best results are shown in bold red and second best in blue.

	Translational error (%)				Rotational error (deg/100 m)			
	FLOAM	ISC-LOAM	LeGO	Ours	FLOAM	ISC-LOAM	LeGO	Ours
K00	0.861	1.020	2.170	<b>0.857</b>	0.349	0.420	1.050	<b>0.348</b>
K01	1.309	2.920	13.400	<b>1.301</b>	<b>0.128</b>	0.630	1.020	<b>0.129</b>
K02	0.952	1.670	2.170	<b>0.947</b>	<b>0.310</b>	0.540	1.010	<b>0.309</b>
K03	1.267	<b>1.150</b>	2.340	1.262	0.227	0.720	1.180	<b>0.226</b>
K04	1.417	1.500	<b>1.270</b>	1.411	0.010	0.560	1.010	<b>0.009</b>
K05	0.835	<b>0.810</b>	1.280	0.834	0.360	0.370	0.740	<b>0.359</b>
K06	0.835	<b>0.760</b>	1.060	0.834	0.332	0.410	0.630	<b>0.331</b>
K07	0.883	<b>0.560</b>	1.120	0.881	0.617	<b>0.430</b>	0.810	<b>0.614</b>
K08	0.869	1.200	1.990	<b>0.864</b>	0.332	0.500	0.940	<b>0.331</b>
K09	1.033	1.400	1.970	<b>1.029</b>	<b>0.317</b>	0.590	0.980	<b>0.318</b>
K10	1.203	1.870	2.210	<b>1.196</b>	<b>0.287</b>	0.620	0.920	<b>0.288</b>
Average	1.042	1.351	2.816	<b>1.038</b>	0.297	0.526	0.935	<b>0.296</b>



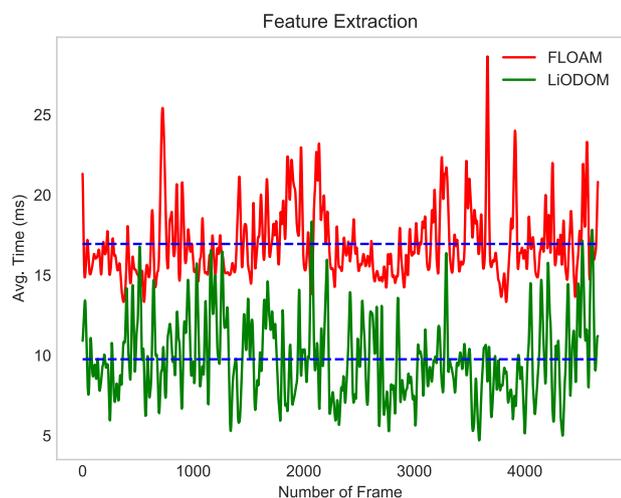
**Fig. 3.** Examples of trajectories estimated for some sequences of the KITTI odometry benchmark. The ground truth is shown as a red dashed line, while FLOAM and LiODOM estimates are respectively shown as blue and green lines. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

sequences. On average, the rotational error of LiODOM is 0.296% deg/100 m, which represents again the lowest average error.

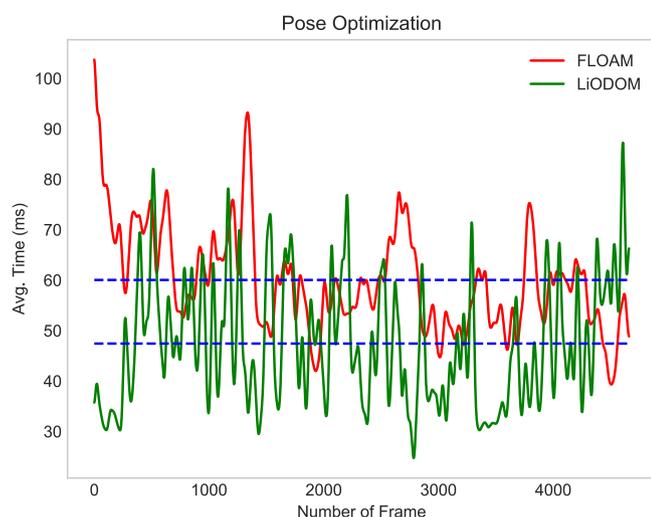
Table 3 reports on the ATE for the KITTI sequences that contain loop closures. Again, results for FLOAM were obtained by ourselves, while results for ISC-LOAM and LeGo-LOAM are reported from, respectively, [37,39]. LiODOM again achieves competitive results in all sequences despite it is actually a pure odometry system and, therefore, does not take any advantage from global map optimization nor from loop closures. In contrast to other solutions based on a sliding window policy, our pose estimation procedure makes use of an adaptive local map covering, in all directions, the surroundings of the LiDAR at its current position. This strategy allows our approach to find associations/matchings in a wider area than just the one covered by the previous frames, hence effectively reducing the pose estimation error. On average,

the ATE for LiODOM results to be 3.535 m, which represents the second best performance among the different methods considered. By way of illustration, Fig. 3 shows the resulting trajectory estimates from our approach and from FLOAM for several KITTI sequences.

In the following, we analyse the computational complexity of LiODOM. To this end, we choose dataset K02, the largest dataset considered in this work. Average processing times for every odometry stage can be found in Figs. 4 and 5, as they are the stages that operate online within a real platform. The capabilities of the mapping module, which is executed as a standalone procedure, are evaluated in the next section. In both figures, we compare LiODOM with FLOAM as a representative of the solutions based on LOAM edges and surfaces as features [4,8,11,34] (notice that the other solutions considered in this section can also be



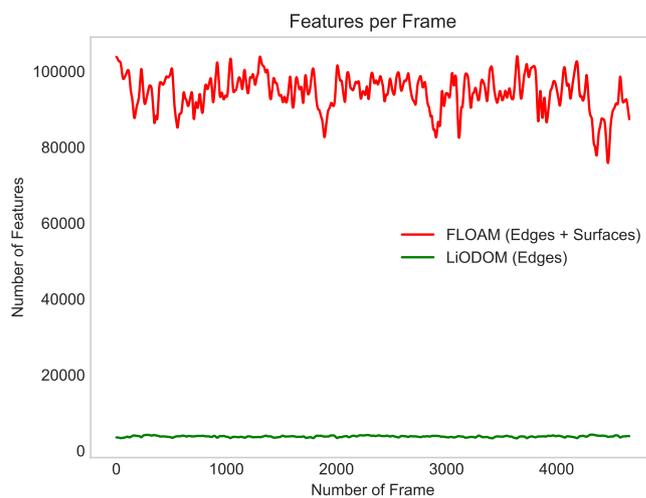
**Fig. 4.** Average processing times of the feature extraction stages of FLOAM and LiODOM. The blue dashed line corresponds to the mean value for each case.



**Fig. 5.** Average processing times of the pose optimization stages of FLOAM and LiODOM. The blue dashed line corresponds to the mean value for each case.

regarded as derivatives of LOAM). Regarding feature extraction, Fig. 4 shows that LiODOM presents lower extraction times than FLOAM. This stage takes, on average, 10.36 ms in contrast to the 16.92 ms required by FLOAM. Alternatively, Fig. 5 shows pose optimization times for each, including the search for correspondences within the respective maps. As can be observed, LiODOM converges, in general, faster than FLOAM, taking, on average, 48.73 ms in contrast to 60.12 ms. Considering the measured feature extraction and pose optimization times, the resulting overall processing times are 59.09 ms per frame for LiODOM and 77.04 ms per frame for FLOAM, which means, respectively, frame rates of around 17 Hz and 13 Hz. This means that LiODOM is about 1.3 times faster than FLOAM.

Fig. 6 reports on the number of features per frame that each solution has to handle. As can be observed, the differences between both approaches are more evident in this matter, since LiODOM only uses edges and, therefore, the required number of features is considerably less, with the corresponding reduction in memory requirements without noticeable differences in performance, slightly outperforming FLOAM for a number of datasets (as shown in Tables 2 and 3).



**Fig. 6.** Number of features per frame.

**Table 3**

Absolute Trajectory Error (m) for the KITTI dataset. Best results are shown in bold red and second best in blue.

	FLOAM	ISC-LOAM	LeGO	Ours
K00	5.137	<b>1.600</b>	6.300	7.135
K02	9.294	<b>4.770</b>	14.700	9.754
K05	2.546	<b>2.490</b>	2.800	<b>0.322</b>
K06	<b>0.934</b>	1.030	<b>0.800</b>	0.956
K07	<b>0.498</b>	<b>0.560</b>	0.700	1.518
K08	4.344	4.880	<b>3.500</b>	4.592
K09	2.144	2.310	<b>2.100</b>	<b>0.470</b>
Average	3.557	<b>2.520</b>	4.414	<b>3.535</b>

**Table 4**

Absolute Trajectory Error (m) for the KITTI dataset with and without the weighting term.

	With	Without
K00	7.135	8.488
K02	9.754	11.613
K05	0.322	4.057
K06	0.956	1.024
K07	1.518	1.638
K08	4.592	5.780
K09	0.470	3.632
Average	3.535	5.176

To finish, we consider the effect of the weighting term defined in Eq. (5), which is related to the range returned by the LiDAR, as this results to be a key component of the pose optimization procedure of LiODOM. In this regard, Table 4 reports on again the ATE for the KITTI sequences that contain loop closures, but with and without using the weighting factor. As can be observed, the weighting term improves the performance of LiODOM in all cases: all in all, the use of the weighting factor results into an average ATE of 3.535 m, in contrast to LiODOM without using this term, which leads to an average ATE of 5.176 m.

#### 6.4. Mapping performance

In this section, we report on several results intended to assess the efficiency of the mapping approach adopted in LiODOM. In a first experiment, we measure the times required to update the global map and to build the local maps using our hashing-based data structure and a KD-tree. The K05 sequence was chosen in this case for computational reasons. The results are shown

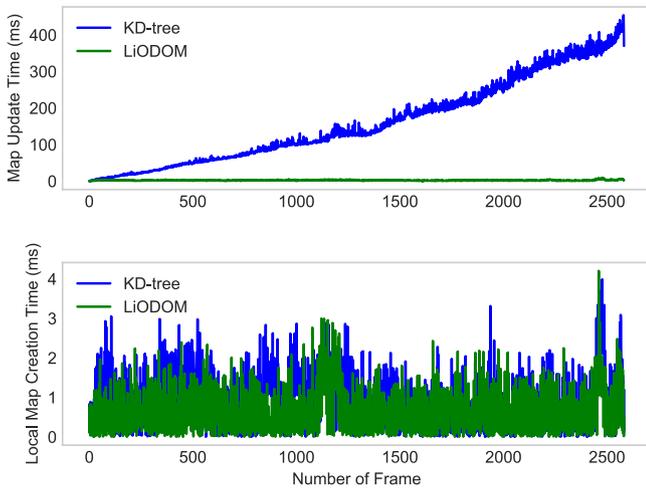


Fig. 7. Performance of the LiODOM mapping structure vs. a KD-tree.

Table 5

Entropy values for the hash table entry lengths for the KITTI sequences including loops.

	Baseline	Eq. (8) (LiODOM)
K00	4.181	5.287
K02	4.409	5.645
K05	3.334	4.933
K06	3.766	4.464
K07	3.407	4.395
K08	3.457	5.321
K09	3.259	4.889
Average	3.688	4.990

in Fig. 7. As can be observed, the time required to update the global map by our approach remains roughly constant along the whole sequence. Contrarily, the running times for the KD-tree approach grow as more frames are processed, which can lead to an impractical operation. This behaviour can be attributed to the fact that, unlike our approach, the whole tree needs to be rebuilt on each update. The differences are less evident as for the times required to build the local maps, where both approaches are very fast, although our approach seems to perform slightly better.

As a last experiment, in this section, we evaluate the quality of LiODOM hash function and the frequency of collisions when storing map points. In the ideal case, the hash function should be such that the number of elements associated to each bucket, i.e. the hash table entry lengths, follow a uniform distribution, what means that the hash function spreads reasonably well the elements between the table entries and therefore the number of collisions when accessing the map for updates is at a minimum. Given a distribution probability, its entropy  $H$  can be used to check how close it is to the uniform distribution. For the case of the hash table,  $H$  would be defined as:

$$H(m) = - \sum_{k=1}^n p(k) \log(p(k)), \quad (10)$$

being  $n$  the number of buckets in the hash table and  $p(k)$  the probability that a new map point lies at the  $k$ th table entry.  $p(k)$  can thus be calculated as the number of data items associated to the entry divided by the total number of items stored in the table. The higher the value of the entropy  $H$ , the more uniformly are spread the points in the hash table and, therefore, the lower the probability of collision. Under this context, in this experiment, we

Table 6

ATE (m) and velocity RMSE (m/s) for the experiments on-board a UAV.

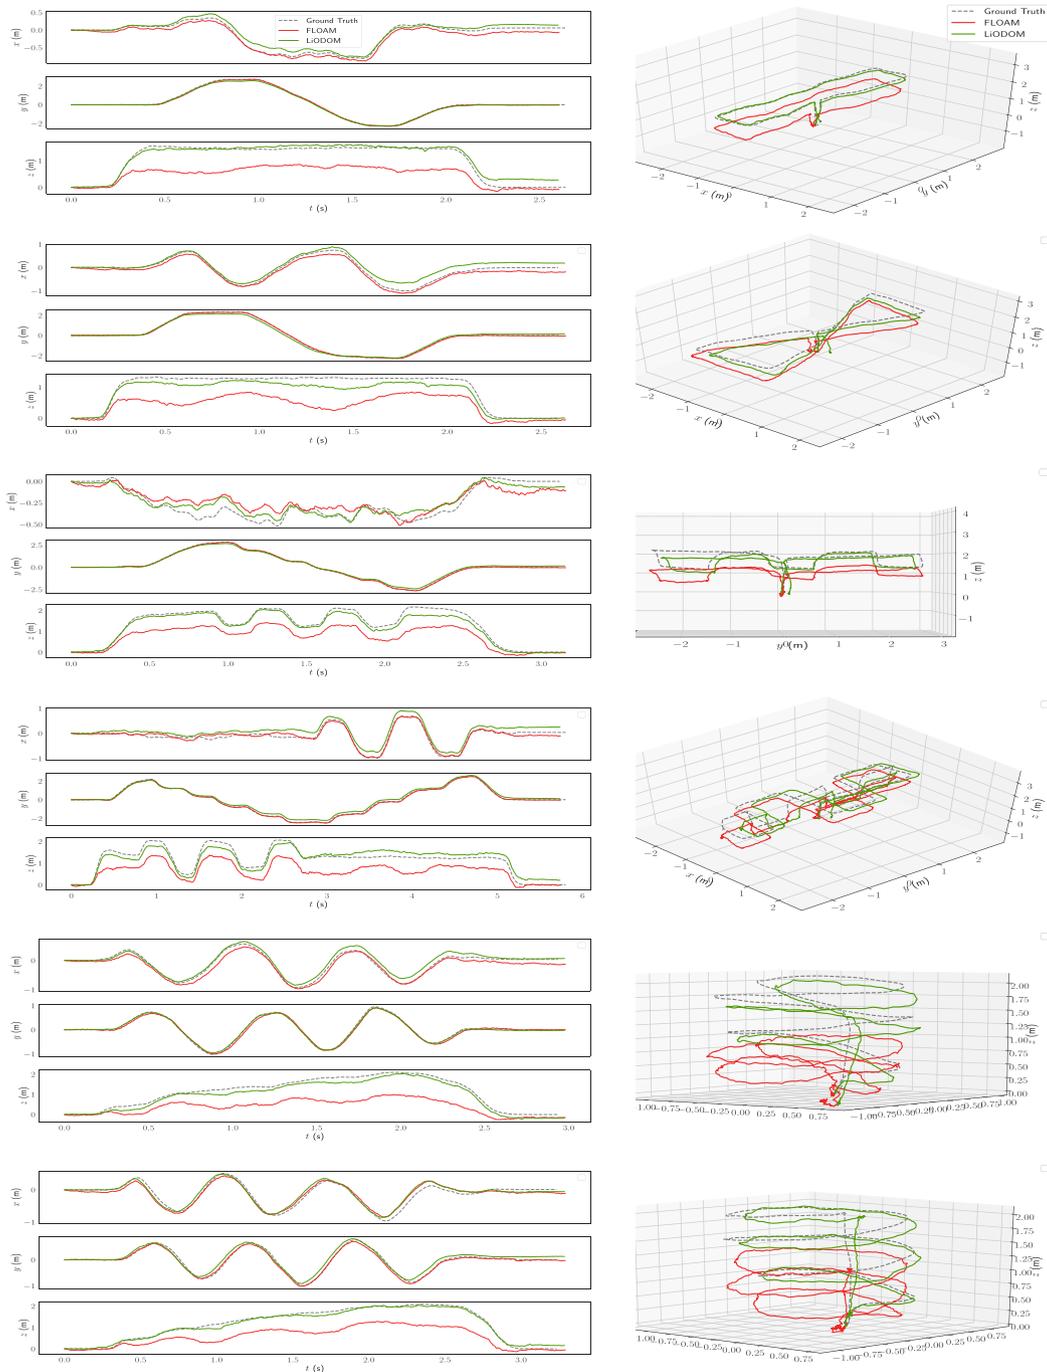
Experiment	ATE		Velocity RMSE (x / y / z)
	LiODOM	FLOAM	
1	0.176776	0.657013	0.042/0.051/0.075
2	0.276960	0.613781	0.071/0.079/0.101
3	0.204852	0.637984	0.048/0.121/0.141
4	0.314232	0.572278	0.088/0.085/0.129
5	0.199806	0.816480	0.086/0.108/0.066
6	0.164287	0.675760	0.063/0.068/0.074

compute the entropy  $H$  of the global map generated by LiODOM after processing a full sequence. We compare the LiODOM hash function, i.e. Eq. (8), with a baseline consisting on simply adding the hashes of the three cell coordinates. We employ again the sequences of the KITTI dataset including loops, since they are typically more prone to produce collisions given that they revisit previously seen places. The results are shown in Table 5. As can be observed, LiODOM in combination with the proposed hash function leads to higher entropy values than the alternative, simpler function in all cases.

### 6.5. Experiments on-board an aerial platform

Finally, we also report on some experiments involving an aerial platform intended for visual inspection tasks [40]. This platform has been recently fitted with an Ouster OS1-64 3D laser scanner that feeds LiODOM. The performance evaluation has been carried out inside the Aerial Robotics laboratory of the University of the Balearic Islands, which is equipped with an OptiTrack Motion Capture system (MOCAP) that supplies ground truth data during the tests. Table 6 shows the ATE obtained using LiODOM and FLOAM for six different experiments mixing diverse motion in the three axes. As can be noticed, the ATE values from LiODOM range from 16 to 31 cm, indicating that position estimates closely resemble the ground truth. The ATE values from FLOAM are higher for all six experiments, ranging between 57 and 68 cm. On the other side, Fig. 8 compares, against the ground truth and along the respective trajectories, the position estimates of LiODOM and FLOAM for all experiments. In this case, estimates for the X and Y axes mostly coincide with the ground truth, while, as also happens for other LiDAR-based odometry frameworks [8,11,34], some drift can be appreciated in the Z-axis estimates. Compared to the position estimates provided by FLOAM, the LiODOM estimates are closer to the ground truth, which is consistent with the lower ATE values shown in Table 6.

Within this robotic system, developed under the Supervised Autonomy paradigm, LiODOM is expected to supply not only pose estimates, but also velocity estimates, which constitute the basis for platform control in this case. To show the performance of LiODOM in this regard, Table 6 additionally reports on the velocity estimation results against the ground truth for the same six experiments as above, in the form of Root Mean Square Errors (RMSE) separately for each axis. The reported values indicate a very high accuracy in the estimation of X and Y velocities, and a slightly larger error for the Z axis. To finish, Fig. 9 compares graphically the vehicle velocities estimated by LiODOM with the values provided by the MOCAP for the six experiments. As also observed for the position estimates, the X and Y velocity estimates coincide almost perfectly with the ground truth, while the Z-axis estimates present a slightly larger error.



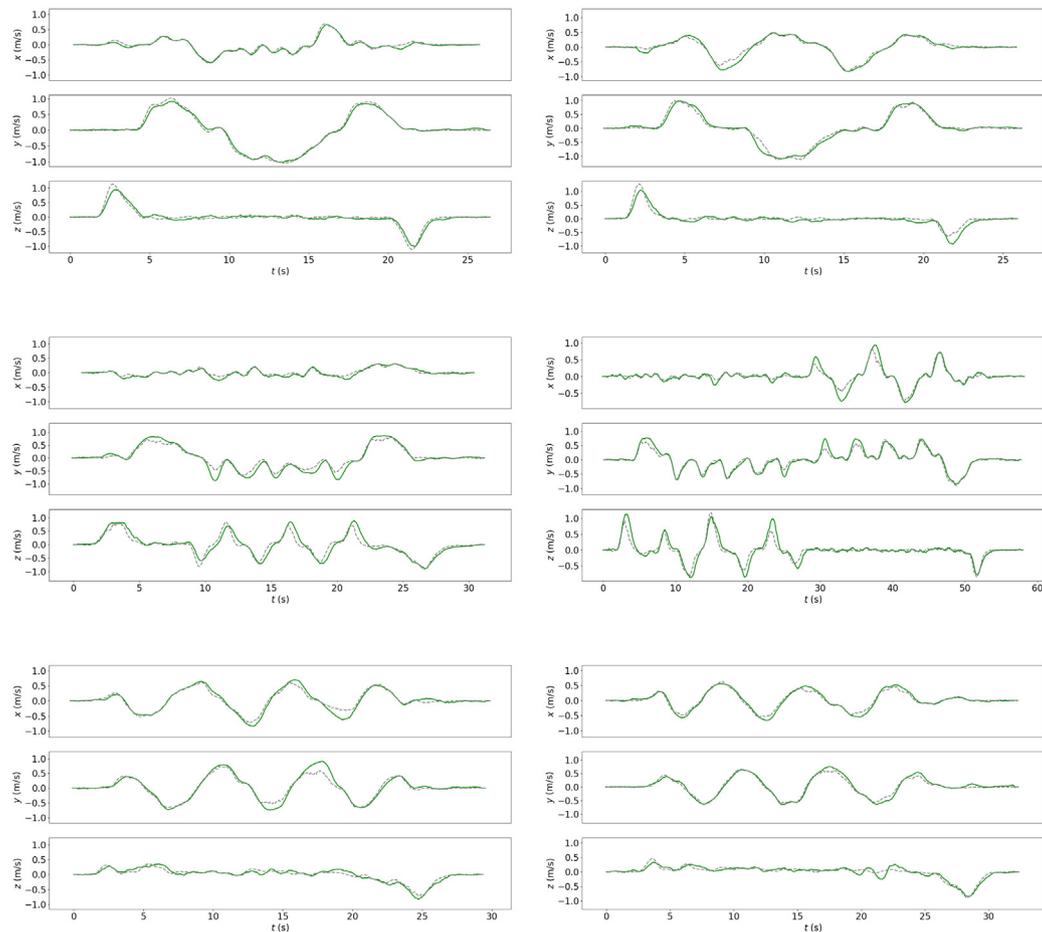
**Fig. 8.** Position estimates and trajectories for the six experiments running LiODOM (green) and FLOAM (red) on-board an aerial platform. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

### 7. Conclusions and future work

This paper proposes LiODOM, a novel LiDAR-only odometry and mapping approach. Our solution fundamentally consists of two parts working concurrently: (1) the odometry module, which is in charge of extracting a set of edges from the input sweep and estimating the current pose of the LiDAR; and (2) the mapping module, which builds and maintains a global map of the environment, and also generates a local map employed for pose estimation. Pose estimation is conceived as an optimization problem which involves a set of weighted point-to-line constraints

between the current sweep and a local map. We have also described a data structure based on a hashing scheme which allows us to rapidly get access to any part of the map and manage it in an efficient way. Furthermore, this structure is also employed to obtain an adaptive local map, used to facilitate data association. Our experiments show that LiODOM compares favourably against other state-of-the-art approaches, and that it can be used for both position and velocity estimation.

Despite its good performance, LiODOM is an odometer and unavoidably drifts. Therefore, we will consider extending the ideas proposed in this paper to develop a complete SLAM/3D



**Fig. 9.** Velocity estimates for the experiments running LiODOM on-board an aerial platform. LiODOM estimates are shown in green, while the ground truth is shown as a grey dashed plot.

reconstruction system, incorporating other motion estimation sensors into a fusion scheme for enhanced performance.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Acknowledgements

This work is partially supported by EU-H2020 project BUG-WRIGHT2 (GA 871260) and by project PGC2018-095709-B-C21 (funded by MCIN/AEI/10.13039/501100011033 and ERDF “A way of making Europe”). This publication reflects only the authors views and the European Union is not liable for any use that may be made of the information contained therein.

### Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.robot.2022.104226>.

### References

- [1] R. Mur-Artal, J.D. Tardos, ORB-SLAM2: An open-source SLAM system for monocular, stereo, and RGB-D cameras, *IEEE Trans. Robot.* 33 (5) (2017) 1255–1262.
- [2] M. Ferrera, A. Eudes, J. Moras, M. Sanfourche, G. Le Besnerais, OV<sup>2</sup>SLAM: A fully online and versatile visual SLAM for real-time applications, *IEEE Robot. Autom. Lett.* 6 (2) (2021) 1399–1406.
- [3] P. Besl, N.D. McKay, A method for registration of 3-D shapes, *IEEE Trans. Pattern Anal. Mach. Intell.* 14 (2) (1992) 239–256.
- [4] J. Zhang, S. Singh, LOAM: Lidar odometry and mapping in real-time, in: *Robot.: Sci. Syst.*, Vol. 2, 2014.
- [5] H. Ye, Y. Chen, M. Liu, Tightly coupled 3D lidar inertial odometry and mapping, in: *IEEE Int. Conf. Robot. Autom.*, 2019, pp. 3144–3150.
- [6] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, R. Daniela, LIO-SAM: Tightly-coupled lidar inertial odometry via smoothing and mapping, in: *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2020, pp. 5135–5142.
- [7] C. Qin, H. Ye, C.E. Pranata, J. Han, S. Zhang, M. Liu, LINS: A lidar-inertial state estimator for robust and efficient navigation, in: *IEEE Int. Conf. Robot. Autom.*, 2020, pp. 8899–8906.
- [8] T. Shan, B. Englot, LeGO-LOAM: Lightweight and ground-optimized lidar odometry and mapping on variable terrain, in: *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2018, pp. 4758–4765.
- [9] J. Lin, F. Zhang, Loam livox: A fast, robust, high-precision LiDAR odometry and mapping package for LiDARs of small FoV, in: *IEEE Int. Conf. Robot. Autom.*, 2020, pp. 3126–3131.
- [10] K. Li, M. Li, U.D. Hanebeck, Towards high-performance solid-state-LiDAR-inertial odometry and mapping, *IEEE Robot. Autom. Lett.* 6 (3) (2021) 5167–5174.
- [11] H. Wang, C. Wang, C. Chen, L. Xie, F-LOAM: Fast LiDAR odometry and mapping, in: *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2020.
- [12] D. Rozenberszki, A.L. Majdik, LOL: Lidar-only odometry and localization in 3D point cloud maps, in: *IEEE Int. Conf. Robot. Autom.*, 2020, pp. 4379–4385.
- [13] S. Yang, X. Zhu, X. Nian, L. Feng, X. Qu, T. Ma, A robust pose graph approach for city scale LiDAR mapping, in: *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2018, pp. 1175–1182.
- [14] M. Demir, K. Fujimura, Robust localization with low-mounted multiple LiDARs in urban environments, in: *IEEE Intell. Transp. Syst.*, 2019, pp. 3288–3293.

- [15] W. Xu, F. Zhang, FAST-LIO: A fast, robust LiDAR-inertial odometry package by tightly-coupled iterated Kalman filter, *IEEE Robot. Autom. Lett.* 6 (2) (2021) 3317–3324.
- [16] C. Le Gentil, T. Vidal-Calleja, S. Huang, IN2LAAMA: Inertial lidar localization autocalibration and mapping, *IEEE Trans. Robot.* 37 (1) (2021) 275–290.
- [17] G. Klein, D. Murray, Parallel tracking and mapping for small AR workspaces, in: *IEEE Int. Symp. on Mixed and Augmented Reality*, 2007, pp. 225–234.
- [18] H. Strasdat, J. Montiel, A.J. Davison, Scale drift-aware large scale monocular SLAM, *Robot.: Sci. Syst.* 2 (3) (2010) 7.
- [19] H. Strasdat, A.J. Davison, J.M. Montiel, K. Konolige, Double window optimisation for constant time visual SLAM, in: *IEEE Int. Conf. Comput. Vision*, 2011, pp. 2352–2359.
- [20] H. Strasdat, J.M. Montiel, A.J. Davison, Visual SLAM: Why filter? *Image Vis. Comput.* 30 (2) (2012) 65–77.
- [21] R. Mur-Artal, J.M.M. Montiel, J.D. Tardos, ORB-SLAM: A versatile and accurate monocular SLAM system, *IEEE Trans. Robot.* 31 (5) (2015) 1147–1163.
- [22] R. Mur-Artal, J.D. Tardos, ORB-SLAM2: An open-source SLAM system for monocular, stereo, and RGB-D cameras, *IEEE Trans. Robot.* 33 (5) (2017) 1255–1262.
- [23] C. Campos, R. Elvira, J.J. Gomez, J.M.M. Montiel, J.D. Tardos, ORB-SLAM3: An accurate open-source library for visual, visual-inertial and multi-map SLAM, *IEEE Trans. Robot.* 37 (6) (2021) 1874–1890.
- [24] R.A. Newcombe, S.J. Lovegrove, A.J. Davison, DTAM: Dense tracking and mapping in real-time, in: *IEEE Int. Conf. Comput. Vision*, 2011, pp. 2320–2327.
- [25] R.A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A.J. Davison, P. Kohi, J. Shotton, S. Hodges, A. Fitzgibbon, KinectFusion: Real-time dense surface mapping and tracking, in: *IEEE Int. Symp. on Mixed and Augmented Reality*, 2011, pp. 127–136.
- [26] F. Steinbrücker, J. Sturm, D. Cremers, Real-time visual odometry from dense RGB-D images, in: *IEEE Int. Conf. Comput. Vision*, 2011, pp. 719–722.
- [27] T. Whelan, S. Leutenegger, R.F. Salas-Moreno, B. Glocker, A.J. Davison, ElasticFusion: Dense SLAM without a pose graph, in: *Robot.: Sci. Syst.*, 2015.
- [28] O. Kähler, V.A. Prisacariu, D.W. Murray, Real-time large-scale dense 3D reconstruction with loop closure, in: *Eur. Conf. Comput. Vision*, 2016, pp. 500–516.
- [29] J. Zhang, S. Singh, Visual-lidar odometry and mapping: Low-drift, robust, and fast, in: *IEEE Int. Conf. Robot. Autom.*, 2015, pp. 2174–2181.
- [30] J. Graeter, A. Wilczynski, M. Lauer, LIMO: Lidar-monocular visual odometry, in: *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2018, pp. 7872–7879.
- [31] Y. Seo, C.-C. Chou, A tight coupling of vision-lidar measurements for an effective odometry, in: *IEEE Intell. Veh. Symp.*, 2019, pp. 1118–1123.
- [32] J. Lin, F. Zhang, A fast, complete, point cloud based loop closure for LiDAR odometry and mapping, 2019, arXiv e-prints: abs/1909.11811. arXiv:1909.11811.
- [33] S. Zhao, H. Zhang, P. Wang, L. Nogueira, S. Scherer, Super odometry: IMU-centric LiDAR-visual-inertial estimator for challenging environments, in: *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2021, pp. 8729–8736.
- [34] H. Wang, C. Wang, L. Xie, Intensity scan context: Coding intensity and geometry relations for loop closure detection, in: *IEEE Int. Conf. Robot. Autom.*, 2020, pp. 2095–2101.
- [35] S. Agarwal, K. Mierle, et al., Ceres solver, <http://ceres-solver.org>.
- [36] A. Geiger, P. Lenz, R. Urtasun, Are we ready for autonomous driving? The KITTI vision benchmark suite, in: *IEEE Conf. Comput. Vision Pattern Recog.*, 2012, pp. 3354–3361.
- [37] L. Li, X. Kong, X. Zhao, W. Li, F. Wen, H. Zhang, Y. Liu, SA-LOAM: Semantic-aided LiDAR SLAM with loop closure, in: *IEEE Int. Conf. Robot. Autom.*, 2021, pp. 7627–7634.
- [38] X. Zheng, J. Zhu, Efficient LiDAR odometry for autonomous driving, 2021, arXiv e-prints: abs/2104.10879. arXiv:2104.10879.
- [39] M. Yokozuka, K. Koide, S. Oishi, A. Banno, LiTAMIN2: Ultra light LiDAR-based SLAM using geometric approximation applied with KL-divergence, 2021, arXiv e-prints: abs/2103.00784. arXiv:2103.00784. URL <https://arxiv.org/abs/2103.00784>.
- [40] F. Bonnín-Pascual, E. Garcia-Fidalgo, J.P. Company-Corcoles, A. Ortiz, MUS-SOL: A micro-uas to survey ship cargo holds, *Remote Sens.* 13 (3419) (2021).



**Emilio Garcia-Fidalgo** is a postdoctoral researcher at the Department of Mathematics and Computer Science, University of the Balearic Islands (UIB). He holds B.Sc., M.Sc. and Ph.D. degrees in Computer Science. His current research interests include localization and mapping, visual place recognition, computer vision, machine learning and unmanned aerial vehicles (UAVs).



**Joan Pep Company-Corcoles** is currently pursuing a Ph.D. degree at the University of the Balearic Islands (UIB). He received the B.Sc. and the M.Sc. degree from the Polytechnic University of Catalonia (UPC). His current research interests include Simultaneous Localization And Mapping (SLAM), place recognition, computer vision and unmanned aerial vehicles (UAVs).



**Francisco Bonnín-Pascual** is a post-doctoral researcher and teaching assistant at the Department of Mathematics and Computer Science of the University of the Balearic Islands (UIB). He holds B.Sc., M.Sc. and Ph.D. degrees in Computer Science. His current research interests include motion estimation and control architectures for aerial robots, and computer vision and machine learning techniques applied to visual inspection.



**Alberto Ortiz** is Full Professor at the Department of Mathematics and Computer Science of the University of the Balearic Islands (UIB). He holds B.Sc. and Ph.D. degrees in Computer Engineering. He is author and co-author of more than 170 publications related with computer vision, machine learning and mobile robotics. His current research interests are machine learning (deep and shallow) and its applications, motion estimation, localization and mapping, visual guidance of mobile robots, including obstacle detection and avoidance, and control architectures for mobile robots.