

Real-Time Truly-Coupled Lidar-Inertial Motion Correction and Spatiotemporal Dynamic Object Detection

Cedric Le Gentil, Raphael Falque, and Teresa Vidal-Calleja

Abstract—Over the past decade, lidars have become a cornerstone of robotics state estimation and perception thanks to their ability to provide accurate geometric information about their surroundings in the form of 3D scans. Unfortunately, most of nowadays lidars do not take snapshots of the environment but *sweep* the environment over a period of time (typically around 100 ms). Such a rolling-shutter-like mechanism introduces motion distortion into the collected lidar scan, thus hindering downstream perception applications. In this paper, we present a novel method for motion distortion correction of lidar data by tightly coupling lidar with Inertial Measurement Unit (IMU) data. The motivation of this work is a map-free dynamic object detection based on lidar. The proposed lidar data undistortion method relies on continuous preintegrated of IMU measurements that allow parameterising the sensors’ continuous 6-DoF trajectory using solely eleven discrete state variables (biases, initial velocity, and gravity direction). The undistortion consists of feature-based distance minimisation of point-to-line and point-to-plane residuals in a non-linear least-square formulation. Given undistorted geometric data over a short temporal window, the proposed pipeline computes the spatiotemporal normal vector of each of the lidar points. The temporal component of the normals is a proxy for the corresponding point’s velocity, therefore allowing for learning-free dynamic object classification without the need for registration in a global reference frame. We demonstrate the soundness of the proposed method and its different components using public datasets and compare them with state-of-the-art lidar-inertial state estimation and dynamic object detection algorithms.

I. INTRODUCTION

A key condition for popularising autonomous systems in society is their ability to robustly perceive their environment and safely interact with humans and other moving agents in their surroundings. Recent advancements in computer vision have provided the robotics community with unprecedented perception abilities. Many of these methods are based on visual images and deep neural networks. However, standard cameras have limitations (e.g., low illumination scenes), not all robotic systems can include a camera (e.g., privacy), and the computational burden of some deep neural networks can be prohibitive for embedded systems. Additionally, these methods sometimes fail to generalise to unseen scenes [1], and for some critical industries, they lack explainability [2]. As illustrated in Fig. 1, this paper focuses on two components of lidar-inertial perception: an ego-motion estimation approach and a dynamic object detection algorithm. This

All the authors are with the Robotics Institute at the University of Technology Sydney, Australia. {cedric.legentil; raphael.falque; teresa.vidalcalleja}@uts.edu.au

Cedric and Teresa are supported by the Australian Research Council Discovery Project under Grant DP210101336.

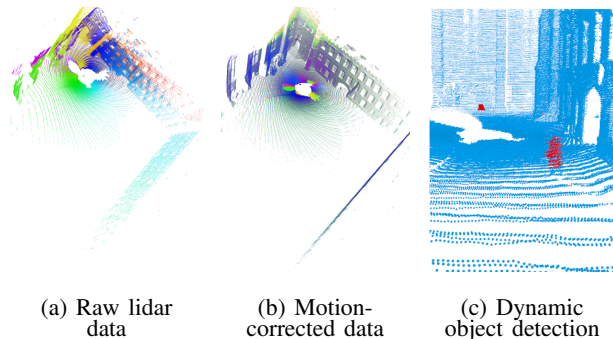


Fig. 1. The proposed method introduces a real-time IMU-lidar motion distortion algorithm that undistorts lidar data (b) to later perform dynamic object detection (c). The colour in (a) and (b) depicts the points’ timestamps, while in (c) corresponds the classification output in (c).

work is based on traditional geometric models and does not require any training phase.

Due to the sweeping nature of most of nowadays lidars, the scans collected from a moving sensor suite are subject to motion distortion. Throughout the years, various techniques have been proposed to address this distortion and perform state estimation. A common approach is to combine lidar and inertial data from an Inertial Measurement Unit (IMU) into a single algorithm [3]. The idea is that the inertial information collected during a lidar scan can be used to correct the motion that happened during the scan. Most of the popular frameworks available for lidar-inertial state estimation rely on an “open-loop” motion distortion correction of the lidar scans in the sense that it is not part of their estimation algorithm, but a preprocessing step that often relies on the current knowledge of the sensor’s pose and velocity [4]. Accordingly, such a method needs roughly known initial conditions to perform optimally. Thus, the claim for “tightly-coupled” lidar-inertial state estimation is only partially true [5]. Our method proposes a truly coupled technique that does not require any initialisation or specific initial conditions. This initialisation-free undistortion is motivated by the task of dynamic object detection without the need for global registration of the lidar data, opening future opportunities for robust lidar-inertial odometry and mapping framework in dynamic environments.

The main idea of the proposed motion distortion correction method is to use our previous work on continuous preintegrated measurements from [6] between timestamps τ_0 and τ_1 as the 6-Degree-of-Freedom (DoF) continuous trajectory representation over the same temporal window. Doing so,

based on the assumption of constant bias between τ_0 and τ_1 , the trajectory is fully defined by the value of the IMU biases (6 values), the velocity at time τ_0 (3 values), and the orientation of the gravity vector at time τ_0 (2 values).

The ability to robustly undistort lidar data is a key condition for model-based scene understanding. In this work, we integrate the proposed motion correction method with our previous work on dynamic object detection [7]. In this previous work, we relied on known poses of the sensor in a global reference frame to extract the spatiotemporal normal vector of the points in the scene. The known sensor poses were provided by the known kinematics of a robotic arm or the output of a visual-inertial framework. In this paper, we show that it is possible to reliably classify the dynamic objects in the scene using only lidar and inertial data¹ by performing motion correction over a short period of time and analysing the spatiotemporal normal vectors without the need for global registration of the lidar data. It makes our method robust to any estimation drift, which contrasts with methods like [8] that require known poses and maintain a global map of the environment.

To summarise, the contributions of this work are as follows: the derivation of a real-time algorithm for initialisation-free motion distortion correction based on lidar-inertial data; its integration with an explainable model-based dynamic object detection; the evaluation of the proposed framework components with real world-data; and the open-source release of the code and dataset with ground truth annotation used for validation².

II. RELATED WORK

A. Lidar inertial state estimation

Provided the ubiquity of IMUs and the popularisation of 3D lidars, a lot of work has focused on the problem of lidar-inertial odometry and state estimation. More than a decade ago, the authors of *Zebedee* [9] introduced a lidar-inertial 3D mapping device that relied on a piece-wise linear motion model as a continuous representation of the sensor trajectory to address the challenging motion of the 2D lidar used. Later, *LOAM* [10] was presented. It introduced a certain amount of concepts that are still the cornerstone of many lidar-inertial frameworks proposed nowadays (including ours). Its formulation relies on scan-to-scan and scan-to-map lidar registration using lidar features corresponding to planar patches and edges in the environment. Despite its ability to use inertial data, *LOAM* algorithm uses lidar and IMU data in a fairly loose manner in the sense that it was mostly used as an open-loop undistortion step and an educated initial guess for the scan-to-scan matching.

Some works attempted to couple the IMU and lidar data by using a factor graph formulation [4], [5] or in an EKF [11]. Recently, *DLIO* [12] introduced a continuous-time undistortion mechanism based on a constant-jerk motion model to interpolate the inertial predictions between IMU

measurements. Despite the claim of “tight” coupling, these methods use the IMU to correct the motion distortion in a one-off state-propagation pre-processing step based on the last state estimate. It means that the position of the lidar during a lidar scan is not estimated using both geometric and inertial information. Therefore, the undistortion accuracy is tied to the accuracy of the previous scan-matching steps. On another hand, methods with a continuous state representation [9], [13] perform “non-rigid” registration of the lidar scans by considering the points independently. A downside of this approach is the size of the estimated state that inherently yields a heavier computation burden in general, as shown in [14].

The concept of preintegration, originally presented in [15], corresponds to the creation of pseudo measurements that combine the information from multiple IMU measurements without the need to know the initial conditions. In the context of lidar-IMU calibration, the work in [16] relied on a continuous representation of the inertial data to virtually upsample the IMU measurement and perform preintegration, thus providing inertial information for each lidar point. This has been later extended with *IN2LAAMA* [17], an offline localisation and mapping framework. With novel continuous preintegrated measurements [18] then [6], *IN2LAAMA* is creating a bridge between discrete and continuous-time state representations with discrete pose, velocity and bias variables for each lidar scan but locally considering continuous trajectories within each of the scans, thus allowing for “non-rigid” scan-to-scan registration. In this paper, we push this notion further by eliminating the concept of scan and removing the pose from the estimated state.

B. Dynamic object detection

As mentioned earlier, a common approach to detect dynamic objects is to leverage semantic segmentation of geometric sensor data. The classic strategy is to use a UNET architecture to classify all the pixels/3D points related to potential dynamic elements. In [19], the authors proposed to use such a method with the Convolutional Neural Network (CNN) Rangenet++ [20] and the KITTI Vision Benchmark dataset [21]. This work was later extended by changing the network input space for a stream of consecutive scans [22]. Alternatively, if depth and color (RGB) images are acquired simultaneously, the RGB data can be used with a standard semantic segmentation network, such as a Region-based Convolutional Neural Network (R-CNN), and then mapped back into the depth information for classifying potential dynamic objects [23].

As mentioned above, these algorithms can classify *potential* dynamic objects. However, they can not separate moving elements from static ones that have the potential to be dynamic (e.g., a parked and a moving vehicle). For this reason, some of the literature focuses on detecting dynamic objects regardless of the semantic information. A straightforward approach is to integrate dynamic object detection into an occupancy grid mapping framework. The occupancy grid can then be used to detect areas that used to

¹Most of the lidars available today include an IMU.

²https://uts-ri.github.io/lidar_inertial_motion_correction/

be *free* and are now *occupied* (i.e., all the newly occupied areas are parts of dynamic objects). This strategy has been proposed in Dynablox [8], where the authors build upon the voxel occupancy map from the Voxelblox framework [24] to solve the dynamic object detection. Similarly, Mersch et al. proposed a framework that segments dynamic objects by looking at the discrepancies between a lidar scan and the map. They propose to use a sparse 4D (i.e., [x, y, z, time]) CNN to segment points that are dynamic [25].

In recent work, we introduced a simple method to detect dynamic elements within range scan data as a byproduct of the spatiotemporal normals computation [7]. While not requiring to solve the mapping problem as in [8], [25], our previous work relied on the accurate knowledge of the sensor poses in a global reference frame. By, integrating [7] over the short temporal windows of the proposed lidar-inertial undistortion method, we demonstrate the ability to recover object dynamicity based on spatiotemporal normal computation without the need for global registration.

III. METHOD

A. Preintegration background

The preintegration concept allows for the aggregation of IMU measurements into pseudo measurements that relate the inertial pose at different timestamps. In the context of motion distortion correction we are only interested in the relative motion between timestamps τ_0 and τ_1 considering the reference frame of the IMU at time τ_0 as the “local” reference frame. We denote it as $\mathfrak{F}_{\tau_0}^I$. Accordingly, the pose of the IMU at τ_0 is the identity. In such conditions, the standard preintegration equations that involve the preintegrated measurements $\Delta \mathbf{R}_{\tau_0}^t$ (rotation) and $\Delta \mathbf{p}_{\tau_0}^t$ (position) can be simplified as

$$\mathbf{R}_{I_0}^t = \Delta \mathbf{R}_{\tau_0}^t \quad (1)$$

$$\mathbf{p}_{I_0}^t = (t - \tau_0) \mathbf{v}_{I_0}^{\tau_0} + \frac{(t - \tau_0)^2}{2} \mathbf{g}_{I_0} + \Delta \mathbf{p}_{\tau_0}^t, \quad (2)$$

where \mathbf{g}_{I_0} is the gravity vector in $\mathfrak{F}_{\tau_0}^I$. The notations $\mathbf{R}_{I_0}^t$, $\mathbf{v}_{I_0}^t$, and $\mathbf{p}_{I_0}^t$ represent the IMU orientation, velocity, and position in $\mathfrak{F}_{\tau_0}^I$, respectively.

For the rest of the paper, we will use the homogenous transformation $\mathbf{T}_{I_0}^t$ to refer to the IMU pose at time t

$$\mathbf{T}_{I_0}^t = \begin{bmatrix} \mathbf{R}_{I_0}^t & \mathbf{p}_{I_0}^t \\ \mathbf{0} & 1 \end{bmatrix}. \quad (3)$$

Note that the preintegrated measurements in (1) and (2) are considered to be bias-free. However, in real-world scenarios, the inertial measurements are affected with slow-varying additive biases. Following the seminal work in [15], most preintegration methods offer a Taylor expansion of the preintegrated measurements to enable post-integration bias correction. Thus the IMU pose (3) is not only a function of the preintegrated measurement, the initial velocity, and gravity vector, but it is also a function of the accelerometer and gyroscope biases. For more details about IMU preintegration and bias correction mechanisms please refer to [6].

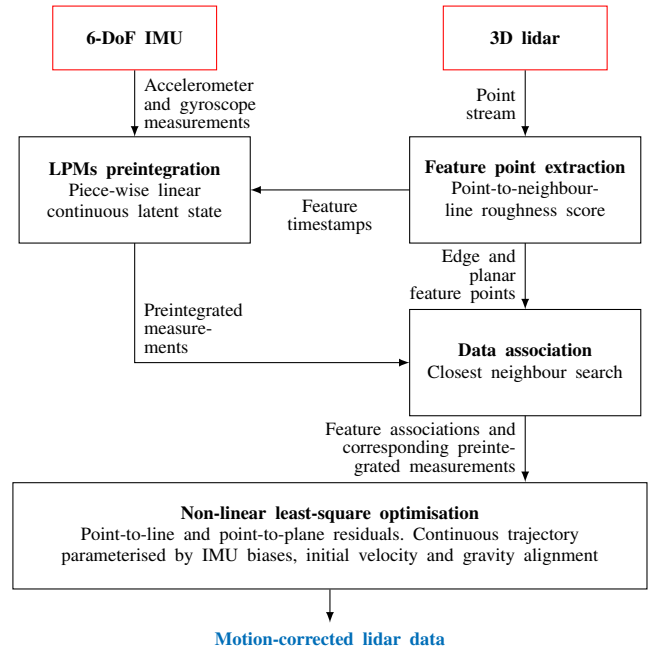


Fig. 2. Overview of the proposed lidar-inertial undistortion pipeline. The sensors’ trajectory is modelled with continuous preintegrated measurements [6] resulting in an estimated state of only 11-DoF.

B. Motion correction

To detect moving objects in the lidar data, we first need to correct for the motion distortion due to the ego-motion of the sensor. Our method relies on both a lidar and a 6-DoF IMU to provide distortion-free point clouds of the sensor’s surroundings based on a non-linear least-square minimisation of distances between lidar points. An overview of the proposed approach is present in Fig. 2. We consider data collected between timestamps τ_0 and τ_1 by lidar and IMU rigidly mounted with \mathbf{T}_I^L their extrinsic calibration. Please note that the proposed method does not rely on the traditional notion of “lidar scan” as the points are considered independently. However, the time window $[\tau_0, \tau_1]$ needs to be long enough to allow for data association as mentioned later in this section.

Given the raw data with timestamps, feature points are extracted from the lidar data. These correspond to planar and edge points. Then, preintegrated measurements are computed for each lidar feature following the Linear Preintegrated Measurements (LPMs) introduced in [6]. Using (3), it is possible to project any lidar point into the sensor’s referential frame at time τ_0 where data association between features can be performed. Each data association lead to a point-to-plane or point-to-line residual (depending on the nature of the features) in our non-linear least-square optimisation. Using the Levenberg-Marquardt algorithm we estimate the state $\mathcal{S} = [\mathbf{b}_f \ \mathbf{b}_\omega \ \mathbf{v}_{I_0}^{\tau_0} \ \mathbf{g}_{I_0}]$, with \mathbf{b}_f and \mathbf{b}_ω the constant accelerometer and gyroscope biases, \mathbf{v}_{τ_0} the IMU velocity at time τ_0 , \mathbf{g}_I the gravity vector in the IMU reference frame at time τ_0 .

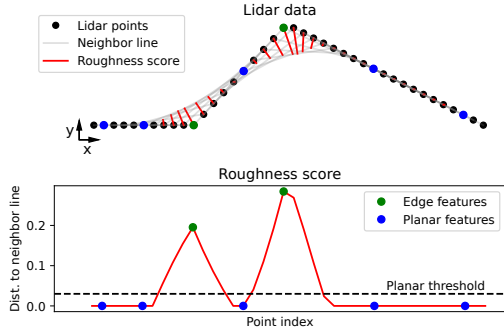


Fig. 3. Illustration of the proposed lidar feature detection as the point-to-line distance with the “neighbour line”.

1) *Feature detection*: Let us denote \mathbf{x}_L^i a lidar point collected at time t_i with $\tau_0 \leq t_i \leq \tau_1$. Most of the lidars available nowadays collect data by swiping the environment with a fixed number of laser beams. Let us use the term “channel” to refer to the successive points collected by one of the lasers.³ The proposed feature detection computes a roughness score for each point in a channel as⁴

$$r(\mathbf{x}_L^i) = \frac{\|(\mathbf{x}_L^i - \mathbf{x}_L^{i-n}) \times (\mathbf{x}_L^i - \mathbf{x}_L^{i+n})\|}{\|\mathbf{x}_L^{i-n} - \mathbf{x}_L^{i+n}\|}, \quad (4)$$

with n a user-defined parameter. This score corresponds to the point-to-line distance between the point under consideration and the line defined by neighbouring points from the same channel. Points with a score under a certain threshold are considered as planar features. To limit the computational burden in the rest of the proposed method, only a random subset of these planar features is used if too many points are planar. The edge features correspond to local maxima in the roughness score that is above the planar threshold as illustrated in Fig. III-B.1. Let us use the generic notation \mathbf{f}_L^i to refer to selected feature points regardless of their nature.

2) *Data association*: As later illustrated in Fig. 5, we divide the time window into K “segments” of the same duration. Each segment requires a spatial overlap with the other segments. Accordingly, the time difference between τ_1 and τ_0 must be superior to K times the duration of the scanning pattern of the lidar used (e.g., with $K = 2$, $\tau_1 - \tau_0 > 200$ ms for a standard spinning lidar operating at 10 Hz). The data association between two successive segments consists first of the following steps:

- Projection of all the features to the IMU frame at time τ_0 leveraging (3): $\begin{bmatrix} \mathbf{f}_{I_0}^i \\ 1 \end{bmatrix} = \mathbf{T}_{I_0}^{t_i}(\mathcal{S}) \mathbf{T}_I^L \begin{bmatrix} \mathbf{f}_L^i \\ 1 \end{bmatrix}$.
- Building KDTrees [26] with the features of each segment (independently for each feature type).
- Closest neighbour search between segment pairs (2 neighbours for edges, and 3 for planar features).

We consider an association of a feature with its neighbours as valid if the Euclidean distance separating them is under a user-defined threshold. We denote $\mathbf{a}_i = \{\mathbf{f}_L^i, \mathbf{f}_L^j, \mathbf{f}_L^k, \mathbf{f}_L^l\}$ one

³For example, there are 16 channels in a Velodyne VLP16, while a Blickfeld Cube1 or a Livox Avia only have 1 channel.

⁴For the sake of notation simplicity, the point indices in (4) reflect the case with only 1 channel.

of such associations with \mathbf{f}_L^i the feature under consideration and $\mathbf{f}_L^j, \mathbf{f}_L^k, \mathbf{f}_L^l$ the target features.

3) *Cost function*: Provided the aforementioned data associations, our estimation problem corresponds to the non-linear least-square optimisation of the cost function as

$$\mathcal{S}^* = \underset{\mathcal{S}}{\operatorname{argmin}} \sum_{\mathbf{a}_i \in \mathcal{A}} (d(\mathbf{a}_i, \mathcal{S}))^2, \quad (5)$$

where $d(\mathbf{a}_i, \mathcal{S})$ is the distance function specific to the type of features. For edges, the distance function is

$$d(\mathbf{a}_i, \mathcal{S}) = \frac{\|(\mathbf{f}_{I_0}^i - \mathbf{f}_{I_0}^j) \times (\mathbf{f}_{I_0}^i - \mathbf{f}_{I_0}^k)\|}{\|\mathbf{f}_{I_0}^j - \mathbf{f}_{I_0}^k\|}, \quad (6)$$

and for planar points

$$d(\mathbf{a}_i, \mathcal{S}) = \frac{(\mathbf{f}_{I_0}^i - \mathbf{f}_{I_0}^j)^\top ((\mathbf{f}_{I_0}^j - \mathbf{f}_{I_0}^k) \times (\mathbf{f}_{I_0}^j - \mathbf{f}_{I_0}^l))}{\|(\mathbf{f}_{I_0}^j - \mathbf{f}_{I_0}^k) \times (\mathbf{f}_{I_0}^j - \mathbf{f}_{I_0}^l)\|}. \quad (7)$$

Note that the state appears on the right-hand side through the transformation of the lidar points into the first IMU frame.

C. Dynamic object detection

This subsection provides the technical details of the integration of our previous work [7] with the proposed undistortion method.

1) *Overview*: Let us define a dynamic object as an element in the environment that has a non-null velocity with respect to an earth-fixed reference frame. Thus, the definition of dynamic object here is agnostic to the type or nature of the objects present in the environment. The goal is to classify the individual points in the distortion-free lidar data based on a per-point *dynamycity score*. As illustrated in Fig. 4, the normal vector to the spatiotemporal surface contains information about the velocity of an object. The temporal component of these spatiotemporal normals is used as our score. In the 1D scenario, it directly links to the object velocity while it corresponds to the projection of the velocity on the spatial normal vector in higher dimensions.

Concretely, the dynamycity score computation for each point consists of the following steps:

- Spatial neighbour search.
- Computation of the local spatiotemporal covariance matrix.
- Normal computation (eigendecomposition, eigenvector with the smallest eigenvalue).
- Score as the temporal component of the normal, and thresholding.

2) *Dynamycity score via spatiotemporal normal*: Let us consider the undistorted lidar points between τ_0 and τ_1 as \mathcal{P} . To provide a score to a point $\mathbf{x}_{I_0}^i$, we first perform a spatial neighbourhood search in \mathcal{P} . Defining \mathcal{N}_i as the set of neighbour points to $\mathbf{x}_{I_0}^i$, the local covariance cov_i is computed as

$$\operatorname{cov}_i = \frac{1}{\|\mathcal{N}_i\|} \sum_{\mathbf{x}_{I_0}^v \in \mathcal{N}_i} \left(\begin{bmatrix} \mathbf{x}_{I_0}^v \\ t_v \end{bmatrix} - \mathbf{m}_i \right) \left(\begin{bmatrix} \mathbf{x}_{I_0}^v \\ t_v \end{bmatrix} - \mathbf{m}_i \right)^\top \quad (8)$$

$$\text{with } \mathbf{m}_i = \frac{1}{\|\mathcal{N}_i\|} \sum_{\mathbf{x}_{I_0}^v \in \mathcal{N}_i} \begin{bmatrix} \mathbf{x}_{I_0}^v \\ t_v \end{bmatrix}.$$

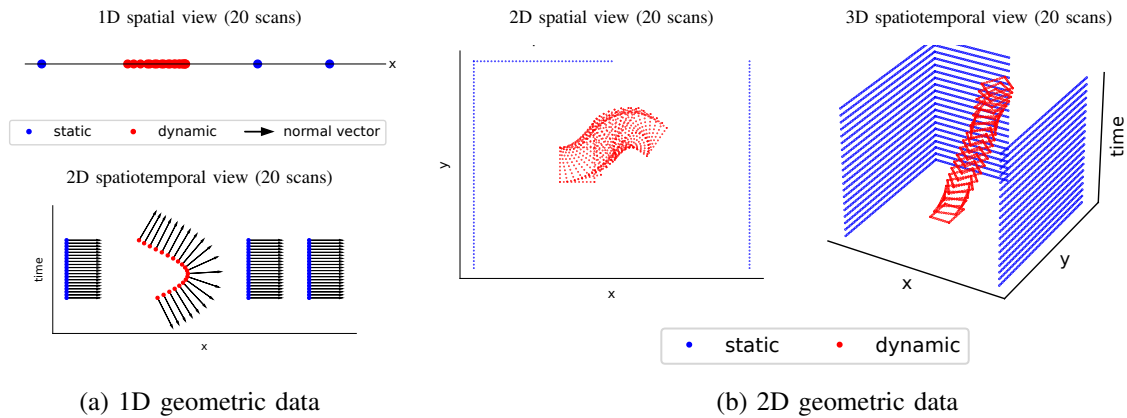


Fig. 4. Intuition for the spatiotemporal normal computation as a proxy for dynamic point detection, 1D in (a) and 2D in (b). The spatiotemporal normal vectors are related to the velocity of an object; for static objects, the temporal component of the normals is null (the normals have been omitted in (b) for the sake of readability) [7].

We define the dynamicity score s_i^j as the absolute value of the temporal component (last component) of the Eigenvector associated with the smallest Eigenvalue of cov_i . This value is related to the projection of the object velocity in the 3D space along the spatial normal vector at that location. The point classification as static or dynamic simply consists of thresholding s_i^j . Please note that this concept is not specific to lidar data and can therefore be applied to data from depth-cameras too as shown in the experiment section of this paper as previously shown in [7].

IV. EXPERIMENT

A. Implementation

Our open-source C++ ROS (1 and 2) implementation of the undistortion has been made available.⁵ We rely on *Ceres* [28] for the non-linear least-square problem in (5). The length of the time window for the motion distortion correction is 450 ms with segments of 150 ms. The undistortion of all points is performed in a sliding window manner with increments of 150 ms, as shown in Fig. 5. The estimated velocity and gravity orientation of a given window are used as initial guess for the following window. To lighten the computational load, the LPMs are computed at a fixed frequency of 1 kHz and interpolation is performed for the feature timestamps in between.

The dynamic object detection is implemented by integrating the framework for spatiotemporal normals computation [7] consisting of four steps that are, point cloud accumulation, downsampling, dynamic score computation, and upsampling. Note that here, there is no need for point cloud accumulation as the undistortion step occurs over several scans, thus alleviating the need for global pose. The parameters have been tuned to work with the density of the pointcloud from the dataset. The downsampling and upsampling modules are based on voxels of size 0.1 m

⁵https://uts-ri.github.io/lidar_inertial_motion_correction/. Please note that the released code also integrates an implementation of [27] for scan-to-map-style odometry and mapping in dynamic environments.

and 0.2 m, respectively. The threshold for classification of dynamic objects is set to 0.4 and the search radius for the spatiotemporal normals is set to 0.3m. Following [7]’s notations, $N = 1$, $d_r = 0.3$, $thr = 0.4$.

B. Motion correction

It is important to note that the proposed motion correction method does not aim at performing global odometry but instead targets local undistortion of the lidar data. Accordingly, standard benchmarking methods for odometry frameworks, like pose trajectory accuracy, are not considered in this section. Instead, this evaluation focuses on the consistency of the lidar data after correction against the 3D map of the environment. We use the *Newer College Dataset* [29] as it includes both inertial and geometric data from an *Ouster OS0* (128 beams) lidar and its embedded IMU, as well as the ground-truth map of the area collected with a survey-grade *Leica BLK360* lidar. To benchmark our motion correction method we choose to compare with state-of-the-art real-time lidar-inertial estimation frameworks that are *Fast-LIO2* [11] and *DLIO* [12] that assert tight coupling of lidar and inertial data in their state estimation pipelines.

1) *Accuracy*: This setup consists of saving to disk all the motion-corrected scans and their associated timestamps before registering them into the ground-truth map using the ground-truth poses and a final ICP alignment. Once the scans are aligned, we query the smallest distance between each lidar point in the scans and the map. As the scans and the map do not have a 100% overlap (some areas seen by the mobile lidar are not observed from any of the survey lidar poses), we compare the mean distances of the 75% best matches. Fig. 6 shows this mean error through time for the three methods on both the *quad medium* and *quad hard* sequences.

Note that these sequences allow for a fair comparison with both *Fast-LIO2* and *DLIO* because the data start with very little motion allowing for close-to-optimal initialisation of the state. One can see that *Fast-LIO2* under-performs compared to our method and *DLIO*. The fact that *DLIO*




Fig. 5. Illustration of the proposed pipeline’s timeline in terms of lidar data used for motion correction and dynamic object detection. As per our implementation, 3 segments (of 150 ms) per temporal window are used with a sliding window increment of 1 segment. For the dynamic detection n , the 3 segments of the motion correction n are used. In each motion correction, the lidar features from the first and last segments in the temporal window are used for data association.

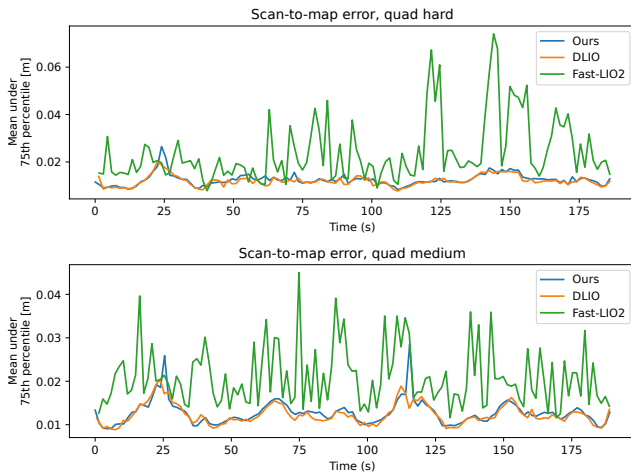


Fig. 6. Comparison of the motion distortion correction accuracy between the proposed method and state-of-the-art lidar-inertial odometry frameworks. Distance between ground-truth map and motion-corrected lidar data after ICP registration.

and the proposed method have similar performance⁶ with a mean error close to both the value of the sensor’s noise (≈ 0.02 m) and the map resolution (0.01 m voxel downsampled) confirms that both methods can recover the motion distortion of the lidar accurately.

2) *Robustness to initial conditions*: In the previous experiments, we have provided the optimal initialisation conditions for *Fast-LIO2* and *DLIO*. Here we explore the scenario where the sensor suite is moving when switched on. In other words, in terms of state estimation, it means that the initial conditions of the lidar dynamics are not accurately known. We compare both our approach and *DLIO* in such scenarios by starting the *Quad hard* sequence at timestamps 120 s and 141 s (these correspond to higher rotational motion).

Fig. 7 shows the “double-wall” problem that occurs when the undistortion is not optimal. In Fig. 8, we report the “gap” within the double walls by computing point-to-plane distances in areas where lidar data self-overlaps between

⁶The occasional spikes in *Ours* are likely due to a sub-optimal result of the ICP refinement.

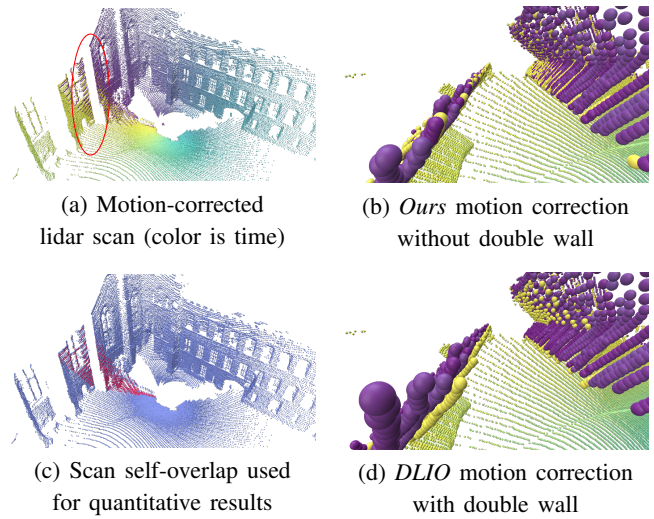


Fig. 7. Illustration of the undistorting abilities of the proposed method where no “double-wall” is visible in the motion-corrected data. (a) shows the scene’s scan, (c) highlights the self-overlap area of the scan, (b) and (d) are closeups of the motion correction results of the proposed method and *DLIO* [12] within the red ellipse in (a).




Fig. 8. Comparison of the motion distortion correction accuracy between the proposed method and state-of-the-art lidar-inertial odometry frameworks.

the beginning and the end of the sweep (cf. Fig. 7(c)). As expected, our method accurately corrects the motion distortion while *DLIO* cannot close the double-wall gap as its undistortion mechanism strongly relies on the knowledge of the initial conditions. These results demonstrate how truly-coupled methods can estimate ego-motion with higher accuracy compared to frameworks based on “one-off” undistortion of the lidar data.

3) *Discussion*: The level of accuracy of *DLIO*’s undistortion has been proven to be sufficient for many lidar-inertial odometry and mapping applications where generally the system is switched on in a relatively static position. However, we showed that its “on-off” undistortion mechanism reaches its limits in more challenging scenarios where the issue of “double walls” arises. In the context of dynamic object detection, this can lead to misdetections: a gap of 0.1 m between surfaces observed approximately 100 ms apart corresponds to an apparent velocity of 1 m/s (slow walking pace). We demonstrated that our approach allows for very accurate motion-distortion correction, thus enabling downstream applications such as dynamic object detection (cf., Section IV-C).

All our experiments run on a consumer-grade laptop with an *Intel i7-1370P* CPU and 32 GB of RAM. All methods display a very low CPU usage with around 3% for *Fast-LIO2*, 6% for *DLIO*, and 7% for ours. It is interesting to note that the sliding-window nature of our implementation is motivated by the task of dynamic object detection that requires significant spatial overlap in the undistorted lidar data for more robust detections but still displays high level of efficiency. For pure odometry, one can imagine even more efficient computations by removing redundant estimations.

C. Dynamic object detection

In this subsection, we evaluate the integration of both our motion distortion correction and dynamic object detection through a comparison with *Dynablox* [8], a state-of-the-art non-learning dynamic object detection framework. This method relies on maintaining a voxel-based occupancy map of the environment, detecting the dynamic points, and clustering them into objects.

To the best of our knowledge, there is no public dataset that provides raw, synchronised lidar and inertial data as well as point-wise labels for point dynamicity. With this in mind, we chose to add the dynamicity information to the *Newer College dataset* (*quad_easy*, *medium*, and *hard*) and we collected a novel dataset referred to as the *Techlab dataset*. The latter contained two labelled sequences that have been collected in our Techlab facility at the University of Technology Sydney with a 16-beam OS1 Ouster lidar and its embedded IMU. A Vicon motion capture system has been used to collect the pose of the lidar relative to an earth-fixed frame. For both datasets, the generation of the ground-truth dynamicity of each point follows a similar two-step process: the first one is mapping the environment without dynamic objects, the second one is the registration of each of the collected lidar points and query the distance between each point and the static map. A point is classified as dynamic if the queried distance is above a certain threshold.⁷ For the *Newer college dataset*, we have manually cleaned the provided map to remove dynamic objects and used the provided groundtruth trajectory for point registration. For the *Techlab dataset* the map has been create with the help of the Vicon system and the lidar data collected in an empty lab. In the first *Techlab* sequence, three people and the sensor carrier are constantly walking around the room, while in the second sequence, people momentarily stop before starting to move again. Note that with the proposed ground-truth generation process, a person that stops will be considered as dynamic despite have a velocity equal to zero.

The evaluation shown in Table I compares the performance of the proposed method and *Dynablox* to which we provided the ground truth localisation as *Dynablox* does not address the problem of scan registration/motion distortion correction. The different metrics are computed for points less than 20m away from the lidar. Across the different metrics,

⁷Some frames are not labelled due to lack of groundtruth pose. Other points are ignored, as not all the areas observed by the lidar are present in the map.

TABLE I
EVALUATION OF THE INTERSECTION OVER UNION (IOU), RECALL, ACCURACY, PRECISION, AND F1 SCORE BETWEEN DYNABLOX [8] AND OUR PROPOSED METHOD ON THE *Techlab* AND THE *Newer College* DATASETS.

<i>Techlab</i>		IoU	Recall	Accuracy	Precision	F-1
<i>Sequence 1</i>	<i>Ours</i>	0.61	0.72	0.99	0.81	0.76
	[8]	0.74	0.74	0.99	0.99	0.85
<i>Sequence 2</i>	<i>Ours</i>	0.45	0.48	0.98	0.85	0.61
	[8]	0.67	0.67	0.99	0.99	0.80

<i>Newer College</i>		IoU	Recall	Accuracy	Precision	F1
<i>quad_easy</i>	<i>Ours</i>	0.80	0.94	0.97	0.85	0.89
	[8]	0.93	0.94	0.99	0.99	0.96
<i>quad_medium</i>	<i>Ours</i>	0.79	0.89	0.95	0.87	0.88
	[8]	0.94	0.95	0.99	0.99	0.97
<i>quad_hard</i>	<i>Ours</i>	0.57	0.76	0.94	0.70	0.73
	[8]	0.90	0.92	0.99	0.98	0.95

For all metrics, higher is better.




Fig. 9. Illustration of the dynamic object detection. Both methods have similar performances.

Dynablox shows better results than our proposed pipeline. However, *Dynablox* implicitly leverages information from all the previous lidar scans while ours only consider the last set of undistorted lidar data (only three scans). Additionally, it is important to note that our method and *Dynablox* have a slightly different definition for dynamic objects. Ours looks at the instant velocity of the points while *Dynablox* is detecting points that are in voxels that have previously been classified as free. Thus, *Dynablox* still considers a person that stopped moving as dynamic in the same way as the ground-truth labels are generated. This explains why the difference between our method and *Dynablox* is smaller on *Techlab 1* and larger on the other data sequences.

According to the different metrics and the qualitative example in Fig. 9, one can see that our method performs well at detecting dynamic points but tends to classify some static points as dynamic (e.g., ground points near the pedestrian’s feet). The degradation in the performance of our method in terms of IoU, recall, and precision for *quad hard* is partly explained by the imperfect ground truth as aforementioned and the fact that the dataset does not contain many dynamic points. Nonetheless, the high accuracy allows for robust removal of dynamic elements from the lidar data. Additionally, as the proposed method only relies on the current set of undistorted lidar data, our approach is not affected

by any trajectory drift while mapping-based frameworks like *Dynablox* are. Thus, future work includes the use of our dynamic-free motion-corrected lidar data in a robust odometry and mapping framework.

V. CONCLUSION

In this paper, we presented a method for lidar-inertial motion distortion correction with application to dynamic object detection. The first part consists of locally estimating the ego-motion of the sensors. It is performed by coupling each lidar point with inertial information via the use of continuous IMU preintegration. With such a formulation, the trajectory of the sensor suite is fully characterised using only 11 variables (IMU biases, initial velocity, and orientation of gravity vector). The estimated state is the result of a non-linear least-square optimisation that aims at minimising point-to-plane and point-to-line distances.

The availability of motion-corrected lidar data allows us to perform dynamic object detection based on the spatiotemporal normals at each point. Concretely, the lidar points are stored all together with their individual timestamps. Similarly to standard geometric normal computation, the method queries the spatial neighbour of each point, but instead of computing the spatial covariance (3×3) of each neighbourhood, our method computes the spatiotemporal covariance (4×4). By performing the eigendecomposition to find the normal vector (eigenvector with the smallest eigenvalue), we use its temporal component as a proxy for the points' velocity in space. Our real-time implementation has been tested in real data and made open-source.

Future work includes the integration of the proposed method in a global odometry/SLAM framework to ensure geometric consistency at large timescales. We will also explore more coupled interaction between state estimation and dynamic object detection to provide a principled approach for outlier rejection in ego-motion estimation.

REFERENCES

- [1] N. Sünderhauf, O. Brock, W. Scheirer, R. Hadsell, D. Fox, J. Leitner, B. Upcroft, P. Abbeel, W. Burgard, M. Milford, and P. Corke, "The limits and potentials of deep learning for robotics," *The International Journal of Robotics Research*, vol. 37, no. 4-5, pp. 405–420, 2018.
- [2] O. Willers, S. Sudholt, S. Raafatnia, and S. Abrecht, "Safety concerns and mitigation approaches regarding the use of deep learning in safety-critical perception tasks," in *Computer Safety, Reliability, and Security. SAFECOMP 2020 Workshops*, A. Casimiro, F. Ortmeier, E. Schoitsch, F. Bitsch, and P. Ferreira, Eds. Cham: Springer International Publishing, 2020, pp. 336–350.
- [3] D. Lee, M. Jung, W. Yang, and A. Kim, "Lidar odometry survey: recent advancements and remaining challenges," *Intelligent Service Robotics*, pp. 1–24, 2024.
- [4] H. Ye, Y. Chen, and M. Liu, "Tightly coupled 3D Lidar inertial odometry and mapping," *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2019-May, pp. 3144–3150, 2019.
- [5] 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 *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 5135–5142.
- [6] C. Le Gentil and T. Vidal-Calleja, "Continuous latent state preintegration for inertial-aided systems," *The International Journal of Robotics Research*, vol. 42, no. 10, pp. 874–900, 2023.
- [7] R. Falque, C. Le Gentil, and F. Sukkar, "Dynamic object detection in range data using spatiotemporal normals," in *Australasian Conference on Robotics and Automation*, 2023.
- [8] L. Schmid, O. Andersson, A. Sulser, P. Pfreundschuh, and R. Siegwart, "Dynablox: Real-time detection of diverse dynamic objects in complex environments," *IEEE Robotics and Automation Letters*, vol. 8, no. 10, pp. 6259–6266, 2023.
- [9] M. Bosse, R. Zlot, and P. Flick, "Zebedee : Design of a spring-mounted 3-D range sensor with application to mobile mapping," *IEEE Transactions on Robotics*, vol. 28, no. October, pp. 1–15, 2012.
- [10] J. Zhang and S. Singh, "Visual-lidar Odometry and Mapping: Low-drift, Robust, and Fast," *IEEE International Conference on Robotics and Automation*, pp. 2174–2181, 2015.
- [11] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang, "Fast-lid2: Fast direct lidar-inertial odometry," *IEEE Transactions on Robotics*, vol. 38, no. 4, pp. 2053–2073, 2022.
- [12] K. Chen, R. Nemiroff, and B. T. Lopez, "Direct lidar-inertial odometry: Lightweight lio with continuous-time motion correction," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 3983–3989.
- [13] T. Y. Tang, D. J. Yoon, and T. D. Barfoot, "A white-noise-on-jerk motion prior for continuous-time trajectory estimation on $\mathbb{SE}(3)$," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 594–601, 2019.
- [14] G. Cioffi, T. Cieslewski, and D. Scaramuzza, "Continuous-time vs. discrete-time vision-based slam: A comparative study," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2399–2406, 2022.
- [15] T. Lupton and S. Sukkariéh, "Visual-inertial-aided navigation for high-dynamic motion in built environments without initial conditions," *IEEE Transactions on Robotics*, vol. 28, no. 1, pp. 61–76, 2012.
- [16] C. Le Gentil, T. Vidal-Calleja, and S. Huang, "3D Lidar-IMU Calibration based on Upsampled Preintegrated Measurements for Motion Distortion Correction," *IEEE International Conference on Robotics and Automation*, 2018.
- [17] —, "IN2LAAMA: INertial Lidar Localisation Autocalibration And Mapping," *IEEE Transactions on Robotics*, 2021.
- [18] C. Le Gentil and T. Vidal-Calleja, "Continuous Integration over $SO(3)$ for IMU Preintegration," *Robotics: Science and Systems XVII*, 2021.
- [19] X. Chen, A. Milioto, E. Palazzolo, P. Giguere, J. Behley, and C. Stachniss, "Suma++: Efficient lidar-based semantic slam," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 4530–4537.
- [20] A. Milioto, I. Vizzo, J. Behley, and C. Stachniss, "Rangenet++: Fast and accurate lidar semantic segmentation," in *2019 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2019, pp. 4213–4220.
- [21] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *2012 IEEE conference on computer vision and pattern recognition*. IEEE, 2012, pp. 3354–3361.
- [22] X. Chen, S. Li, B. Mersch, L. Wiesmann, J. Gall, J. Behley, and C. Stachniss, "Moving object segmentation in 3d lidar data: A learning-based approach exploiting sequential data," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 6529–6536, 2021.
- [23] M. Henein, J. Zhang, R. Mahony, and V. Ila, "Dynamic slam: The need for speed," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2123–2129.
- [24] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.
- [25] M. Benedikt, G. Tiziano, C. Xieyuanli, V. Ignacio, B. Jens, and S. Cyrill, "Building Volumetric Beliefs for Dynamic Environments Exploiting Map-Based Moving Object Segmentation," *IEEE Robotics and Automation Letters (RA-L)*, vol. 8, no. 8, pp. 5180–5187, 2023.
- [26] J. L. Bentley, "Multidimensional binary search trees used for associative searching," *Communications of the ACM*, vol. 18, no. 9, pp. 509–517, sep 1975.
- [27] C. Le Gentil, O.-L. Ouabi, L. Wu, C. Pradalier, and T. Vidal-Calleja, "Accurate gaussian-process-based distance fields with applications to echolocation and mapping," *IEEE Robotics and Automation Letters*, vol. 9, no. 2, pp. 1365–1372, 2024.
- [28] S. Agarwal and K. Mierle, "Ceres Solver." [Online]. Available: <http://ceres-solver.org>
- [29] L. Zhang, M. Camurri, D. Wisth, and M. Fallon, "Multi-camera lidar inertial extension to the newer college dataset," 2021.