

ALICE-LRI: A General Method for Lossless Range Image Generation for Spinning LiDAR Sensors without Calibration Metadata

Samuel Soutullo^{a,*}, Miguel Yermo^{a,b}, David L. Vilariño^b, Óscar G. Lorenzo^{a,b}, José C. Cabaleiro^{a,b}, Francisco F. Rivera^{a,b}

^aCentro Singular de Investigación en Tecnoloxías Intelixentes (CiTIUS), Rúa de Jenaro de la Fuente Domínguez, Santiago de Compostela, 15782, A Coruña, Spain

^bDepartamento de Electrónica e Computación, Universidade de Santiago de Compostela, Rúa Lope Gómez de Marzoa, Santiago de Compostela, 15782, A Coruña, Spain

Abstract

3D LiDAR sensors are essential for autonomous navigation, environmental monitoring, and precision mapping in remote sensing applications. To efficiently process the massive point clouds generated by these sensors, LiDAR data is often projected into 2D range images that organize points by their angular positions and distances. While these range image representations enable efficient processing, conventional projection methods suffer from fundamental geometric inconsistencies that cause irreversible information loss, compromising high-fidelity applications. We present ALICE-LRI (Automatic LiDAR Intrinsic Calibration Estimation for Lossless Range Images), the first general, sensor-agnostic method that achieves lossless range image generation from spinning LiDAR point clouds without requiring manufacturer metadata or calibration files. Our algorithm automatically reverse-engineers the intrinsic geometry of any spinning LiDAR sensor by inferring critical parameters including laser beam configuration, angular distributions, and per-beam calibration corrections, enabling lossless projection and complete point cloud reconstruction with zero point loss. Comprehensive evaluation across the complete KITTI and DurLAR datasets demonstrates that ALICE-LRI achieves perfect point preservation, with zero points lost across all point clouds. Geometric accuracy is maintained well within sensor precision limits, establishing geometric losslessness with real-time performance. We also present a compression case study that validates substantial downstream benefits, demonstrating significant quality improvements in practical applications. This paradigm shift from approximate to lossless LiDAR projections opens new possibilities for high-precision remote sensing applications requiring complete geometric preservation.

Keywords: LiDAR, Lossless Range image, Sensor calibration, Point cloud reconstruction, Parameter estimation

1. Introduction

3D LiDAR sensors have become a key component of modern perception systems, enabling accurate mapping [1, 2], navigation [3, 4, 5], and scene understanding in autonomous vehicles [6, 7, 8], as well as in mobile robotics [9, 10]. A widely used class of LiDARs—commonly referred to as rotating or spinning LiDAR sensors—operates by continuously rotating a set of vertically aligned laser beams. As the sensor spins, each beam sweeps out a horizontal arc, producing structured 3D scans of the surrounding environment.

To efficiently process the resulting point clouds, it is common to project them onto 2D range images: dense, grid-like structures where each pixel encodes the distance (range) to the nearest surface along a fixed direction. This projection follows a spherical pattern defined by the sensor geometry, mapping azimuthal angles to horizontal image coordinates and elevation angles to vertical coordinates. Range images are widely used in object detection [11], semantic segmentation [12], upsampling [13], and point cloud compression [14], as they enable

efficient computation through their regular grid structure and compatibility with established image processing techniques.

This 3D-to-2D projection is often treated as bijective, under the assumption that each point corresponds to a unique combination of range and angular directions. In ideal spinning LiDAR sensors this assumption holds, as the vertical and horizontal angles take discrete values defined by the beam layout and the rotational sampling rate. This enables exact reconstruction of each 3D point from its position and value in the range image. However, this model does not accurately reflect the behavior of real sensors, whose laser beams originate from slightly different positions and orientations due to mechanical constraints. These variations are typically compensated through factory calibration, which adjusts the measured ranges and angles to compute accurate 3D coordinates. However, the presence of this calibration breaks the assumptions of the spherical projection model, leading to projection mismatches, pixel collisions, and quantization artifacts when computing range images.

In downstream applications, such artifacts are often ignored or mitigated employing ad-hoc techniques. However, in tasks that demand high geometric fidelity even small projection errors can accumulate and degrade performance. A detailed analysis by Wu et al. [15] demonstrates that standard projection strategies introduce measurable distortions, highlighting the need for

*Corresponding author. Postal address: Centro Singular de Investigación en Tecnoloxías Intelixentes (CiTIUS), Rúa de Jenaro de la Fuente Domínguez, 15782 Santiago de Compostela, A Coruña.

Email address: s.soutullo@usc.es (Samuel Soutullo)

more accurate range image generation. While previous work has proposed models that account for laser-specific offsets [16], these approaches rely on manufacturer-supplied lookup tables (LUTs) or packet-level data, limiting their generality. Despite the fundamental importance of this problem for LiDAR-based applications, the challenge of generating truly lossless range images from calibrated point clouds without sensor metadata has been largely overlooked in the literature. To the best of our knowledge, this work presents the first general solution that enables lossless range image generation directly from calibrated point clouds in the absence of sensor metadata.

We introduce ALICE-LRI (Automatic LiDAR Intrinsic Calibration Estimation for Lossless Range Images), a method that infers key geometric parameters of spinning LiDAR sensors directly from point clouds without requiring factory metadata. These inferred parameters are then used to generate range images that are fully consistent with the intrinsic geometry of the sensor and allow exact reconstruction of the original point cloud. The result is a sensor-agnostic, lossless projection algorithm suitable for any application where geometric fidelity is essential. To facilitate adoption, we publish ALICE-LRI as an open-source C++/Python library.¹

The key contributions of this work are:

- The formulation of the fundamental problem of lossless range image generation along with an analysis of how real-world sensor geometry breaks conventional projection methods.
- ALICE-LRI, to the best of our knowledge, the first method to achieve lossless range image generation from calibrated spinning LiDAR point clouds without requiring manufacturer metadata or calibration files by automatically inferring sensor intrinsics directly from raw data.
- A comprehensive evaluation on the KITTI and DurLAR datasets demonstrating the geometric losslessness of our approach, and validation of practical downstream benefits through a compression case study.

The remainder of the paper is organized as follows. Section 2 reviews relevant prior work on range image applications and LiDAR calibration. Section 3 formulates the lossless range image generation problem and analyzes the gap between ideal and real-world sensor models. Section 4 presents the ALICE-LRI method for automatic parameter estimation and lossless projection. Section 5 describes the experimental methodology, datasets, and evaluation metrics. Section 6 reports comprehensive evaluation results, including parameter estimation accuracy, reconstruction quality, and runtime performance. Section 7 demonstrates the application and practical benefits of ALICE-LRI through a point cloud compression case study. Finally, Section 8 concludes the paper, and Section 9 outlines future research directions.

2. Related Work

This section reviews prior work across two fronts: the downstream use cases that rely on range image projections—highlighting the role of projection accuracy in each—and the body of research on LiDAR sensor geometry and calibration.

2.1. Range Image Use Cases

Segmentation. Segmentation pipelines frequently rely on range images to reduce computation and exploit spatial regularities. Classical methods such as the one proposed by Bogoslavskyi and Stachniss [17] perform geometric clustering on cylindrical projections. More recent approaches like SqueezeSeg [12] and RangeNet++ [8] apply convolutional neural networks (CNNs) directly to spherical range images, leveraging their grid structure for efficient inference. RangeNet++, in particular, includes a specialized post-processing module to correct range image discretization errors and smooth out the blurry outputs introduced by CNN inference. Other works explicitly tackle the limitations of the projection process itself: Triess et al. [18] introduce a scan unfolding strategy that reconstructs laser beam structure from point clouds; however, this approach is ad-hoc for the KITTI dataset, as it leverages the known order of points in the files. Kong et al. [19] identify common artifacts—such as many-to-one conflicts and pixel holes—that arise from discretization and quantization. These efforts highlight the sensitivity of segmentation to projection fidelity and the reliance of current methods on dataset-specific heuristics or other strategies that do not generalize across sensors.

Object Detection. Range images are widely used for efficient 3D object detection, enabling the use of fast 2D CNNs in place of computationally expensive 3D convolutions. Early work projected LiDAR scans onto cylindrical grids encoding distance and height [20], while other methods fused range view with bird’s-eye and RGB imagery for enhanced proposal generation [21]. Detection pipelines such as LaserNet [11] and FVNet [22] perform inference directly in the range domain. However, object detection models often overlook the impact of projection fidelity on detection quality, despite facing similar challenges to segmentation. This suggests an opportunity to improve detection robustness by adopting more principled range image generation techniques.

Upsampling. Range images often also serve as the backbone for LiDAR super-resolution, where low-density scans are up-sampled into denser point clouds. Some methods follow a 3D–2D–3D pipeline: sparse scans are projected to 2D, enhanced via neural models, and then backprojected [13, 23, 24], despite the 3D–2D–3D pipeline being prone to introducing quantization artifacts. Some recent approaches have begun to question this dependency on discrete projections [25], but most still assume access to a consistent and accurate mapping between 3D points and 2D pixels.

¹<https://github.com/alice-lri/alice-lri>

Odometry and Localization. Range images also play a central role in LiDAR-based odometry and localization systems, where they serve as structured intermediates for geometric alignment and pose estimation. Behley et al. [5] proposed a surfel-based SLAM system that leverages dense range images to enable efficient frame-to-model ICP registration. Similarly, Chen et al. [3] perform pixel-wise comparisons between live and synthetic range images rendered from a mesh map, enabling GPS-free localization via Monte Carlo inference. More recently, Lichtenfeld et al. [4] introduced a real-time odometry pipeline for dynamic environments that projects both point clouds and residuals onto cylindrical range images. This allows rapid segmentation and tracking of moving objects, which are then filtered out to maintain a clean static map. Across these works, range images consistently boost efficiency and robustness, but their performance hinges on projection quality.

Compression. A significant family of LiDAR compression methods relies on 2D range images to exploit spatial and temporal redundancies inherent in the structured representation. Projection fidelity is especially vital in this class of methods, where information loss directly affects point cloud reconstruction. Early systems [26] encoded raw packet data into structured range images, enabling efficient compression via JPEG or MPEG. Later methods improved spatial and temporal redundancy exploitation by modeling motion or beam trajectories [27, 28]. Other approaches segment range images into planar patches or predict inter-frame deltas to compress efficiently while preserving geometry [29, 30]. Recent methods like RID-DLE [14] rely on neural predictors and delta encoding in the range domain. Yet all of these methods assume that the input range image is a faithful representation of the 3D scan—a violated condition when projecting from calibrated point clouds using classical methods. As a result, authors either sidestep the issue by operating on raw packet data (when available), or accept minor geometric distortion.

While the preceding use cases highlight the importance of accurate range image projections, achieving this fidelity requires knowledge of the intrinsic geometry of the sensor. To better understand this limitation, we briefly review existing work on LiDAR sensor geometry and calibration.

2.2. LiDAR Sensor Geometry and Calibration

Due to manufacturing tolerances and mechanical constraints, real-world LiDAR sensors exhibit geometric non-idealities. These include spatial offsets, angular deviations, and range biases that deviate from theoretical design specifications. LiDAR calibration is a well-established research area concerned with accurately estimating these intrinsic geometric parameters. These parameters are incorporated into the computation of 3D point cloud coordinates, ensuring that output point clouds accurately represent real-world geometry.

Early methods such as [31] introduced optimization-based procedures to refine beam orientations and offsets by minimizing point-to-plane distances across multiple scans of structured environments. More extensive models—such as the one proposed in [32] for the Velodyne HDL-64E—incorporated up to

six parameters per laser and used planar least-squares adjustments to significantly reduce residual errors and reveal systematic biases across the sensor array. Later, Glennie et al. [33] evaluated the factory calibration and long-term stability of the Ouster VLP-16. They showed that even modest improvements in calibration yield noticeable gains in accuracy, though with limited reusability over time. Complementary work by Chan and Lichti [34] extended calibration to use both planar and cylindrical environmental features for the Velodyne HDL-32E. They explicitly modeled additional horizontal angle and range offsets to better capture beam-specific distortions.

Critically, existing calibration research treats this as a forward problem: estimating parameters from raw sensor data to produce geometrically accurate point clouds. However, in many practical scenarios, LiDAR sensors output already-calibrated point cloud coordinates with the calibration parameters embedded in the data processing pipeline. Similarly, most publicly available datasets provide only these final calibrated point clouds with raw data discarded. ALICE-LRI addresses the complementary inverse problem: given a calibrated point cloud, we reverse-engineer the calibration parameters that were applied during its generation. This inverse perspective enables lossless range image generation for many real-world scenarios that lack access to raw sensor packets or calibration metadata. This is a fundamental capability that has received little attention despite its practical importance.

3. Problem Formulation

This section formulates the core challenge of projecting LiDAR point clouds into 2D range images while preserving all original information. We establish the mathematical foundations by examining an ideal sensor model that enables perfect lossless projection, then analyze how real-world sensor geometry breaks this idealized framework, and finally identify the key parameters that must be recovered to restore losslessness.

3.1. Range Image Representation

Range images provide a structured 2D representation of LiDAR scans by mapping spherical coordinates (r, θ, ϕ) to a regular grid, where each pixel (u, v) encodes the range r along a discretized angular direction. The projection assumes a spherical sensor model with L laser beams at fixed elevation angles and H uniformly sampled azimuth angles per revolution, creating a mapping between 3D points and image pixels.

A lossless projection requires that this mapping is bijective: each pixel corresponds to a unique angular direction, and each measured point maps to exactly one pixel. Figure 1 illustrates this discretization of the sensing sphere into a dense grid structure.

Note that projection fidelity depends critically on how accurately the assumed sensor model reflects the true device geometry. To understand when lossless projection is theoretically possible, we first examine an ideal sensor model that perfectly satisfies the spherical assumption.




Figure 1: LiDAR point cloud projection onto a 2D range image grid based on discretized angular coordinates.

3.2. Ideal Sensor Model

Under ideal conditions, a spinning LiDAR sensor exhibits perfect geometric regularity that enables exact lossless projection. We define an ideal sensor as having L laser beams, all originating from a common sensor origin, with each beam oriented at a vertical angle $\varphi^{(l)}$ where $l \in \{0, \dots, L-1\}$. During rotation, the sensor collects H uniformly spaced azimuthal samples per revolution, where each sample $h \in \{0, \dots, H-1\}$ corresponds to an azimuthal angle

$$\theta^{(h)} = \frac{h}{H} \cdot 2\pi. \quad (1)$$

For every pair (l, h) , the sensor measures a range value $r^{(l, h)}$ corresponding to the distance to the first obstacle encountered in the direction defined by the beam.

Under this model, a point $p_i = (x_i, y_i, z_i)$ captured by the laser beam l at horizontal sample h is uniquely determined from its corresponding triplet $(r^{(l, h)}, \varphi^{(l)}, \theta^{(h)})$, such that

$$\begin{cases} x_i = r^{(l, h)} \cdot \cos(\varphi^{(l)}) \cdot \cos(\theta^{(h)}) \\ y_i = r^{(l, h)} \cdot \cos(\varphi^{(l)}) \cdot \sin(\theta^{(h)}) \\ z_i = r^{(l, h)} \cdot \sin(\varphi^{(l)}). \end{cases} \quad (2)$$

Conversely, given a point $p_i = (x_i, y_i, z_i)$, its spherical coordinates are computed as

$$\begin{cases} r_i = \sqrt{x_i^2 + y_i^2 + z_i^2} \\ \varphi_i = \arcsin\left(\frac{z_i}{r_i}\right) \\ \theta_i = \arctan\left(\frac{y_i}{x_i}\right). \end{cases} \quad (3)$$

Since φ_i and θ_i in the ideal model are always drawn from discrete sets of known angles, they must satisfy $\varphi_i \in \{\varphi^{(l)}\}, \forall i$ and $\theta_i \in \{\theta^{(h)}\}, \forall i$. The range r_i can take any real, positive value.

Under these conditions, generating a lossless range image of size $H \times L$ is straightforward. First, compute the spherical

coordinates of each point as defined by Equation (3). Then, map each tuple $(r_i, \varphi_i, \theta_i)$ to a pixel (u_i, v_i) in the image using the indices (l, h) of its associated beam angles as

$$\begin{cases} u_i := h & \text{such that } \theta_i = \theta^{(h)} \\ v_i := l & \text{such that } \varphi_i = \varphi^{(l)}. \end{cases} \quad (4)$$

The corresponding pixel at position (u_i, v_i) is assigned the value r_i . Pixels for which no point is present are marked with a sentinel value (e.g., zero or NaN) to indicate absence of measurement.

This ideal model demonstrates that lossless projection is theoretically achievable when sensor geometry exactly matches the spherical assumption. However, real sensors deviate significantly from this idealized framework.

3.3. Real-World Sensor Model

Practical LiDAR sensors exhibit geometric non-idealities that fundamentally break the lossless projection capability. Manufacturing tolerances and mechanical constraints cause each laser beam to originate from slightly different physical locations, introducing spatial offsets relative to the sensor origin. Additionally, beams may exhibit unique range deviations due to mechanical and optical factors, and azimuthal offsets from the nominal rotation angle.


These non-idealities render the equations from the ideal model (Equation (2)) insufficient. In response, manufacturers typically define device-specific calibration models, which vary in format but share a common structure: point coordinates depend on beam-specific angles, positional offsets, and range corrections.

We adopt a general formulation that captures these factors. Each laser beam l is characterized by five key parameters, as illustrated in Figure 2: a vertical angle $\varphi^{(l)}$, spatial offsets $(o_x^{(l)}, o_y^{(l)})$, an azimuthal offset $\theta_{\text{off}}^{(l)}$, and a horizontal resolution $H^{(l)}$ representing the number of azimuth samples per revolution for beam l . Although many practical sensors employ uniform horizontal resolution across all beams, we allow beam-specific resolutions for generality. Spatial corrections along the forward-backward axis are omitted because they can be absorbed into the horizontal and vertical offsets $(o_x^{(l)}, o_y^{(l)})$, affecting only the measured range r_i without influencing the angular computations central to our projection method. Without loss of generality, we assume that beam indices are ordered by increasing vertical angle,


$$\varphi^{(0)} < \varphi^{(1)} < \dots < \varphi^{(L-1)}. \quad (5)$$

Under this calibrated model, a 3D point $p_i = (x_i, y_i, z_i)$ measured by beam l includes additional residual terms:

$$\begin{cases} x_i = r^{(l, h^{(l)})} \cdot \cos\left(\varphi^{(l)} + \varphi_{\text{res}}^{(l, h^{(l)})}\right) \cdot \cos\left(\theta^{(h^{(l)})} + \theta_{\text{off}}^{(l)} + \theta_{\text{res}}^{(l, h^{(l)})}\right) \\ y_i = r^{(l, h^{(l)})} \cdot \cos\left(\varphi^{(l)} + \varphi_{\text{res}}^{(l, h^{(l)})}\right) \cdot \sin\left(\theta^{(h^{(l)})} + \theta_{\text{off}}^{(l)} + \theta_{\text{res}}^{(l, h^{(l)})}\right) \\ z_i = r^{(l, h^{(l)})} \cdot \sin\left(\varphi^{(l)} + \varphi_{\text{res}}^{(l, h^{(l)})}\right), \end{cases} \quad (6)$$



(a) Side view of a LiDAR sensor illustrating two laser beams: an ideal beam (blue, dashed) originating from the sensor's origin, and a real beam (red, solid) with an offset origin. The elevation angles along the ideal beam ($\varphi^{(l)}$) remain constant with respect to the origin of coordinates, while those along the real beam vary ($\varphi_1 \neq \varphi_2$) due to the offset $o_y^{(l)}$ that introduces residual angles $\varphi_{\text{res}}^{(l,h^{(l)})}$ such that $\varphi_1 = \varphi^{(l)} + \varphi_{\text{res}}^{(l,1)}$ and $\varphi_2 = \varphi^{(l)} + \varphi_{\text{res}}^{(l,2)}$.



(b) Top view of a LiDAR sensor illustrating two laser beams: an ideal beam (blue, dashed) originating from the sensor's center, and a real beam (red, solid) with a horizontal offset $o_x^{(l)}$ and an azimuthal offset $\theta_{\text{off}}^{(l)}$. The azimuthal angles along the ideal beam remain constant with respect to the origin of coordinates ($\theta^{(h^{(l)})}$), while those along the real beam vary due to the offset that introduces residual angles $\theta_{\text{res}}^{(l,h^{(l)})}$. Further, the azimuthal offset $\theta_{\text{off}}^{(l)}$ shifts the angle of the real beam by a constant amount.

Figure 2: Illustration of geometric differences between the ideal sensor model (perfectly aligned beams from a common origin) and the real-world model (beams with spatial and angular offsets), highlighting the calibration parameters required for lossless projection.

where the residual angles $\varphi_{\text{res}}^{(l,h^{(l)})}$ and $\theta_{\text{res}}^{(l,h^{(l)})}$ are derived from the spatial offsets ($o_x^{(l)}, o_y^{(l)}$) such that

$$\begin{cases} \varphi_{\text{res}}^{(l,h^{(l)})} = \arcsin\left(\frac{o_y^{(l)}}{r^{(l,h^{(l)})}}\right) \\ \theta_{\text{res}}^{(l,h^{(l)})} = \arcsin\left(\frac{o_x^{(l)}}{\rho^{(l,h^{(l)})}}\right), \end{cases} \quad (7)$$

where $\rho^{(l,h^{(l)})}$ is the distance in the XY plane of point p_i , measured by beam l at horizontal sample $h^{(l)}$, defined as

$$\rho^{(l,h^{(l)})} = \sqrt{x_i^2 + y_i^2} = r^{(l,h^{(l)})} \cdot \cos(\varphi^{(l)} + \varphi_{\text{res}}^{(l,h^{(l)})}). \quad (8)$$

The azimuthal sampling angles $\theta^{(h^{(l)})}$ are defined similarly to the ideal model, but with beam-specific horizontal resolutions $H^{(l)}$ such that

$$\theta^{(h^{(l)})} = \frac{h^{(l)}}{H^{(l)}} \cdot 2\pi. \quad (9)$$

Because of these additional corrections, using the point coordinates to compute the spherical angles φ_i and θ_i as shown in Equation (3), no longer yields results in the discrete sets $\{\varphi^{(l)}\}$ and $\{\theta^{(h^{(l)})}\}$ defined in the ideal model. Instead, they are continuous values defined as:

$$\begin{cases} \varphi_i = \varphi^{(l)} + \varphi_{\text{res}}^{(l,h^{(l)})} \\ \theta_i = \theta^{(h^{(l)})} + \theta_{\text{off}}^{(l)} + \theta_{\text{res}}^{(l,h^{(l)})}. \end{cases} \quad (10)$$

As a result, even under the best circumstances, a point will not be projected to the exact center of its corresponding pixel, reducing geometric precision. This misalignment often leads to multiple points being projected onto the same pixel (collisions),

or to empty pixels appearing where measurements should exist (holes). These effects fundamentally prevent standard projection strategies from achieving losslessness.

3.4. Lossless Range Image Generation

The prior discussion highlights that the main obstacle to lossless projection is the presence of unknown beam-specific parameters $\{\varphi^{(l)}, o^{(l)}, \theta_{\text{off}}^{(l)}, H^{(l)}\}$. If known, these values can invert the calibration transformations applied to the observed spherical angles, recovering the discrete angular sets required for bijective mapping. In particular, given a point p_i measured by beam l at horizontal sample $h^{(l)}$, it is possible to compute the corrected angles (φ'_i, θ'_i) using the observed angles (φ_i, θ_i) and the residual terms from Equation (7) as:

$$\begin{cases} \varphi'_i = \varphi_i - \varphi_{\text{res}}^{(l,h^{(l)})} \\ \theta'_i = \theta_i - \theta_{\text{off}}^{(l)} - \theta_{\text{res}}^{(l,h^{(l)})}. \end{cases} \quad (11)$$

With these corrections, $\varphi'_i \in \{\varphi^{(l)}\}$ and $\theta'_i \in \{\theta^{(h^{(l)})}\}$. This condition guarantees bijective mapping and enables lossless projection. Moreover, the range image dimensions are determined by L (vertical resolution) and the least common multiple of all $H^{(l)}$ (horizontal resolution).

Consequently, recovering these sensor-specific parameters is essential for lossless range image generation. This motivates our approach: we propose to infer the factory-calibrated geometric parameters directly from point cloud data, enabling the generation of range images that are intrinsically consistent with sensor geometry.

4. ALICE-LRI

To solve the problem formulated in Section 3.4, we present ALICE-LRI, a general method that automatically recovers the




Figure 3: General outline of ALICE-LRI for estimating LiDAR intrinsic parameters and generating lossless range images. Gray diamond-shaped nodes represent data inputs and outputs; rectangular nodes represent processing steps. Solid lines denote execution flow, while dashed lines indicate data flow. The method first estimates sensor parameters from the input point cloud, then uses these parameters to generate range images from subsequent point clouds.

sensor parameters needed for the angle corrections in Equation (11). The method operates directly on point cloud data without requiring sensor metadata or scene knowledge. As illustrated in Figure 3, ALICE-LRI comprises two main stages: parameter estimation and lossless projection.

In the first stage, the intrinsic parameters are automatically estimated from the point cloud data, including the number of laser beams L , the vertical angles $\varphi^{(l)}$, the vertical offsets $o_y^{(l)}$, the horizontal angular resolutions $H^{(l)}$, the horizontal offsets $o_x^{(l)}$, and the azimuthal offsets $\theta_{\text{off}}^{(l)}$. In the second stage, the estimated parameters are used to project the point cloud onto a range image. Since parameter estimation needs to be performed only once per device, the recovered parameters can be subsequently reused to generate range images from multiple point clouds acquired with the same hardware.

The core of the method lies in the parameter estimation stage, which is further subdivided into two phases: vertical parameter estimation and horizontal parameter estimation. The vertical phase determines the number of beams L , assigns points to each beam, and simultaneously estimates $\varphi^{(l)}$ and $o_y^{(l)}$. The horizontal phase then leverages these assignments to estimate $H^{(l)}$, $o_x^{(l)}$, and $\theta_{\text{off}}^{(l)}$. The following subsections provide a detailed description of these steps.

4.1. Vertical Parameter Estimation: L , $\varphi^{(l)}$, and $o_y^{(l)}$

As discussed in Section 3.3, the vertical angle φ_i of each point is affected by a beam-specific vertical offset. Substituting the angular residual $\varphi_{\text{res}}^{(l,h^{(l)})}$ from Equation (7) into Equation (10) yields

$$\varphi_i = \varphi^{(l)} + \arcsin\left(\frac{o_y^{(l)}}{r_i}\right). \quad (12)$$

The objective of the vertical parameter estimation is to recover the set of scanlines that best fit the observed data under this model.

To this end, we employ an iterative consensus-based algorithm that progressively identifies candidate scanlines, assigns points to them, and refines their parameters. At each iteration, the algorithm performs four steps: (1) proposes candidate scanline parameters $(\varphi^{(l)}, o_y^{(l)})$, (2) selects points consistent with

these parameters, (3) fits the scanline to refine the values of $(\varphi^{(l)}, o_y^{(l)})$, and (4) resolves conflicts with previously accepted scanlines. The process continues until all points are assigned to scanlines or the search space is exhausted. An overview of the proposed algorithm is shown in Figure 4. The subsequent subsections describe each component of the algorithm in detail.

4.1.1. Candidate Scanline Parameter Identification via the Hough Transform

The first step in each iteration is to identify candidate scanline parameters $(\varphi^{(l)}, o_y^{(l)})$. To this end, we adopt the Hough Transform [35, 36], a classical feature extraction technique widely used to detect parametric curves in noisy data. While originally developed for line and circle detection in 2D images [37, 38], the method has since been extended to 3D point clouds for extracting structures such as building roofs [39], powerlines [40], and trees [41].

The central principle of the Hough Transform is to map each observation in the measurement space to a curve in a discretized parameter space. Observations belonging to the same structure yield curves that intersect at a common point in the parameter space. The coordinates of this intersection define the underlying structure. In our case, each observation (r_i, φ_i) is mapped to a curve in the parameter space (φ', o_y) obtained by rearranging Equation (12):

$$\varphi'_i = \varphi_i - \arcsin\left(\frac{o_y}{r_i}\right). \quad (13)$$

Thus, points from the same scanline l generate curves that intersect at the true parameters $(\varphi^{(l)}, o_y^{(l)})$. To make these intersections explicit, we discretize the parameter space into an accumulator $\mathbf{A}(\tilde{\varphi}', \tilde{o}_y)$ where

$$\tilde{\varphi}' = \tilde{\varphi}'^{\min} + j \cdot \Delta\tilde{\varphi}', \quad j = 0, 1, 2, \dots, \left\lfloor \frac{\tilde{\varphi}'^{\max} - \tilde{\varphi}'^{\min}}{\Delta\tilde{\varphi}'} \right\rfloor \quad (14)$$

and

$$\tilde{o}_y = \tilde{o}_y^{\min} + k \cdot \Delta\tilde{o}_y, \quad k = 0, 1, 2, \dots, \left\lfloor \frac{\tilde{o}_y^{\max} - \tilde{o}_y^{\min}}{\Delta\tilde{o}_y} \right\rfloor \quad (15)$$

with step sizes $\Delta\tilde{\varphi}'$ and $\Delta\tilde{o}_y$. For each observation (r_i, φ_i) and




Figure 4: Flowchart of the vertical parameter estimation algorithm (Sec. 4.1). The algorithm iteratively identifies candidate scanline parameters using the Hough Transform, selects points consistent with these parameters, fits the scanline to refine parameter values, and resolves conflicts with previously accepted scanlines. The process continues until all points are assigned to scanlines or the search space is exhausted. Key processing steps are annotated with references to their corresponding subsections.

each discretized $\tilde{o}_y$, we compute the corresponding $\tilde{\phi}'$ and increment the corresponding accumulator cell:

$$\forall i, \tilde{o}_y: \mathbf{A}(\tilde{\phi}', \tilde{o}_y) \leftarrow \mathbf{A}(\tilde{\phi}', \tilde{o}_y) + 1$$

$$\text{where } \tilde{\phi}' = \left\lfloor \frac{\phi_i - \arcsin\left(\frac{\tilde{o}_y}{r_i}\right)}{\Delta\tilde{\phi}'} + 0.5 \right\rfloor \cdot \Delta\tilde{\phi}'. \quad (16)$$

In our implementation, the parameter space is discretized with $\tilde{o}_y^{\max} = \min(\min_i r_i, 0.5 \text{ m})$, $\tilde{o}_y^{\min} = -\tilde{o}_y^{\max}$ and $\Delta\tilde{o}_y = 10^{-3} \text{ m}$. The angular range is fixed to $\tilde{\phi}' \in [-\pi/2, \pi/2]$ with $\Delta\tilde{\phi}' = 10^{-4} \text{ rad}$. These values are physically motivated: the vertical offset is bounded by the physical dimensions of LiDAR sensors (where offsets beyond $\pm 0.5 \text{ m}$ would be implausible given typical sensor sizes), and it is further clipped by the minimum observed point range to ensure the arcsine argument remains valid. The vertical angles are constrained to $\pm 90^\circ$ since this represents the theoretical maximum range for laser emission. These values provide sufficient resolution to capture realistic LiDAR sensor parameters while maintaining computational tractability.

After all points are processed, peaks in $\mathbf{A}$ correspond to consistent parameter combinations and are selected as candidate scanlines. The iterative phase of the algorithm then begins: at each iteration, the cell with the highest vote count is selected as the next candidate. If accepted, the votes of its contributing points are removed from the accumulator. This strategy, first described as adaptive cluster detection [42], prevents previously assigned points from influencing later iterations and enables the progressive extraction of further scanlines.

To improve robustness and efficiency, we extend the standard Hough Transform with two modifications. First, we introduce a vote-for-discontinuities strategy: when a vote skips multiple cells in the accumulator (for example, due to steep gradients in the mapping), intermediate cells are also incremented, thereby reducing the risk of fragmented responses (see Figure 5). Second, we maintain a hash accumulator $\mathbf{B}(\tilde{\phi}', \tilde{o}_y)$ that stores a compact hash of the contributing point indices on each cell. This allows us to quickly determine whether different cells are supported by the same set of points, which is useful for equivalence checks in later stages.

The hashing scheme uses a 64-bit Knuth-style hash [43] applied to each point index and combined by bitwise XOR:

$$\kappa(i) := ((i+1)c) \bmod 2^{64}, \quad c = 11400714819323198485,$$

$$\mathbf{B}(\tilde{\phi}', \tilde{o}_y) = \bigoplus_{i \in \mathbf{S}(\tilde{\phi}', \tilde{o}_y)} \kappa(i), \quad (17)$$

where $\mathbf{S}(\tilde{\phi}', \tilde{o}_y)$ is the set of point indices that voted for cell $(\tilde{\phi}', \tilde{o}_y)$. Cells with identical values therefore share the same contributing set, up to a negligible collision probability.

Despite its utility, the Hough Transform also presents limitations. It identifies candidate parameter locations but does not directly evaluate their quality of fit. Noise may generate spurious peaks or fragment a single scanline into multiple candidates, while sparse scanlines may yield weak responses. Unrelated points can occasionally accumulate votes for an incorrect parameter combination. In addition, achieving higher precision requires finer discretization of the accumulator, which increases




Figure 5: Comparison of basic Hough Transform strategy versus vote-for-discontinuities strategy. The basic strategy (a) misses the intersection peak due to vote fragmentation across cells, while the discontinuities strategy (b) successfully detects it by incrementing intermediate cells when votes skip multiple bins due to steep gradients in the parameter mapping.

both memory and computational costs. For these reasons, we employ the Hough Transform only as the initial stage of our algorithm. The candidate parameters it produces are then refined and validated in subsequent steps.

4.1.2. Scanline Point Selection via Error Bounds

Once candidate scanline parameters $(\varphi^{(l)}, o_y^{(l)}) = (\tilde{\varphi}', \tilde{o}_y)$ are identified using the Hough Transform, the next step is to select points that are consistent with these parameters. This filtering stage is essential for the subsequent fitting process, as it ensures that only relevant points contribute to the model.

Due to the finite precision of point cloud coordinates—arising from sensor quantization, rounding, or floating-point representation—points do not generally lie exactly on the theoretical scanline defined by Equation (12). Some sensors and datasets quantize their coordinates to discrete levels; for example, the KITTI raw dataset [44] quantizes coordinates to the nearest millimeter. This coordinate quantization directly propagates as error in the vertical angles φ_i . To account for this, we derive an angular error bound $\delta\varphi_i$ from the coordinate quantization level. That is, a point i is considered part of scanline l if

$$\left| \varphi_i - \varphi^{(l)} - \arcsin\left(\frac{o_y^{(l)}}{r_i}\right) \right| \leq \delta\varphi_i. \quad (18)$$

Since we operate without prior knowledge of the sensor or dataset specifications, we must infer the coordinate quantization level directly from the data itself. We achieve this by examining the minimum non-zero differences between coordinate components across all points in the point cloud.

Assume each Cartesian coordinate is quantized with some unknown step $\Delta_\alpha > 0$, i.e. $x_i, y_i, z_i \in \Delta_\alpha \mathbb{Z}$. Then every pairwise difference along an axis is an integer multiple of Δ_α , so the minimum nonzero spacing along that axis is at least Δ_α . Let

$$\Delta_x^{\min} := \min \{ |x_i - x_j| \neq 0 \}$$

be the minimum nonzero difference between the x -coordinates (analogously $\Delta_y^{\min}, \Delta_z^{\min}$). We assume that at least one axis exhibits two occupied adjacent quantization levels in the frame². Under this assumption, we can reliably estimate the quantization level as the smallest of these minimum differences:

$$\hat{\Delta}_\alpha := \min \{ \Delta_x^{\min}, \Delta_y^{\min}, \Delta_z^{\min} \},$$

and $\varepsilon = \hat{\Delta}_\alpha/2$ as the error bound between the observed and the ideal (non-quantized) coordinates. To guard against non-quantized data while covering potential errors introduced by floating-point representation, we clamp $\varepsilon \leftarrow \max(\varepsilon, 10^{-6} \text{ m})$.

This upper bound on coordinate error ε can then be propagated to derive $\delta\varphi_i$. Let $\mathbf{u}_i = (x_i^*, y_i^*, z_i^*)$ represent the ideal coordinates and $\mathbf{v}_i = (x_i^* \pm \varepsilon, y_i^* \pm \varepsilon, z_i^* \pm \varepsilon)$ the observed coordinates. Defining $\Delta\mathbf{u} = (\pm\varepsilon, \pm\varepsilon, \pm\varepsilon)$, we have $\mathbf{v}_i = \mathbf{u}_i + \Delta\mathbf{u}$. By the triangle inequality,

$$\|\mathbf{v}_i\| = \|\mathbf{u}_i + \Delta\mathbf{u}\| \leq \|\mathbf{u}_i\| + \|\Delta\mathbf{u}\| = \|\mathbf{u}_i\| + \sqrt{3}\varepsilon.$$

Thus, with $r_i^* = \|\mathbf{u}_i\|$ and $r_i = \|\mathbf{v}_i\|$,

$$\delta r = |r_i^* - r_i| \leq \sqrt{3}\varepsilon. \quad (19)$$

Defining the angular error bound $\delta\varphi_i$ as the difference between the observed and ideal vertical angles, we have:

$$\delta\varphi_i = \left| \arcsin\left(\frac{z_i^*}{r_i^*}\right) - \arcsin\left(\frac{z_i}{r_i}\right) \right|.$$

Since $\arcsin(z_i/r_i) = \arctan(z_i/\rho_i)$ with $\rho_i = \sqrt{x_i^2 + y_i^2}$,

$$\delta\varphi_i = \left| \arctan\left(\frac{z_i^*}{\rho_i^*}\right) - \arctan\left(\frac{z_i}{\rho_i}\right) \right|.$$

Applying similar reasoning as in Equation (19), we bound

$$\delta\rho = |\rho_i^* - \rho_i| \leq \sqrt{2}\varepsilon.$$

By the Mean Value Theorem, for some ξ between z_i^*/ρ_i^* and z_i/ρ_i ,

$$\arctan\left(\frac{z_i^*}{\rho_i^*}\right) - \arctan\left(\frac{z_i}{\rho_i}\right) = \frac{1}{1+\xi^2} \left(\frac{z_i^*}{\rho_i^*} - \frac{z_i}{\rho_i} \right).$$

Since $\frac{1}{1+\xi^2} \leq 1$, it follows that

$$\delta\varphi_i \leq \left| \frac{z_i^*}{\rho_i^*} - \frac{z_i}{\rho_i} \right| = \frac{|z_i^*\rho_i - z_i\rho_i^*|}{|\rho_i\rho_i^*|}.$$

We now derive an upper bound for $\delta\varphi_i$ by establishing an upper bound on the numerator and a lower bound on the denominator. For the numerator, write

$$\rho_i^* = \rho_i + \Delta\rho_i, \quad \Delta\rho_i \in [-\delta\rho, \delta\rho] = [-\sqrt{2}\varepsilon, \sqrt{2}\varepsilon],$$

²In practice, where a typical LiDAR frame contains around 10^5 points, this assumption holds with overwhelming probability, as shown in our experiments (see Section 6).




Figure 6: Flowchart of the scanline fitting via weighted least squares procedure (Sec. 4.1.3). The algorithm performs iterative WLS fitting to refine vertical parameters $(\hat{\phi}^{(l)}, \hat{\delta}_y^{(l)})$ and compute point memberships. A heuristic fallback is employed to handle cases where the fitting process fails to converge.

so that

$$\begin{aligned}
 |z_i^* \rho_i - z_i \rho_i^*| &= |z_i^* \rho_i - z_i(\rho_i + \Delta \rho_i)| \\
 &= |(z_i^* - z_i) \rho_i - z_i \Delta \rho_i| \\
 &\leq |(z_i^* - z_i) \rho_i| + |z_i \Delta \rho_i| \quad (\text{since } |a - b| \leq |a| + |b|) \\
 &\leq \varepsilon \rho_i + |z_i| \sqrt{2} \varepsilon \quad (\text{since } |z_i^* - z_i| \leq \varepsilon, |\Delta \rho_i| \leq \sqrt{2} \varepsilon) \\
 &\leq \varepsilon (\rho_i + \sqrt{2} |z_i|).
 \end{aligned}$$

For the denominator we have

$$\rho_i^* = \rho_i + \Delta \rho_i \geq \rho_i - |\Delta \rho_i| \geq \rho_i - \delta \rho,$$

hence

$$|\rho_i \rho_i^*| = \rho_i |\rho_i + \Delta \rho_i| \geq \rho_i (\rho_i - \delta \rho) \geq \rho_i^2 - \sqrt{2} \varepsilon \rho_i.$$

Combining the bounds for the numerator and the denominator, we obtain

$$\delta \varphi_i \leq \frac{\varepsilon (\rho_i + \sqrt{2} |z_i|)}{\rho_i^2 - \sqrt{2} \varepsilon \rho_i}. \quad (20)$$

This bound provides a principled criterion to identify points that belong to a candidate scanline (as defined by Equation (18)) while explicitly quantifying the uncertainty in φ_i . Both aspects are crucial for the subsequent fitting stage.

4.1.3. Scanline Fitting via Weighted Least Squares

At this stage of the algorithm, we are given a candidate set of points that may form a scanline—i.e., a set of points likely to originate from the same laser beam. Now the objective is to estimate the scanline’s vertical parameters: the true vertical angle $\varphi^{(l)}$ and the vertical offset $o_y^{(l)}$. This is achieved by fitting these points to the model defined in Equation (12). Figure 6 shows the flowchart of the fitting process.

Note that fitting this model directly is non-trivial as it involves a nonlinear arcsine function, which does not have a

closed-form least-squares solution and typically requires iterative optimization techniques. However, we exploit the fact that, in practice, the vertical offset $o_y^{(l)}$ is small relative to the range r_i of each point (e.g., $o_y^{(l)}$ in centimeters vs. r_i in meters). This allows us to apply the approximation

$$\arcsin(x) \approx x \quad \text{for small } x. \quad (21)$$

Applying this approximation to Equation (12), we obtain

$$\varphi_i \approx \varphi^{(l)} + \frac{o_y^{(l)}}{r_i}.$$

Introducing $q_i = 1/r_i$, we rewrite the equation as:

$$\varphi_i \approx \varphi^{(l)} + o_y^{(l)} q_i. \quad (22)$$

This linear form enables us to fit the data (φ_i, q_i) with a straight line, where the intercept corresponds to $\varphi^{(l)}$ and the slope to $o_y^{(l)}$. However, unlike ordinary least squares, we cannot assume uniform noise across observations: as established in Equation (18), each measurement φ_i comes with its own error bound $\delta \varphi_i$, leading to heteroscedastic residuals.

To account for this non-uniform uncertainty, a weighted least squares (WLS) framework is used. We assign each observation a weight $w_i = 1/\delta \varphi_i^2$, which down-weights noisier points relative to more reliable ones. While $\delta \varphi_i$ is not a true standard deviation, Popoviciu’s inequality on variances [45] guarantees that for any bounded error $\Delta \varphi_i \in [-\delta \varphi_i, \delta \varphi_i]$, the variance satisfies $\text{Var}(\Delta \varphi_i) \leq \delta \varphi_i^2$. This result justifies treating $\delta \varphi_i^2$ as a conservative upper bound on the residual variance, providing a principled basis for WLS under a heteroscedastic noise model.

Given a set of n pairs (q_i, φ_i) , the model assumes the linear relation described in Equation (22) with the optimal slope $\hat{o}_y^{(l)}$ and intercept $\hat{\varphi}^{(l)}$ obtained by minimizing the weighted residual

sum of squares

$$\min_{\hat{\delta}_y^{(l)}, \hat{\phi}^{(l)}} \sum_{i=1}^n w_i R_i^2, \quad \text{where } R_i = \phi_i - \hat{\delta}_y^{(l)} q_i - \hat{\phi}^{(l)}.$$

This yields closed-form estimators computed using weighted moments:

$$\hat{\delta}_y^{(l)} = \frac{SS_{xy} - S_x S_y}{SS_{xx} - S_x^2}, \quad \hat{\phi}^{(l)} = \frac{S_{xx} S_y - S_x S_{xy}}{SS_{xx} - S_x^2},$$

where $S = \sum w_i$, $S_x = \sum w_i q_i$, $S_y = \sum w_i \phi_i$, $S_{xx} = \sum w_i q_i^2$, and $S_{xy} = \sum w_i q_i \phi_i$. To quantify the uncertainty in the estimated parameters, confidence intervals are computed based on the variances of the estimators

$$\text{Var}(\hat{\delta}_y^{(l)}) = \sigma^2 \frac{S}{SS_{xx} - S_x^2}, \quad \text{Var}(\hat{\phi}^{(l)}) = \sigma^2 \frac{S_{xx}}{SS_{xx} - S_x^2},$$

where

$$\sigma^2 = \frac{1}{n-2} \sum_{i=1}^n w_i R_i^2$$

is the weighted residual variance.

We employ the Student's t -distribution to construct 95% confidence intervals, using $n-2$ degrees of freedom to reflect the loss of two degrees of freedom from estimating $\hat{\delta}_y^{(l)}$ and $\hat{\phi}^{(l)}$. This choice is justified by the small-sample regime, since some scanlines may contain only a few tens of points. Under the assumption of normally distributed residuals within the error bounds, the t -distribution is more appropriate than the asymptotic normal approximation.

To evaluate model plausibility, we compute a log-likelihood under the assumption of Gaussian errors with heteroscedastic variance, using the conservative bounds $\delta \phi_i^2$ as proxies for the variances:

$$\log \mathcal{L} = -\frac{1}{2} \sum_{i=1}^n \left[\log(2\pi \delta \phi_i^2) + \frac{R_i^2}{\delta \phi_i^2} \right]$$

This surrogate log-likelihood captures both the fit quality and the uncertainty structure of the data. Higher values indicate that the observed residuals are small relative to the conservative noise bounds, making the model more statistically plausible. For practical purposes, we define an uncertainty score as the negative log-likelihood,

$$\mathcal{U} = -\log \mathcal{L}, \quad (23)$$

which enables direct comparison between alternative fits: lower scores indicate better alignment with the observed data and its bounded-error structure.

A single WLS fit would be sufficient if the initial estimates from the Hough transform were already accurate and the corresponding candidate points were correctly assigned. In practice, however, this is not always the case, so a more robust iterative refinement is required. After the initial WLS fit, the estimated parameters $(\hat{\delta}_y^{(l)}, \hat{\phi}^{(l)})$ are used to reassess point membership based on the criterion in Equation (18), yielding an updated

point set. The model is then re-fitted to this refined set, with the process repeated until point assignments stabilize or a maximum of 10 iterations is reached. This iterative scheme ensures convergence to a stable, self-consistent solution both for parameters and point memberships.

Fallback: Heuristic Estimation. In some edge cases, the WLS fitting procedure may fail to converge. Namely when the number of candidate points is insufficient or when the convergence is not achieved within the iteration limit. As a last resort for these rare scenarios where the mathematical fitting is insufficient, we introduce a heuristic fallback that leverages assumptions based on the physical arrangement of the laser beams of real LiDAR sensors. Specifically, it assumes that the vertical offset of the laser beam can be approximated as the average offset of the two neighboring lasers, that is $\hat{\delta}_y^{(l)} = \frac{1}{2} (\hat{\delta}_y^{(l-1)} + \hat{\delta}_y^{(l+1)})$. This is reasonable, as physically adjacent lasers tend to have similar vertical angles. This method is only triggered when some reliable scanlines have already been computed, as it depends on existing parameter estimates. Once the offset $\hat{\delta}_y^{(l)}$ is heuristically determined, the vertical angle $\hat{\phi}^{(l)}$ is directly obtained by rearranging Equation (22) and computing the mean value.

$$\hat{\phi}^{(l)} = \frac{1}{N} \sum_{i=1}^N (\phi_i - \hat{\delta}_y^{(l)} q_i) \quad (24)$$

After either the WLS or the heuristic fit, the algorithm recomputes which points fall within the estimated scanline bounds. If the point set is empty the scanline candidate is marked as invalid and rejected. To avoid redundant processing of identical scanlines detected in multiple Hough cells, once a scanline is rejected we employ the hash accumulator of Section 4.1.1, resetting all cells with identical hashes to zero:

$$\forall (u, v) : \mathbf{B}(u, v) = \mathbf{B}(\phi', \delta_y) \implies \mathbf{A}(u, v) = 0, \quad (25)$$

thereby ensuring that identical point sets are not repeatedly re-evaluated. Conversely, if either the WLS or the heuristic fit is successful, the scanline is passed to the conflict resolution step.

4.1.4. Scanline Conflict Resolution

Up to this point, the algorithm has focused on estimating individual scanlines in isolation. Accepting each candidate greedily would risk inconsistencies, such as multiple scanlines claiming the same points or geometrically implausible intersections. These issues are especially likely in later iterations, when candidate scanlines are supported by fewer points and thus become more sensitive to noise or imperfect fitting. To address this, we introduce a conflict resolution step that enforces global consistency and supports backtracking, enabling previously accepted scanlines to be discarded in favor of better alternatives. Its objectives are twofold: (1) to resolve conflicts between the current scanline and those already accepted, thereby preventing inconsistencies in the final set, and (2) to allow recovery from sub-optimal local decisions, increasing the likelihood of a globally




Figure 7: Flowchart of the scanline conflict resolution process (Sec. 4.1.4). The algorithm checks for intersections between the current scanline candidate and previously accepted scanlines. If intersections exist, uncertainty scores are compared to determine which scanlines to accept or reject. Rejected scanlines are stored for potential recovery if their conflicting counterparts are later invalidated.

consistent solution. An overview of the conflict resolution process is shown in Figure 7.

Conflicts may arise as point intersections, where at least one point is simultaneously claimed by multiple scanlines, or as geometric intersections, where fitted scanlines (including their confidence intervals) overlap within the valid range of the point cloud. Both checks are performed. If no intersections are found, the scanline is accepted: its parameters are stored, points are assigned, and their votes removed from the Hough accumulator (Section 4.1.1); otherwise, conflicts must be resolved.

The conflict resolution process relies on the uncertainty scores computed during fitting, as given by Equation (23). A new scanline with higher uncertainty than any conflicting counterpart is deemed less reliable and rejected. Conversely, if it is more reliable than all conflicting counterparts, it is accepted, and the conflicting scanlines are invalidated. Rejected candidates are not permanently discarded: their parameters and rejection metadata—including the identities of the scanlines causing invalidation—are stored for potential recovery. After invalidating conflicting scanlines, the algorithm iterates over previously rejected candidates. If any no longer exhibit conflicts, they are reintroduced into the Hough accumulator for potential future selection. This backtracking mechanism allows the algorithm to recover from earlier suboptimal choices and promotes convergence to a globally consistent solution.

Once conflict resolution is complete, the new scanline is either accepted or rejected, and the algorithm continues accordingly. When scanlines are rejected, the same hash-based cell reset procedure described in Sec. 4.1.3 is applied to avoid redundant processing of identical scanlines (see Equation (25)).

4.2. Horizontal Parameter Estimation: $H^{(l)}$, $o_x^{(l)}$, and $\theta_{\text{off}}^{(l)}$

After estimating the vertical parameters, the algorithm has identified the number of laser beams L , the vertical angles $\varphi^{(l)}$, the vertical offsets $o_y^{(l)}$, and the assignments of each point in the cloud to its originating laser beam. The next stage is to estimate the horizontal parameters: the horizontal angular resolution $H^{(l)}$, the horizontal offset $o_x^{(l)}$, and the azimuthal offset $\theta_{\text{off}}^{(l)}$ of each beam. Expanding the second expression in Equa-

tion (10), the observed azimuth θ_i can be expressed as

$$\theta_i = \theta^{(h^{(l)})} + \theta_{\text{off}}^{(l)} + \theta_{\text{res}}^{(l, h^{(l)})} = \frac{h^{(l)}}{H^{(l)}} \cdot 2\pi + \theta_{\text{off}}^{(l)} + \arcsin\left(\frac{o_x^{(l)}}{\rho_i}\right). \quad (26)$$

The objective of this step is to recover the set $\{H^{(l)}, o_x^{(l)}, \theta_{\text{off}}^{(l)}\}$ for each beam l such that the scanline observations are best fitted. We treat each scanline independently, without assuming correlations across beams, which enhances robustness to variations in sensor design. An overview of the algorithm is shown in Figure 8.

Using Equation (21), a simplification is obtained by exploiting that $o_x^{(l)} \ll \rho_i$. Introducing $\omega_i = 1/\rho_i$, we approximate

$$\theta_i \approx \frac{h^{(l)}}{H^{(l)}} \cdot 2\pi + \theta_{\text{off}}^{(l)} + o_x^{(l)} \cdot \omega_i. \quad (27)$$

Since $H^{(l)}$ is an integer, with a minimum value equal to the number of points in scanline l and a practical maximum on the order of 10^3 for LiDAR sensors [46, 47, 48, 49], it is feasible to perform an exhaustive search over candidate resolutions $\hat{H}^{(l)}$. For each candidate, we compute the corresponding values of $\hat{o}_x^{(l)}$, $\hat{\theta}_{\text{off}}^{(l)}$, and a loss metric $L(\hat{H}^{(l)})$. The parameter set that minimizes this loss is then selected. In our implementation, we set the maximum $\hat{H}^{(l)}$ to 10^4 , though this upper bound can be adjusted if future sensors exceed this resolution.

4.2.1. Computation of $\hat{o}_x^{(l)}$, $\hat{\theta}_{\text{off}}^{(l)}$, and $L(\hat{H}^{(l)})$ for a given $\hat{H}^{(l)}$

For each candidate resolution $\hat{H}^{(l)}$, the computation proceeds in three stages: (1) quantify the discrepancy $\Delta\theta_i$ between the observed azimuth and the ideal angle implied by $\hat{H}^{(l)}$ as a function of ω_i , which under the correct resolution forms a piecewise linear function; (2) extract local linear pieces to obtain a prior slope estimate; and (3) refine the offsets by compensating for periodic discontinuities and regressing a continuous linear model whose residuals define the normalized loss $L(\hat{H}^{(l)})$.

For a candidate resolution $\hat{H}^{(l)}$, the ideal azimuths are restricted to the discrete set

$$\left\{ \frac{2\pi}{\hat{H}^{(l)}} \cdot k \mid k \in \{0, 1, \dots, \hat{H}^{(l)} - 1\} \right\}.$$




Figure 8: Flowchart of the horizontal parameter estimation algorithm (Sec. 4.2). The algorithm performs an exhaustive search over candidate horizontal resolutions $\hat{H}^{(l)}$ for each scanline, computing horizontal offset $\hat{o}_x^{(l)}$ and azimuthal offset $\hat{\theta}_{\text{off}}^{(l)}$, and selects the parameters that minimize the loss function. For scanlines with insufficient points, a heuristic fallback reuses parameters from previously processed beams.

Hence each observed angle θ_i can be associated with the nearest element of this set via an integer index k_i , and the discrepancy is defined as

$$\Delta\theta_i = \theta_i - \frac{k_i}{\hat{H}^{(l)}} \cdot 2\pi, \quad k_i = \left\lfloor \frac{\theta_i}{2\pi} \cdot \hat{H}^{(l)} + 0.5 \right\rfloor. \quad (28)$$

Substituting θ_i from Equation (27) into k_i and simplifying we obtain

$$k_i = \left\lfloor \frac{h^{(l)}}{H^{(l)}} \cdot \hat{H}^{(l)} + \frac{\theta_{\text{off}}^{(l)} + o_x^{(l)} \cdot \omega_i}{2\pi} \cdot \hat{H}^{(l)} + 0.5 \right\rfloor.$$

Assuming the candidate resolution $\hat{H}^{(l)}$ is correct ($\hat{H}^{(l)} = H^{(l)}$), we obtain

$$k_i = h^{(l)} + \left\lfloor \frac{\theta_{\text{off}}^{(l)} + o_x^{(l)} \cdot \omega_i}{2\pi} \cdot \hat{H}^{(l)} + 0.5 \right\rfloor,$$

where we have used that $h^{(l)} \in \mathbb{Z}$ and is therefore unaffected by the rounding operation. We express the rounded value as an integer correction term

$$k_i = h^{(l)} + \Delta h_i^{(l)}, \quad \Delta h_i^{(l)} := \left\lfloor \frac{H^{(l)}}{2\pi} (\theta_{\text{off}}^{(l)} + o_x^{(l)} \omega_i) + 0.5 \right\rfloor \in \mathbb{Z}.$$

Since the only non-constant factor in $\Delta h_i^{(l)}$ is $\omega_i > 0$, $\Delta h_i^{(l)}$ is an integer-valued, piecewise-constant function of ω_i :

$$\Delta h_i^{(l)} = f(\omega_i), \quad f: \mathbb{R}_{>0} \rightarrow \mathbb{Z} \text{ piecewise constant.}$$

Substituting k_i into Equation (28) gives

$$\Delta\theta_i = \underbrace{\frac{h^{(l)}}{H^{(l)}} \cdot 2\pi + \theta_{\text{off}}^{(l)} + o_x^{(l)} \cdot \omega_i}_{\theta_i \text{ (from Eq. (27))}} - \frac{h^{(l)} + \Delta h_i^{(l)}}{H^{(l)}} \cdot 2\pi,$$

which simplifies to

$$\Delta\theta_i = -\Delta h_i^{(l)} \cdot \frac{2\pi}{H^{(l)}} + \theta_{\text{off}}^{(l)} + o_x^{(l)} \cdot \omega_i. \quad (29)$$

Thus when $\hat{H}^{(l)} = H^{(l)}$, $\Delta\theta_i$ is a piecewise linear function of ω_i , with slope $o_x^{(l)}$ and periodic intercept of period $T^{(l)} = 2\pi/H^{(l)}$.

We estimate the horizontal offset $o_x^{(l)}$ by exploiting the piecewise linearity of $\Delta\theta_i$. We define two segmentation thresholds, τ_ω and $\tau_{\Delta\theta}$, to partition the $(\omega_i, \Delta\theta_i)$ space into contiguous segments. Consecutive points are assigned to the same segment if the difference in ω_i is smaller than τ_ω and the difference in $\Delta\theta_i$ is smaller than $\tau_{\Delta\theta}$. This procedure isolates regions corresponding to the individual linear portions of the piecewise function. In our implementation, the thresholds are chosen as $\tau_\omega = T^{(l)}/4$ rad (one quarter of the piecewise function's period, ensuring reliable detection of discontinuities) and $\tau_{\Delta\theta} = 10^{-2}\text{m}^{-1}$ (empirically validated for typical LiDAR ranges of 1–100 m).

Within each segment, we perform a linear regression to estimate a local slope. These slope estimates are then aggregated using a weighted median, where each segment contributes proportionally to its number of points. In this way, larger and more reliable segments exert greater influence on the aggregated estimate $\hat{o}_x^{(l)}$. This initial estimate serves as a prior for the subsequent optimization steps. We first define a reference line with zero intercept and slope $\hat{o}_x^{(l)}$, and compute the residuals of each point $(\omega_i, \Delta\theta_i)$ with respect to this line:

$$R_i = \Delta\theta_i - \hat{o}_x^{(l)} \cdot \omega_i = -\Delta h_i^{(l)} \cdot \frac{2\pi}{H^{(l)}} + \theta_{\text{off}}^{(l)}.$$

Because $\Delta h_i^{(l)}$ is integer-valued, it contributes multiples of $2\pi/H^{(l)}$. Hence the residuals are congruent with the azimuthal

offset modulo this period:

$$R_i \equiv \theta_{\text{off}}^{(l)} \pmod{2\pi/H^{(l)}}.$$

The azimuthal offset $\theta_{\text{off}}^{(l)}$ can therefore be estimated by applying a circular mean to the residuals modulo $2\pi/H^{(l)}$:

$$\hat{\theta}_{\text{off}}^{(l)} = \text{CircMean}\left(R_i \bmod \frac{2\pi}{H^{(l)}}\right).$$

At this point we obtain preliminary estimates $\hat{\theta}_x^{(l)}$ and $\hat{\theta}_{\text{off}}^{(l)}$. To refine them, we compute centered residuals by subtracting the estimated offset:

$$R'_i = R_i - \hat{\theta}_{\text{off}}^{(l)} \approx -\Delta h_i^{(l)} \cdot \frac{2\pi}{H^{(l)}}.$$

From these, the correction terms $\Delta h_i^{(l)}$ are estimated as

$$\hat{\Delta h}_i^{(l)} = -\left\lfloor \frac{H^{(l)}}{2\pi} R'_i + 0.5 \right\rfloor.$$

Assuming a correct estimation, we have $\hat{\Delta h}_i^{(l)} = \Delta h_i^{(l)}$. This equality holds under the assumptions of $\hat{H}^{(l)} = H^{(l)}$ and $\hat{\theta}_x^{(l)} \approx \theta_x^{(l)}$. We transform the original piecewise-linear $\Delta\theta_i$ into a continuous linear form by compensating for the periodic discontinuities:

$$\hat{\Delta\theta}_i = \Delta\theta_i + \hat{\Delta h}_i^{(l)} \cdot \frac{2\pi}{H^{(l)}}.$$

Using the definition of $\Delta\theta_i$ from Equation (29), this simplifies to

$$\hat{\Delta\theta}_i = \theta_{\text{off}}^{(l)} + \theta_x^{(l)} \cdot \omega_i. \quad (30)$$

This linear form explicitly shows that $\hat{\Delta\theta}_i$ depends linearly on ω_i , with slope $\theta_x^{(l)}$ and intercept $\theta_{\text{off}}^{(l)}$. Consequently, performing linear regression on the transformed pairs $(\omega_i, \hat{\Delta\theta}_i)$ yields refined estimates $\hat{\theta}_x^{(l)}$ and $\hat{\theta}_{\text{off}}^{(l)}$. The loss metric for the fit is quantified by the squared sum of the residuals,

$$L(\hat{H}^{(l)}) = \sum_i (R'_i)^2 \cdot (\hat{H}^{(l)})^2, \quad (31)$$

where

$$R'_i = \hat{\Delta\theta}_i - \hat{\theta}_{\text{off}}^{(l)} - \hat{\theta}_x^{(l)} \cdot \omega_i.$$

Note that the factor $(\hat{H}^{(l)})^2$ normalizes the loss across candidate resolutions. Since the residuals R'_i are bounded by the period $2\pi/H^{(l)}$, their magnitudes decrease with increasing resolution. Without normalization, higher $\hat{H}^{(l)}$ values would be artificially favored.

4.2.2. Heuristic Fallback

For scanlines containing insufficient points, the estimation procedure described above may be unreliable due to limited data. To address this, such scanlines are deferred until all others have been processed, after which their parameters are estimated using a heuristic fallback strategy.

This strategy reuses information from previously estimated scanlines. Specifically, we collect the values of $H^{(l)}$ and $\theta_x^{(l)}$ obtained from other beams, and for each candidate pair $(\hat{H}^{(l)}, \hat{\theta}_x^{(l)})$ we compute a heuristic loss function $M(\hat{H}^{(l)}, \hat{\theta}_x^{(l)})$. The parameter set minimizing this loss is selected as the final estimate.

We begin by computing corrected angles by removing the effect of the horizontal offset:

$$\theta_i^{\text{corrected}} = \theta_i - \theta_{\text{res}}^{(l, H^{(l)})} \approx \theta_i - \hat{\theta}_x^{(l)} \cdot \omega_i.$$

Substituting in Equation (27), if the candidate $\hat{\theta}_x^{(l)}$ is correct, the corrected angles satisfy

$$\theta_i^{\text{corrected}} \approx \frac{h^{(l)}}{H^{(l)}} \cdot 2\pi + \theta_{\text{off}}^{(l)},$$

so the corresponding ideal quantized angles implied by $\hat{H}^{(l)}$ are:

$$\theta_i^* = \left\lfloor \frac{\theta_i^{\text{corrected}}}{2\pi/\hat{H}^{(l)}} + 0.5 \right\rfloor \cdot \frac{2\pi}{\hat{H}^{(l)}}.$$

For correct parameters, the corrected angles should align with the ideal quantization grid up to a constant $\theta_{\text{off}}^{(l)}$, i.e.,

$$\forall i: \quad \theta_i^* \approx \theta_i^{\text{corrected}} - \theta_{\text{off}}^{(l)}.$$

This observation allows us to estimate the azimuthal offset as the mean deviation between corrected and ideal angles

$$\hat{\theta}_{\text{off}}^{(l)} = \frac{1}{n} \sum_{i=1}^n (\theta_i^{\text{corrected}} - \theta_i^*).$$

Finally, the quality of each candidate parameter set is quantified by the mean absolute deviation

$$M(\hat{H}^{(l)}, \hat{\theta}_x^{(l)}) = \frac{1}{n} \sum_{i=1}^n \left| \theta_i^{\text{corrected}} - \theta_i^* - \hat{\theta}_{\text{off}}^{(l)} \right| \cdot \hat{H}^{(l)}. \quad (32)$$

As in the main estimation procedure, the multiplication by $\hat{H}^{(l)}$ serves as a normalization factor. Since the residual magnitudes shrink with increasing resolution due to the shorter period $2\pi/H^{(l)}$, this scaling ensures that heuristic loss values remain comparable across candidate resolutions.

The candidate pair $(\hat{H}^{(l)}, \hat{\theta}_x^{(l)})$ minimizing M is selected as the final estimate. This fallback provides reasonable and consistent parameter estimates even with limited data.

4.3. Range Image Generation and Point Cloud Reconstruction

Once all intrinsic parameters have been estimated, ALICE-LRI can generate lossless range images from point clouds and subsequently reconstruct the original point clouds without information loss. This bidirectional transformation constitutes the core objective of our approach.

4.3.1. Point Cloud to Range Image

Given estimated intrinsics $\{\varphi^{(l)}, o_y^{(l)}, H^{(l)}, o_x^{(l)}, \theta_{\text{off}}^{(l)}\}_{l=0}^{L-1}$ and a point cloud $\{(x_i, y_i, z_i)\}$, we first compute spherical coordinates $(r_i, \varphi_i, \theta_i)$ as in Equation (3). Subsequently, each point is assigned to the candidate beam $\hat{l}$ whose vertical parameters best explain its elevation angle φ_i :

$$\hat{l}_i = \arg \min_l \left| \varphi_i - \varphi^{(l)} - \arcsin \left(\frac{o_y^{(l)}}{r_i} \right) \right|.$$

The corrected azimuthal angle θ'_i is computed using the second expression in Equation (11):

$$\theta'_i = \theta_i - \arcsin \left(\frac{o_x^{(\hat{l}_i)}}{r_i \cos(\varphi_i)} \right) - \theta_{\text{off}}^{(\hat{l}_i)}$$

The height of the range image is L , and the width is $H = \text{LCM}\{H^{(l)}\}$ (least common multiple) to ensure a uniform grid across all beams even when their resolutions differ. The horizontal (u_i) and vertical (v_i) indices are computed as

$$\begin{cases} u_i = \left\lfloor \frac{\theta'_i}{2\pi} \cdot H + 0.5 \right\rfloor \\ v_i = L - \hat{l}_i - 1 \quad (\text{see Eq. (5)}), \end{cases}$$

where uniqueness of (u_i, v_i) is guaranteed under correct parameters. Finally, each r_i is stored on pixel (u_i, v_i) .

4.3.2. Range Image to Point Cloud

Given estimated intrinsics $\{\varphi^{(l)}, o_y^{(l)}, o_x^{(l)}, \theta_{\text{off}}^{(l)}\}_{l=0}^{L-1}$ and a range image of size $H \times L$, each non-empty pixel (u_i, v_i) with range r_i can be back-projected to a 3D point. First, the beam index is identified as $\hat{l}_i = L - v_i - 1$ and then the original observed angles are recovered by adding back the angular residuals:

$$\begin{cases} \varphi_i = \varphi^{(\hat{l}_i)} + \arcsin \left(\frac{o_y^{(\hat{l}_i)}}{r_i} \right) \\ \theta_i = \frac{2\pi u_i}{H} + \theta_{\text{off}}^{(\hat{l}_i)} + \arcsin \left(\frac{o_x^{(\hat{l}_i)}}{r_i \cos \varphi_i} \right). \end{cases}$$

Finally,

$$x_i = r_i \cos(\varphi_i) \cos(\theta_i), \quad y_i = r_i \cos(\varphi_i) \sin(\theta_i), \quad z_i = r_i \sin(\varphi_i).$$

Under the assumption of correctly estimated intrinsics, this reconstruction is lossless up to floating-point precision.

5. Materials and Methods

This section details the experimental methodology used to evaluate the proposed approach. We describe the datasets employed, the experimental setup, and the metrics used for quantitative and qualitative assessments.

5.1. Datasets

The proposed method was evaluated on two publicly available datasets: KITTI [44] and DurLAR [50]. These datasets were selected to represent two distinct sensor configurations with contrasting characteristics, enabling thorough evaluation of the proposed approach. Their main characteristics are provided in Table 1.

KITTI features point clouds acquired with a Velodyne HDL-64E sensor with 64 scanlines and substantial beam offsets (vertical: 100 mm to 210 mm, horizontal: -26 mm to 26 mm). The dataset presents challenges primarily due to these significant geometric distortions.

DurLAR provides point clouds from an Ouster OS1-128 sensor with 128 scanlines and minimal offsets (vertical: 25 mm to 40 mm, horizontal: -1 mm to 1 mm). Despite the smaller geometric distortions, DurLAR poses distinct challenges due to its higher scanline count and extended vertical field of view ($[-22.5^\circ, 22.5^\circ]$). The wider angular coverage increases the likelihood of missing returns at high elevation angles where ground surfaces are absent, resulting in sparse or missing scanlines and varying point densities.

These complementary characteristics allow comprehensive evaluation of the proposed method: KITTI tests robustness against significant sensor non-idealities and geometric distortions, while DurLAR evaluates performance with compact sensors under challenging data sparsity conditions at extreme viewing angles.

5.2. Experimental Setup

All experiments were conducted on the CESGA Finisterre III [51] supercomputing cluster to efficiently process the entire KITTI and DurLAR datasets, which contain 193,801 point clouds. This allowed us to analyze both datasets in their entirety rather than using only small subsets. For each point cloud, results—including estimated sensor parameters, range images, and evaluation metrics—were stored in local SQLite [52] databases to ensure reproducibility and facilitate downstream analysis. After processing, these databases were aggregated and analyzed offline on a local workstation using custom scripts.

Runtime performance analysis was conducted separately on a local workstation to provide timing measurements under controlled conditions. The workstation specifications include an Intel Core i7-13700K processor and 64 GB of RAM, running Ubuntu 22.04.5 LTS.

5.3. Evaluation Metrics

The evaluation described in Section 6 considers two main aspects: (i) the accuracy of the estimated sensor parameters, and (ii) the quality of the reconstructed point clouds obtained from the range images. In addition, the application presented in Section 7 demonstrates the practical benefits through a compression case study. Each aspect is assessed using appropriate quantitative metrics, as detailed below.

Table 1: Characteristics of the datasets used for evaluation, showing sensor specifications and data distribution.

Dataset	Sequences	Point Clouds	Scanlines	Horizontal Resolution	Vertical Angles	Vertical Offsets	Horizontal Offsets
KITTI [44]	151	47,885	64	4000	$[-24.9^\circ, 2.0^\circ]$	[100 mm, 210 mm]	$[-26 \text{ mm}, 26 \text{ mm}]$
DurLAR [50]	5	145,916	128	2048	$[-22.5^\circ, 22.5^\circ]$	[25 mm, 40 mm]	$[-1 \text{ mm}, 1 \text{ mm}]$

(i) *Sensor Parameter Estimation.* For parameter estimation accuracy, we use standard classification metrics (overall accuracy, precision, recall, F1-score) and mean absolute error (MAE) for continuous parameters. We report both macro-averaged (mP, mR, mF1) and weighted-averaged (wP, wR, wF1) variants of classification metrics. Macro-averaged metrics compute the unweighted mean across all classes, while weighted-averaged metrics weight each class by its support to reflect class imbalance.

(ii) *Point Cloud Reconstruction Quality.* To evaluate the fidelity of reconstructed point clouds, we employ geometric similarity metrics that compare points in the original set P with their counterparts in the reconstruction $\hat{P}$.

- **Chamfer Distance (CD):** The Chamfer Distance quantifies point cloud similarity by measuring bidirectional nearest-neighbour distances. Despite its widespread use in 3D vision, implementations vary considerably in the literature.

Some researchers utilize the non-squared Euclidean distance [53], while others employ squared distances to penalize outliers more severely [54]. Additionally, the directional terms may be either summed [55] or averaged, as in the Point Cloud Utils library [56].

We adopt the non-squared, averaged Chamfer Distance formulation

$$\text{CD}(P, \hat{P}) = \frac{\text{MAE}(P, \hat{P}) + \text{MAE}(\hat{P}, P)}{2},$$

where

$$\text{MAE}(A, B) = \frac{1}{|A|} \sum_{a \in A} \min_{b \in B} \|a - b\|_2.$$

This approach provides an interpretable metric that directly expresses reconstruction error in physical units (meters), representing the average displacement between corresponding points.

- **Peak Signal-to-Noise Ratio (PSNR):** Adapted from image processing, PSNR quantifies reconstruction quality by comparing maximum signal power to error power:

$$\text{PSNR}(P, \hat{P}) = 10 \cdot \log_{10} \left(\frac{\text{MAX}^2}{\text{MSE}(P, \hat{P})} \right),$$

where

$$\text{MSE}(A, B) = \frac{1}{|A|} \sum_{a \in A} \min_{b \in B} \|a - b\|_2^2,$$

and MAX represents the maximum representable range value (that is, 120 m for KITTI [46] and 170 m for DurLAR [47]).

- **Sampling Error (SE):** Point preservation is quantified through sampling error:

$$\text{SE}(P, \hat{P}) = \frac{||P| - |\hat{P}||}{|P|}, \quad 0 \leq \text{SE} \leq 1.$$

This metric measures the relative difference in cardinality between the original and the reconstructed point clouds.

6. Evaluation

This section evaluates ALICE-LRI on the KITTI and DurLAR datasets. Our analysis combines quantitative metrics with qualitative visualizations to assess estimation accuracy, reconstruction fidelity, and computational performance. To ensure reproducibility, we provide all code and experiment configurations publicly.³

We apply the exact same methodology across both datasets without any dataset-specific customization, demonstrating the generality and robustness of ALICE-LRI. We focus on three key aspects of evaluation: (1) parameter estimation accuracy, (2) range image reconstruction quality, and (3) runtime performance for real-time applications. These evaluations collectively demonstrate the technical accuracy of ALICE-LRI and validate its effectiveness for practical use.

6.1. Parameter Estimation

First, we evaluate the accuracy of ALICE-LRI in estimating the intrinsic parameters of spinning LiDAR sensors. We assess the estimation of the number of scanlines L , the vertical angles $\varphi^{(l)}$, spatial offsets $o^{(l)}$, horizontal resolutions $H^{(l)}$, and azimuthal offsets $\theta_{\text{off}}^{(l)}$.

We run ALICE-LRI on all frames in each dataset and compare the estimated parameters against reference values derived from manufacturer specifications and available sensor calibration data. To ensure consistent evaluation, the reference values for $\varphi^{(l)}$, $o^{(l)}$, $H^{(l)}$, and $\theta_{\text{off}}^{(l)}$ are fixed per dataset according to the known sensor configuration. However, the number of scanlines L may vary between frames because some laser beams occasionally fail to produce returns, requiring frame-specific verification. Since calibration metadata is available for both datasets, we developed dataset-specific verification scripts. These scripts

³<https://github.com/alice-lri/alice-lri-experiments>

Table 2: Scanlines count estimation metrics on KITTI and DurLAR datasets. For each subset (all point clouds and those with $n^{(l)} \geq 64$), the number of point clouds (# Samples), incorrect predictions (# Incorrect), overall accuracy (OA), macro and weighted precision (mP, wP), recall (mR, wR), and F1-score (mF1, wF1) are reported.

Dataset	Subset	Scanlines Count								
		# Samples	# Incorrect	OA (%)	mP (%)	mR (%)	mF1 (%)	wP (%)	wR (%)	wF1 (%)
KITTI	All	47,885	0	100.00	100.00	100.00	100.00	100.00	100.00	100.00
	$n^{(l)} \geq 64$	47,543	0	100.00	100.00	100.00	100.00	100.00	100.00	100.00
DurLAR	All	145,916	130	99.91	92.98	93.67	93.31	99.92	99.91	99.92
	$n^{(l)} \geq 64$	130,757	0	100.00	100.00	100.00	100.00	100.00	100.00	100.00

automatically determine the effective scanline count in each frame by checking which of the predefined scanlines receive valid point returns. This procedure ensures an unambiguous, frame-specific reference for L , consistent with the documented sensor configuration.

6.1.1. Scanlines Count

We treat the problem of estimating the number of scanlines L as a multi-class classification task, where each sample is a point cloud and each class corresponds to a different number of scanlines.

Table 2 summarizes the scanline count estimation results on the KITTI and DurLAR datasets. The reported metrics are: number of samples (point clouds), number of incorrect predictions, OA (overall accuracy), mP (macro precision), mR (macro recall), mF1 (macro F1-score), wP (weighted precision), wR (weighted recall), and wF1 (weighted F1-score), all expressed as percentages except for the sample counts. To provide a comprehensive assessment, we partition each dataset into two subsets that highlight performance under different conditions. The “all” subset encompasses every point cloud in the dataset, while the “ $n^{(l)} \geq 64$ ” subset contains only point clouds with at least 64 points per scanline—a reasonable threshold that maintains most of the dataset while excluding severely degraded scans.

The results demonstrate that ALICE-LRI achieves perfect scanline count prediction on KITTI (100% overall accuracy) and near-perfect results on DurLAR (99.91% overall accuracy) when considering all point clouds in the datasets. Perfect classification (100% accuracy) on both datasets is achieved when considering only point clouds with sufficient point density (≥ 64 points per scanline). These results validate the robustness of ALICE-LRI across different sensor configurations and environmental conditions.

The difference between macro metrics (mP, mR, mF1) and their weighted counterparts (wP, wR, wF1) on the DurLAR dataset primarily arises from the relationship between scanline completeness and point density. Point clouds containing all scanlines typically have many points per scanline, making them easier to classify and thus dominant in the weighted metrics. In contrast, point clouds with missing scanlines often contain remaining scanlines with very few points—sometimes only one or two, especially those adjacent to the missing ones. Consequently, these cases are more challenging for the method, and since macro metrics assign equal weight to all classes, they are

more affected by such difficult, low-density instances. Therefore, the observed gap between macro and weighted metrics reflects the underlying data distribution rather than any limitation or bias of the method.

6.1.2. Per-Beam Parameters

We now evaluate the accuracy of the estimated per-beam parameters. For each detected scanline, the estimated horizontal resolutions $\hat{H}^{(l)}$, vertical angles $\hat{\phi}^{(l)}$, spatial offsets $\hat{\delta}^{(l)}$, and azimuthal offsets $\hat{\theta}_{\text{off}}^{(l)}$ are compared against the corresponding reference values. Horizontal resolution estimation is treated as a binary classification problem (considered correct if it matches the reference value, and incorrect otherwise) whereas the remaining parameters are evaluated as continuous quantities.

Table 3 presents the horizontal resolution estimation results across both datasets. The method achieves near-perfect accuracy of 99.99% on both KITTI (84 incorrect) and DurLAR (237 incorrect) when considering all point clouds in the datasets. When focusing on point clouds with sufficient point density ($n^{(l)} \geq 64$), performance improves further, reaching 100% accuracy on DurLAR (0 incorrect) and 99.99% accuracy on KITTI (4 incorrect). These results confirm the robustness of our horizontal resolution estimation approach across different sensor configurations and data quality conditions.

Table 3: Horizontal resolution estimation accuracy on KITTI and DurLAR. We report the total number of scanlines (# Samples), incorrect predictions (# Incorrect), and overall accuracy (OA) for all point clouds and for those with $n^{(l)} \geq 64$.

Dataset	Subset	Horizontal Resolution		
		# Samples	# Incorrect	OA (%)
KITTI	All	3,064,085	84	99.99
	$n^{(l)} \geq 64$	3,042,253	4	99.99
DurLAR	All	18,637,656	237	99.99
	$n^{(l)} \geq 64$	16,736,809	0	100.00

Table 4 summarizes the estimation errors for the remaining per-beam parameters using mean absolute error (MAE) and maximum error (MAX) metrics. The results demonstrate high accuracy, particularly on point clouds with adequate point density ($n^{(l)} \geq 64$). For these well-populated scanlines, vertical

Table 4: Per-beam parameter estimation errors (MAE and MAX) for vertical angles, vertical and horizontal offsets, and azimuthal offsets on KITTI and DurLAR. Results are shown for all point clouds and for those with $n^{(l)} \geq 64$.

Dataset	Subset	Vert. Angle $\phi^{(l)}$ (deg)		Vert. Offset $o_y^{(l)}$ (mm)		Horiz. Offset $o_x^{(l)}$ (mm)		Azim. Offset $\theta_{\text{off}}^{(l)}$ (deg)	
		MAX	MAE	MAX	MAE	MAX	MAE	MAX	MAE
KITTI	All	0.123479	0.000413	137.338350	0.057363	202.204798	0.040357	0.081500	0.000872
	$n^{(l)} \geq 64$	0.049864	0.000412	4.006125	0.056994	19.806281	0.038512	0.081500	0.000870
DurLAR	All	3.490045	0.000033	323.204766	0.005478	17.107649	0.003304	0.174496	0.000032
	$n^{(l)} \geq 64$	0.012686	0.000008	0.102721	0.000780	0.010094	0.002778	0.000117	0.000028

angle estimation achieves sub-degree accuracy, with MAE values of $4.12 \cdot 10^{-4}^\circ$ on KITTI and $8 \cdot 10^{-6}^\circ$ on DurLAR, and maximum errors below 0.050° and 0.013° respectively. Spatial offset estimation shows remarkable sub-millimeter precision, with MAE values of $5.7 \cdot 10^{-2}$ mm (vertical) and $3.9 \cdot 10^{-2}$ mm (horizontal) on KITTI, and $7.8 \cdot 10^{-4}$ mm and $2.8 \cdot 10^{-3}$ mm on DurLAR. The corresponding maximum errors are well-contained at 4.0 mm and 19.8 mm for KITTI, and 0.10 mm and 0.01 mm for DurLAR. The high maximum horizontal offset error of 19.8 mm on KITTI stems from cases with incorrect resolution estimation, as resolution and horizontal offset estimation are interdependent; it reduces to 6.3 mm when considering only correct resolution estimates. Azimuthal offset estimation is also highly accurate, with MAE values of $8.7 \cdot 10^{-4}^\circ$ on KITTI and $2.8 \cdot 10^{-5}^\circ$ on DurLAR, and maximum errors of 0.082° and $1.2 \cdot 10^{-4}^\circ$ respectively.

When considering all point clouds, MAE values remain low, indicating that the higher maximum errors are caused by a small number of outliers in sparse scans rather than systematic failure. The primary sources of error are insufficient point density, which degrades statistical estimation, and incorrect point-to-scanline assignments in severely degraded scans. Nevertheless, the method demonstrates robust performance on well-populated scanlines, confirming its effectiveness for practical use.

6.1.3. Ablation Study

ALICE-LRI incorporates several algorithmic components that vary in mathematical rigor: while core elements like the Hough Transform and weighted least squares fitting have strong theoretical foundations, other components such as the heuristic fallbacks rely on empirical assumptions about LiDAR sensor design. To assess the contribution of each component, we conduct a comprehensive ablation study over the entire KITTI and DurLAR datasets by selectively disabling individual algorithm features. This analysis is crucial for determining whether performance gains arise from principled mathematical modeling or from heuristics that may implicitly encode sensor-specific knowledge.

We focus our ablation analysis on scanline count estimation and horizontal resolution accuracy as they represent the most critical parameters for the internal consistency of the algorithm and parameter interdependencies. Scanline count errors cascade through the entire parameter estimation pipeline, leading to incorrect point-to-scanline assignments that compromise

subsequent parameter fitting stages. Similarly, horizontal resolution accuracy is fundamental for proper horizontal and azimuthal offsets estimation, as these parameters are intrinsically coupled in the fitting process—incorrect resolution estimates prevent accurate spatial offset determination. These parameters also exhibit the clearest binary success/failure characteristics, making them ideal metrics for assessing component contributions. While other parameters such as vertical angles and spatial offsets are equally important for reconstruction quality, their continuous nature makes ablation analysis more complex and less interpretable in terms of algorithmic component impact.

We evaluate four key components: (1) *Hough Continuity*, which implements the vote-for-discontinuities strategy to improve robustness to steep gradients in the parameter space; (2) *Conflict Resolution*, which enforces global consistency through backtracking and prevents inconsistent scanline assignments; (3) *Vertical Heuristics*, which provide fallback parameter estimation when weighted least squares fitting fails due to insufficient data; and (4) *Horizontal Heuristics*, which estimate horizontal parameters for sparse scanlines using values from other beams. Note that, while the first two components enhance the core mathematical framework, the latter two represent domain-specific heuristics that exploit typical LiDAR sensor configurations.

When *Hough Continuity* is disabled, the accumulator uses a discontinuous voting scheme as shown in Figure 5a. Disabling *Conflict Resolution* reduces the algorithm to a greedy approach where scanlines are immediately accepted upon successful fitting unless they contain already assigned points, in which case they are permanently discarded. When *Vertical Heuristics* are disabled, scanlines that fail weighted least squares estimation due to insufficient data are entirely discarded. Finally, disabling *Horizontal Heuristics* enforces the standard fitting procedure across all scanlines irrespective of point density. Scanlines that fail to converge receive invalid parameter values. The computational overhead of these individual components is negligible compared to the total parameter estimation time, allowing us to focus solely on their algorithmic contributions without performance considerations.

Table 5 summarizes the features enabled in each experiment configuration (E0-E7), along with their impact on scanline count estimation and resolution accuracy across both datasets. Configuration E0 represents the complete method with all components active, while subsequent configurations (E1-E7) sys-

Table 5: Ablation study results showing the impact of different algorithm components on estimation accuracy. For each configuration, we report the number of incorrect scanline counts and resolution errors across KITTI and DurLAR datasets. Values not in parentheses represent all point clouds, while values in parentheses represent point clouds with $n^{(l)} \geq 64$. The full method (E0) includes all components, while subsequent configurations systematically disable specific features.

	Algorithm Components				KITTI		DurLAR	
	Hough Continuity	Conflict Resolution	Vertical Heuristics	Horizontal Heuristics	# Incorrect Scanlines Count	# Incorrect Resolutions	# Incorrect Scanlines Count	# Incorrect Resolutions
E0	✓	✓	✓	✓	0 (0)	84 (4)	130 (0)	237 (0)
E1		✓	✓	✓	0 (0)	84 (4)	145 (0)	260 (0)
E2	✓		✓	✓	2 (0)	86 (4)	2110 (24)	3263 (154)
E3	✓	✓		✓	136 (3)	414 (7)	4137 (0)	13741 (0)
E4	✓	✓	✓		0 (0)	299 (4)	130 (0)	12807 (0)
E5	✓	✓			136 (3)	455 (7)	4137 (0)	14204 (0)
E6			✓	✓	5 (2)	90 (5)	2294 (7)	2734 (3)
E7					140 (5)	461 (9)	5666 (7)	17694 (4)

tematically disable specific features to isolate their individual contributions. Values not in parentheses represent all point clouds in the dataset, while values in parentheses represent point clouds with $n^{(l)} \geq 64$ points per scanline. This experimental design allows us to distinguish between improvements achieved through rigorous mathematical methods versus those dependent on heuristic assumptions, thereby validating the fundamental soundness of our approach.

The ablation results show that the full method (E0) achieves the best performance, with errors increasing as key components are disabled. Removing core mathematical elements or heuristics leads to substantial drops in accuracy, especially for sparse data. This effect is more pronounced on DurLAR, which presents greater challenges than KITTI. Crucially, disabling heuristics has a very limited impact on dense point clouds in both datasets: even when both vertical and horizontal heuristics are removed, performance on dense data remains nearly perfect.

On DurLAR, configuration E6 (disabling both Hough continuity and conflict resolution) outperforms E2 (disabling only conflict resolution) on dense data. This counterintuitive result stems from specific consecutive frames where Hough continuity creates false accumulator peaks. While conflict resolution can address these issues when enabled, its absence in E2 causes failures that are mitigated in E6 by also disabling the problematic Hough continuity component.

The results reveal a clear pattern: the core mathematical framework demonstrates robust performance on well-populated scanlines, while heuristic components serve as essential fallbacks for sparse data scenarios. The stark contrast between performance on all point clouds versus dense subsets when no heuristics are used confirms that estimation accuracy fundamentally depends on point density. When sufficient data is available, the underlying mathematical formulation operates reliably even without domain-specific optimizations, demonstrating the robustness of our approach.

6.2. Range Image Error Analysis

We now turn to the core objective that motivates this work: assessing whether our inferred sensor parameters enable more

accurate range image generation compared to standard approaches. While the parameter estimation accuracy demonstrated in the previous subsection validates the technical correctness of our inference algorithm, the true test lies in whether these parameters translate into measurably improved range image quality and lossless reconstruction capability. This evaluation directly addresses the central promise of our approach: enabling lossless range image generation from calibrated point clouds without requiring manufacturer metadata.

For each sequence in the datasets, we first apply our algorithm to the initial point cloud to infer the sensor parameters, then reuse these parameters consistently for all subsequent point clouds in the sequence. This mirrors typical real-world deployment scenarios where practitioners encounter an unknown sensor and must establish its configuration for subsequent lossless range image generation.

As the baseline, we employ the Projection-By-Elevation-Angle (PBEA) method described by Wu et al. [15]. PBEA assumes the ideal model described in Section 3.2 and operates by uniformly sampling the elevation angle space between predefined minimum and maximum values, creating an equidistant grid in the vertical dimension while maintaining uniform azimuthal sampling in the horizontal dimension. This approach requires only the specification of vertical and horizontal range image resolutions along with elevation angle bounds. It is the most widely adopted method in practice due to its simplicity and lack of sensor-specific requirements. Furthermore, Wu et al. demonstrated that PBEA outperforms alternative approaches such as Projection-By-Laser-ID (PBID) when sufficient resolution is available, establishing it as a strong and fair baseline for our comparison.

We evaluate range image generation quality using both quantitative metrics and qualitative visual analysis. The quantitative analysis measures reconstruction accuracy through objective metrics, while the qualitative analysis examines presence of artifacts and structural preservation to assess practical performance.

6.2.1. Quantitative Analysis

To assess reconstruction quality, we project each point cloud to a range image using our method, then unproject it back to 3D space and compute three key metrics against the original point cloud: Chamfer Distance (CD) and Peak Signal-to-Noise Ratio (PSNR) for geometric fidelity, and sampling error (SE) representing the fraction of points lost during the projection-unprojection process.

Table 6 presents a quantitative comparison between our proposed method and the PBEA baseline across multiple range image resolutions, evaluated on the complete datasets without excluding any frames. Traditional range image generation methods like PBEA require practitioners to manually select appropriate resolutions through trial and error, as the optimal resolution depends on unknown sensor characteristics. To provide a fair comparison, we evaluate PBEA using the native resolution of each sensor (4000×64 for KITTI’s HDL-64E and 2048×128 for DurLAR’s Ouster OS1-128).

At the native resolution, our results show that ALICE-LRI significantly outperforms PBEA across all metrics. On KITTI at 4000×64 , PBEA achieves 0.027 m average Chamfer Distance with 8.69% sampling error, while ALICE-LRI delivers $3.80 \cdot 10^{-4}$ m CD with 0% sampling error. On DurLAR at 2048×128 , PBEA shows 0.024 m average CD and 3.52% sampling error compared to ALICE-LRI’s $1 \cdot 10^{-6}$ m CD and 0% sampling error. Crucially, for ALICE-LRI, the maximum sampling error is exactly zero across the entire KITTI and DurLAR datasets, that is, no points are ever lost during projection and unprojection.

A common approach to mitigate reconstruction loss in range image projections involves increasing the resolution until the desired quality is achieved. As demonstrated in Table 6, higher resolutions indeed improve PBEA reconstruction quality. However, even with substantial resolution increases, the performance gap relative to ALICE-LRI persists. Scaling PBEA resolution by a factor of 32 in each dimension—from 4000×64 to 128000×2048 for KITTI and from 2048×128 to 65536×4096 for DurLAR (representing a 1024-fold increase in total pixels)—reduces the Chamfer Distance to the millimeter/sub-millimeter scale. Nevertheless, this approach still yields non-zero sampling error. This demonstrates that even significant increases in resolution cannot fully address the fundamental limitations of the PBEA approach.

Figure 9 complements the quantitative results presented in Table 6 by providing frame-by-frame analysis of reconstruction quality across representative sequences. The plots reveal the consistent performance advantage of ALICE-LRI over the PBEA baseline, with ALICE-LRI (solid lines) maintaining sub-millimeter Chamfer Distance values across all frames. In contrast, the PBEA baseline (dashed lines) exhibits significantly higher reconstruction errors, even at increased resolutions. This frame-level analysis confirms that ALICE-LRI’s superior performance is not driven by outliers but represents consistent improvement across diverse scanning conditions and scene contents.

Overall, the quantitative analysis confirms that ALICE-LRI

achieves lossless geometric reconstruction. No points are lost in any frame, and the Chamfer Distance remains below the intrinsic noise level of each LiDAR sensor, surpassing their physical precision.

6.2.2. Qualitative Analysis

The quantitative metrics presented above provide a comprehensive assessment of reconstruction quality, but visual inspection offers complementary insights into the practical implications of using ALICE-LRI. Figure 10 shows the differences between the original point cloud and the reconstructions obtained using both the PBEA baseline and ALICE-LRI. The visualizations reveal that while the baseline method introduces geometric distortions and point cloud artifacts, ALICE-LRI preserves the original geometric structure with high fidelity.

ALICE-LRI not only significantly improves 3D point cloud reconstruction but also enhances the quality of range image representations. Figure 11 provides a side-by-side comparison of range images generated using both the baseline PBEA and ALICE-LRI. The baseline PBEA method produces range images with missing data regions (highlighted areas). These artifacts arise from the mismatch between the idealized spherical projection model and the actual sensor geometry. In contrast, ALICE-LRI generates range images with smooth, continuous patterns that accurately reflect the underlying sensor scanning behavior. The absence of holes and irregular discontinuities in range images directly correlates with the zero sampling error achieved in our quantitative evaluation, confirming the lossless nature of our approach. This improvement in range image quality has implications beyond reconstruction fidelity: cleaner range images without artifacts can potentially benefit downstream applications such as semantic segmentation, object detection, and scene understanding, though such evaluations lie outside the scope of this work.

6.3. Runtime Performance


Range images often serve as critical intermediate representations in numerous real-time LiDAR processing applications due to their regular grid structure, which enables efficient computation. Since efficiency is critical in such contexts, we now evaluate the runtime performance of ALICE-LRI.

Runtime measurements were obtained by processing representative frames from each dataset on the local workstation described in Section 5.2. For KITTI, we used one frame from each of 10 different sequences, while for DurLAR, we used two frames from each of the 5 available sequences. This selection ensures variability across different scene conditions. The implementation operates in single-threaded mode and timing measurements exclude all I/O operations.

Table 7 presents the runtime analysis across both datasets. Parameter estimation requires an average of 31.3 seconds for KITTI and 41.3 seconds for DurLAR. Importantly, the parameter estimation phase is performed only once per sensor, making an execution time below one minute highly acceptable for practical deployment scenarios. For the per-frame operations, KITTI projection requires an average of 9.3 ms per frame, while

Table 6: Point cloud reconstruction quality metrics comparing ALICE-LRI against PBEA baseline at native sensor resolutions and progressively increased resolutions. Results are computed over the complete datasets without frame exclusion. For each method and resolution, average (AVG) and worst-case (MAX/MIN) values for Chamfer Distance (CD), Peak Signal-to-Noise Ratio (PSNR), and Sampling Error (SE) are reported.

Dataset	Method	CD (m)		PSNR (dB)		SE (%)	
		AVG	MAX	AVG	MIN	AVG	MAX
KITTI	PBEA (4000×64)	0.027419	0.051127	63.056212	45.769613	8.689915	17.652985
	PBEA (8000×128)	0.012521	0.022010	72.296912	51.615480	0.518697	3.034373
	PBEA ($16,000 \times 256$)	0.006156	0.010768	79.090223	52.142493	0.126372	0.755860
	PBEA ($32,000 \times 512$)	0.003055	0.005281	85.929010	52.152183	0.034610	0.196997
	PBEA ($64,000 \times 1024$)	0.001520	0.002664	92.654247	61.124456	0.009202	0.091909
	PBEA ($128,000 \times 2048$)	0.000758	0.001335	99.450472	63.225820	0.002254	0.033769
	ALICE-LRI (4000×64)	0.000380	0.000423	109.331016	108.196745	0.000000	0.000000
DurLAR	PBEA (2048×128)	0.023616	0.050328	70.359540	38.323632	3.518374	11.450134
	PBEA (4096×256)	0.010813	0.021681	79.518543	39.413287	0.303959	1.271056
	PBEA (8192×512)	0.005354	0.010646	86.099955	48.081959	0.083244	0.525120
	PBEA ($16,384 \times 1024$)	0.002657	0.005335	92.606193	49.798801	0.013953	0.193679
	PBEA ($32,768 \times 2048$)	0.001324	0.002615	98.833528	50.249024	0.003500	0.084627
	PBEA ($65,536 \times 4096$)	0.000662	0.001458	104.882723	50.249114	0.001278	0.038590
	ALICE-LRI (2048×128)	0.000001	0.000006	157.268976	140.288547	0.000000	0.000000



(a) Chamfer Distance performance across the first 1000 point clouds from the 2011_09_30_drive_0018 sequence of the KITTI dataset.

(b) Chamfer Distance performance across the first 1000 point clouds from the DurLAR_20211209 sequence.

Figure 9: Chamfer Distance of ALICE-LRI and PBEA across different range image resolutions. ALICE-LRI (solid lines) consistently achieves lower reconstruction error compared to PBEA (dashed lines). Results show frame-by-frame reconstruction quality, where lower values indicate better geometric fidelity.




Figure 10: Qualitative comparison of 3D point cloud reconstruction quality. The highlighted regions (red boxes) illustrate the reconstruction fidelity achieved by ALICE-LRI compared to the baseline PBEA approach. ALICE-LRI maintains structural integrity while the baseline introduces visible artifacts and geometric distortions.




Figure 11: Comparison of range image quality between baseline and ALICE-LRI methods. The baseline method (top) exhibits visible discontinuities and missing data regions, while ALICE-LRI (bottom) produces clean, artifact-free range images that accurately represent the sensor scanning pattern. The highlighted regions show the improved continuity and data preservation achieved by ALICE-LRI.

unprojection completes in 3.8 ms. The DurLAR dataset exhibits slightly longer processing times due to its higher scanline count: 20.0 ms for projection and 4.8 ms for unprojection.

Table 7: Runtime performance analysis showing mean execution times ($\pm$ standard deviation) for parameter estimation, projection, and unprojection phases on the KITTI and DurLAR datasets. Parameter estimation times represent one-time offline estimation costs, while projection and unprojection times indicate per-frame processing requirements for real-time applications.

Dataset	Estimation (s)	Proj. (ms)	Unproj. (ms)
KITTI	31.3 ± 0.8	9.3 ± 0.3	3.8 ± 0.1
DurLAR	41.3 ± 2.8	20.0 ± 2.8	4.8 ± 0.6

For real-time applications, in the worst case, both the projection and unprojection phases must be executed for every point cloud frame. ALICE-LRI’s combined projection and unprojection times are 13.1 ms for KITTI and 24.8 ms for DurLAR. Considering that both datasets employ sensors operating at 10 Hz, corresponding to a 100 ms frame period, these times represent only 13.1% and 24.8% of the available budget. This demonstrates that ALICE-LRI comfortably meets real-time constraints while leaving substantial computational headroom for downstream processing.

7. Application

We now examine the practical implications of ALICE-LRI for downstream applications. To illustrate the real-world utility of the proposed approach, we present a case study on point cloud compression using the Real-Time Spatio-Temporal Point Cloud Compression (RTST) [29] algorithm, which relies on range images as intermediate representations. We selected RTST because it meets two key criteria: (1) it uses range images as an intermediate representation for compression, and (2) it provides open-source code with clearly identifiable projection components that can be replaced with ALICE-LRI. This allows for a direct comparison of the impact of ALICE-LRI on compression performance while keeping all other algorithmic components unchanged.

To ensure a fair comparison, we evaluate the original RTST implementation against a modified version in which only the range image projection and unprojection components are replaced with ALICE-LRI. Although the RTST algorithm is general and applicable to different sensor configurations, the publicly released code by its authors is tailored specifically for the KITTI dataset. To maintain experimental integrity, we chose not to modify the original implementation beyond the projection components. Therefore, this compression evaluation is conducted exclusively on the KITTI dataset.

The RTST algorithm includes an error threshold parameter that controls compression aggressiveness, enabling us to evaluate performance across different compression-quality trade-offs. We compute aggregate metrics over the KITTI dataset, measuring Compression Ratio (CR), Chamfer Distance (CD), Peak Signal-to-Noise Ratio (PSNR), and Sampling Error (SE). The Compression Ratio is defined as:

$$CR = \frac{\text{Original Size}}{\text{Compressed Size}}$$

Table 8 presents the results of this study. While compression ratios show mixed results, the point cloud quality metrics demonstrate clear advantages for the modified version leveraging ALICE-LRI. At low error thresholds (0.001 to 0.100), where reconstruction loss is primarily attributed to projection-unprojection rather than compression artifacts, ALICE-LRI achieves substantial improvements: CD reduces from 0.0293 m to 0.0020 m at threshold 0.001, PSNR increases from 62.43 dB to 93.82 dB, and SE decreases from 10.09 % to 0.0009 %. As error thresholds increase, the performance gap narrows for CD and PSNR since compression artifacts dominate over projection losses. Still, the SE difference remains substantial across all tested configurations, suggesting that most point losses originate from the range image projection-unprojection process rather than the compression algorithm itself.

While the compression ratio results might suggest that ALICE-LRI only improves point cloud quality without enhancing compression efficiency, this interpretation overlooks the practical implications for users. In real-world scenarios, practitioners typically prioritize achieving specific compression ratios while maintaining acceptable point cloud quality rather than optimizing for particular error threshold values. By achieving superior point cloud quality at equivalent compression ratios, ALICE-LRI enables users to select more aggressive compression settings without experiencing significant quality degradation.

Figure 12 illustrates this relationship by plotting Compression Ratio against Chamfer Distance for both methods across all evaluated error thresholds. Each point represents one execution of the compression algorithm for either the original or the ALICE-LRI modified method. Linear regression analysis reveals that for a given Chamfer Distance, the modified method achieves higher compression ratios on average compared to the baseline. This demonstrates that users achieve better compression efficiency without sacrificing point cloud quality by leveraging ALICE-LRI’s range image generation approach.

These results validate that ALICE-LRI not only enhances range image quality as an intermediate representation but also provides tangible benefits for downstream applications. In the compression domain, ALICE-LRI delivers improvements in both compression efficiency and point cloud fidelity, demonstrating the practical value of accurate sensor parameter estimation for real-world LiDAR processing workflows.

7.1. Runtime Performance Analysis

Beyond quality improvements, a critical concern for practical deployment is whether ALICE-LRI maintains the computational efficiency required for real-time applications. Section 6.3 demonstrated substantial computational headroom for ALICE-LRI’s projection and unprojection operations. To further validate real-time viability, we analyze the runtime performance impact of integrating ALICE-LRI into the RTST compression workflow.

Table 8: Point cloud compression performance comparison between the original RTST algorithm and the ALICE-LRI modified version. Results show Compression Ratio, Chamfer Distance (CD), Peak Signal-to-Noise Ratio (PSNR), and Sampling Error (SE) across different error thresholds. Bold values indicate superior performance for each metric.

Error Threshold	Compression Ratio		CD (m)		PSNR (dB)		SE (%)	
	Original	Modified	Original	Modified	Original	Modified	Original	Modified
0.001	10.7899	10.3009	0.0293	0.0020	62.4274	93.8161	10.0882	0.0009
0.010	10.4894	10.1392	0.0293	0.0021	62.4270	93.5504	10.0881	0.0009
0.050	12.2672	11.8046	0.0315	0.0059	62.3669	81.2060	10.0877	0.0008
0.100	15.0507	14.9141	0.0362	0.0130	62.1628	74.8480	10.0873	0.0007
0.250	19.5483	21.3198	0.0474	0.0301	61.4164	68.2742	10.0869	0.0005
0.500	23.4855	26.6180	0.0627	0.0478	60.2123	64.5307	10.0867	0.0004
0.750	26.4736	30.3682	0.0760	0.0624	59.1101	62.3571	10.0866	0.0004
1.000	28.7996	33.5210	0.0878	0.0754	58.2241	60.8477	10.0865	0.0003




Figure 12: Compression Ratio versus Chamfer Distance trade-off analysis. Each point represents one execution of the compression algorithm. Blue circles denote the original method and red triangles denote the ALICE-LRI modified method. The blue solid and red dashed lines show the corresponding linear regression fits, demonstrating that the ALICE-LRI modified method achieves higher compression ratios for equivalent reconstruction quality.

Table 9 presents a detailed comparison of execution times between the original RTST algorithm and our ALICE-LRI modified version across different error thresholds. The results show that ALICE-LRI introduces a modest computational overhead of 2-8 ms for encoding and 2-3 ms for decoding operations. Even at the highest error thresholds, where compression times are longest, the total execution times remain well within the computational headroom available for real-time applications, confirming that practitioners can adopt ALICE-LRI without sacrificing temporal performance in time-critical applications.

The minimal runtime overhead observed validates the practical feasibility of ALICE-LRI integration in production systems where both quality and performance are critical requirements.

8. Conclusion

This work addressed a fundamental limitation in LiDAR-based perception systems: the inability to generate lossless

range images from calibrated point clouds without sensor-specific metadata. We presented ALICE-LRI (Automatic LiDAR Intrinsic Calibration Estimation for Lossless Range Images), which automatically infers key geometric parameters of spinning LiDAR sensors and enables truly lossless range image generation.

Experimental evaluation on KITTI and DurLAR datasets demonstrated ALICE-LRI effectiveness across different sensor configurations. Critically, ALICE-LRI achieved zero point loss with geometric accuracy orders of magnitude below sensor precision during projection, while the baseline PBEA method exhibited substantial geometric errors (0.027 m Chamfer Distance on KITTI) and significant sampling losses (8.69% on KITTI).

The algorithm demonstrated robustness through its iterative approach combining Hough Transform with weighted least squares fitting and backtracking-based conflict resolution. Comprehensive ablation studies revealed that the core mathematical framework performs reliably on well-populated data even without domain-specific heuristics, thereby validating the fundamental soundness of the approach. Beyond accuracy, ALICE-LRI maintained computational efficiency suitable for real-time applications, with projection and unprojection operations completing in milliseconds and achieving substantial computational headroom. In practical applications, ALICE-LRI enabled better compression efficiency without sacrificing point cloud quality, achieving higher compression ratios at equivalent reconstruction fidelity.

ALICE-LRI’s sensor-agnostic design enables deployment across diverse LiDAR platforms without any manual tuning or manufacturer metadata. This establishes a new paradigm prioritizing geometric fidelity over computational simplicity, addressing a critical gap in 3D perception systems. The open-source implementation facilitates adoption for applications demanding high-precision geometric modeling.

The implications extend beyond range image generation to any application requiring accurate geometric modeling of LiDAR data. As autonomous vehicles, robotics, and mapping applications increasingly demand high-fidelity 3D perception, ALICE-LRI provides a foundation for more reliable and accurate LiDAR processing pipelines. The ability to achieve loss-

Table 9: Mean runtime per frame for the original RTST algorithm and the ALICE-LRI modified version across different error thresholds. Execution times are reported in milliseconds. The minimal overhead introduced by ALICE-LRI ensures real-time viability while providing substantial quality improvements.

Error Threshold	Encoding Time (ms)			Decoding Time (ms)		
	Original	Modified	Overhead	Original	Modified	Overhead
0.001	20.79	23.24	2.45	3.84	6.70	2.86
0.010	19.71	26.00	6.29	4.08	6.78	2.70
0.050	22.57	28.70	6.13	4.85	7.34	2.49
0.100	30.25	37.67	7.43	4.05	7.55	3.51
0.250	43.27	49.89	6.63	4.47	7.52	3.05
0.500	56.98	65.27	8.30	4.03	7.50	3.46
0.750	67.08	72.41	5.33	4.83	7.50	2.66
1.000	75.18	80.06	4.88	4.17	7.32	3.14

less range image generation without manufacturer metadata removes a significant barrier to deploying LiDAR-based systems across diverse sensor platforms and application domains.

9. Future Work

Several promising research directions emerge from our findings. Below, we outline key areas for future investigation that could further advance LiDAR-based perception and geometric modeling.

Our current approach focuses on intrinsic sensor parameters—geometric distortions arising from the physical design and layout of the laser beams. However, in datasets such as KITTI odometry, the ideal spherical projection model is affected not only by intrinsic factors but also by extrinsic corrections applied to compensate for ego-motion during data acquisition. When LiDAR sensors are mounted on moving platforms (e.g., vehicles or robots), point clouds are often motion-corrected to account for sensor displacement and rotation during scanning. These corrections introduce additional geometric transformations that deviate from the static sensor model. Future work could extend our framework to simultaneously estimate intrinsic sensor parameters and extrinsic motion effects. While access to uncorrected raw data or known extrinsic parameters is often available in practice, developing methods that directly account for ego-motion corrections would increase the generality of our approach. This would enable lossless range image generation from a broader class of datasets.

Beyond compression, the impact of lossless range images on downstream perception tasks needs further investigation. Evaluating their effects on semantic segmentation, object detection, and scene understanding could quantify the benefits of eliminating projection artifacts and preserving geometric fidelity. The inferred sensor parameters also open the door to sensor-aware point cloud upsampling and reconstruction. Once sensor geometry is accurately characterized, additional points can be generated in a manner consistent with the physical scanning pattern and measurement characteristics. This capability would be particularly valuable for enhancing sparse automotive LiDAR data and for generating training datasets that accurately

reflect sensor-specific properties for machine learning applications.

Finally, even though our current implementation demonstrates real-time viability, it remains single-threaded. Exploiting the inherently parallel nature of per-point projections through multi-threading and GPU acceleration could dramatically reduce execution times, freeing computational resources for downstream tasks and enabling deployment in high-throughput systems.

Acknowledgments

The authors acknowledge CESGA (Centro de Supercomputación de Galicia) for providing access to the Finisterrae-III supercomputer, which enabled the computational resources necessary for this research. This research was funded by the Agencia Estatal de Investigación (Spain) (MCIN/AEI/10.13039/501100011033) codes PID2019-104834GB-I00, PID2022-141623NB-I00, and PREP2022-000375, the Xunta de Galicia - Consellería de Cultura, Educación, Formación Profesional e Universidades (Centro de investigación de Galicia accreditation 2024-2027 ED431G-2023/04 and Reference Competitive Group accreditation ED431C-2022/016), the European Union (European Regional Development Fund - ERDF/EU).

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used *ChatGPT* in order to assist with drafting, rephrasing, and improving the clarity and structure of the text. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

- [1] N. Haala, et al., Mobile LiDAR mapping for 3D point cloud collection in urban areas—a performance test, in: International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, volume 37, 2008, pp. 1119–1127.

- [2] J. Zhang, S. Singh, LOAM: LiDAR odometry and mapping in real-time, in: *Robotics: Science and Systems (RSS)*, volume 2, 2014.
- [3] X. Chen, I. Vizzo, T. Labe, J. Behley, C. Stachniss, Range image-based LiDAR localization for autonomous vehicles, in: *Proceedings of the 2021 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2021, pp. 5802–5808.
- [4] J. Lichtenfeld, K. Daun, O. von Stryk, Efficient dynamic LiDAR odometry for mobile robots with structured point clouds, in: *Proceedings of the 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2024, pp. 10137–10144.
- [5] J. Behley, C. Stachniss, Efficient surfel-based SLAM using 3D laser range data in urban environments, in: *Proceedings of Robotics: Science and Systems*, Pittsburgh, Pennsylvania, 2018. doi:10.15607/RSS.2018.XIV.016.
- [6] Y. Li, J. Ibanez-Guzman, LiDAR for autonomous driving: The principles, challenges, and trends for automotive LiDAR and perception systems, *IEEE Signal Processing Magazine* 37 (2020) 50–61.
- [7] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, J. Gall, SemanticKITTI: A dataset for semantic scene understanding of LiDAR sequences, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [8] A. Milioto, I. Vizzo, J. Behley, 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.
- [9] U. Weiss, P. Biber, Plant detection and mapping for agricultural robots using a 3D LiDAR sensor, *Robotics and Autonomous Systems* 59 (2011) 265–273.
- [10] D. Hutabarat, et al., LiDAR-based obstacle avoidance for the autonomous mobile robot, in: *2019 12th International Conference on Information & Communication Technology and System (ICTS)*, IEEE, 2019, pp. 137–142.
- [11] G. P. Meyer, A. Laddha, E. Kee, C. Vallespi-Gonzalez, C. K. Wellington, LaserNet: An efficient probabilistic 3D object detector for autonomous driving, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12677–12686.
- [12] B. Wu, A. Wan, X. Yue, K. Keutzer, SqueezeSeg: Convolutional neural nets with recurrent CRF for real-time road-object segmentation from 3D LiDAR point cloud, in: *2018 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2018, pp. 1887–1893.
- [13] T. Shan, J. Wang, F. Chen, P. Szenher, B. Englot, Simulation-based LiDAR super-resolution for ground vehicles, *Robotics and Autonomous Systems* 134 (2020) 103647.
- [14] X. Zhou, C. R. Qi, Y. Zhou, D. Anguelov, RIDDLE: LiDAR data compression with range image deep delta encoding, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 17212–17221.
- [15] T. Wu, H. Fu, B. Liu, H. Xue, R. Ren, Z. Tu, Detailed analysis on generating the range image for LiDAR point cloud processing, *Electronics* 10 (2021) 1224.
- [16] W. Dong, K. Ryu, M. Kaess, J. Park, Revisiting LiDAR registration and reconstruction: A range image perspective, *arXiv preprint arXiv:2112.02779* (2021).
- [17] I. Bogoslavskyi, C. Stachniss, Efficient online segmentation for sparse 3D laser scans, *PFG–Journal of Photogrammetry, Remote Sensing and Geoinformation Science* 85 (2017) 41–52.
- [18] L. T. Triess, D. Peter, C. B. Rist, J. M. Zollner, Scan-based semantic segmentation of LiDAR point clouds: An experimental study, in: *2020 IEEE Intelligent Vehicles Symposium (IV)*, IEEE, 2020, pp. 1116–1121.
- [19] L. Kong, Y. Liu, R. Chen, Y. Ma, X. Zhu, Y. Li, Y. Hou, Y. Qiao, Z. Liu, Rethinking range view representation for LiDAR segmentation, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 228–240.
- [20] B. Li, T. Zhang, T. Xia, Vehicle detection from 3D LiDAR using fully convolutional network, *arXiv preprint arXiv:1608.07916* (2016).
- [21] X. Chen, H. Ma, J. Wan, B. Li, T. Xia, Multi-view 3D object detection network for autonomous driving, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1907–1915.
- [22] J. Zhou, X. Tan, Z. Shao, L. Ma, FVNet: 3D front-view proposal generation for real-time object detection from point clouds, in: *2019 12th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*, IEEE, 2019, pp. 1–8.
- [23] C. Chen, W. Ge, SRMamba: Mamba for super-resolution of LiDAR point clouds, *arXiv preprint arXiv:2505.10601* (2025).
- [24] Y. Kwon, M. Sung, S.-E. Yoon, Implicit LiDAR network: LiDAR super-resolution via interpolation weight prediction, in: *2022 International Conference on Robotics and Automation (ICRA)*, IEEE, 2022, pp. 8424–8430.

- [25] D. Tian, D. Zhao, D. Cheng, J. Zhang, LiDAR super-resolution based on segmentation and geometric analysis, *IEEE Transactions on Instrumentation and Measurement* 71 (2022) 1–17.
- [26] C. Tu, E. Takeuchi, C. Miyajima, K. Takeda, Compressing continuous point cloud data using image compression methods, in: *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, IEEE, 2016, pp. 1712–1719.
- [27] C. Tu, E. Takeuchi, A. Carballo, K. Takeda, Real-time streaming point cloud compression for 3D LiDAR sensor using U-Net, *IEEE Access* 7 (2019) 113616–113625.
- [28] C. Tu, E. Takeuchi, A. Carballo, C. Miyajima, K. Takeda, Motion analysis and performance improved method for 3D LiDAR sensor data compression, *IEEE transactions on intelligent transportation systems* 22 (2019) 243–256.
- [29] Y. Feng, S. Liu, Y. Zhu, Real-time spatio-temporal LiDAR point cloud compression, in: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2020, pp. 10766–10773.
- [30] Q. Wang, L. Jiang, X. Sun, J. Zhao, Z. Deng, S. Yang, An efficient LiDAR point cloud map coding scheme based on segmentation and frame-inserting network, *Sensors* 22 (2022) 5108.
- [31] N. Muhammad, S. Lacroix, Calibration of a rotating multi-beam LiDAR, in: *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 2010, pp. 5648–5653.
- [32] C. Glennie, D. D. Lichti, Static calibration and analysis of the Velodyne HDL-64E S2 for high accuracy mobile scanning, *Remote sensing* 2 (2010) 1610–1624.
- [33] C. L. Glennie, A. Kusari, A. Facchin, Calibration and stability analysis of the VLP-16 laser scanner, *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* 40 (2016) 55–60.
- [34] T. O. Chan, D. D. Lichti, Feature-based self-calibration of Velodyne HDL-32E LiDAR for terrestrial mobile mapping applications, in: *The 8th International Symposium on Mobile Mapping Technology*, Tainan, Taiwan, 2013, pp. 1–3.
- [35] P. V. Hough, Method and means for recognizing complex patterns, 1962. US Patent 3,069,654.
- [36] R. O. Duda, P. E. Hart, Use of the Hough transformation to detect lines and curves in pictures, *Communications of the ACM* 15 (1972) 11–15.
- [37] J. Wang, P. J. Howarth, Use of the Hough transform in automated lineament, *IEEE transactions on geoscience and remote sensing* 28 (1990) 561–567.
- [38] N. Aggarwal, W. C. Karl, Line detection in images through regularized Hough transform, *IEEE transactions on image processing* 15 (2006) 582–591.
- [39] F. Tarsha-Kurdi, T. Landes, P. Grussenmeyer, Hough-transform and extended RANSAC algorithms for automatic detection of 3D building roof planes from LiDAR data, in: *ISPRS Workshop on Laser Scanning 2007 and SilviLaser 2007*, volume 36, 2007, pp. 407–412.
- [40] M. Yermo, R. Laso, O. G. Lorenzo, T. F. Pena, J. C. Cabaleiro, F. F. Rivera, D. L. Vilariño, Powerline detection and characterization in general-purpose airborne LiDAR surveys, *IEEE journal of selected topics in applied earth observations and remote sensing* 17 (2024) 10137–10157.
- [41] A. H. Safaie, H. Rastiveis, A. Shams, W. A. Sarasua, J. Li, Automated street tree inventory using mobile LiDAR point clouds based on Hough transform and active contours, *ISPRS Journal of Photogrammetry and Remote Sensing* 174 (2021) 19–34.
- [42] T. Risse, Hough transform for line recognition: Complexity of evidence accumulation and cluster detection, *Computer Vision, Graphics, and Image Processing* 46 (1989) 327–345.
- [43] D. E. Knuth, *The art of computer programming. Vol.3: Sorting and searching*, 1973.
- [44] A. Geiger, P. Lenz, C. Stiller, R. Urtasun, Vision meets robotics: The KITTI dataset, *The International Journal of Robotics Research* 32 (2013) 1231–1237.
- [45] T. Popoviciu, Sur les équations algébriques ayant toutes leurs racines réelles, *Mathematica* 9 (1935) 20.
- [46] Velodyne Lidar, Inc., User’s Manual and Programming Guide: HDL-64E S3 High-Definition LiDAR Sensor, 2019. URL: https://www.researchgate.net/profile/Joerg_Fricke/post/How_the_LiDARs_photodetector_distinguishes_lasers_returns/attachment/5fa947b8543da600017dcf9b/AS%3A955957442002980%401604929423693/download/HDL-64E_S3_UsersManual.pdf, accessed: 2025-10-07.
- [47] Ouster, Inc., OS1 LIDAR Sensor Datasheet, San Francisco, CA, USA, 2025. URL: <https://data.ouster.io/downloads/datasheets/datasheet-rev7p1-os1.pdf>, revision 7.1, Accessed: 2025-10-07.
- [48] Hesai Technology, Pandar128E3X User Manual: 128-Channel Mechanical LiDAR, 2025. URL: <https://www.hesaitech.com/product/pandar128/>, version 128-en-241220, Accessed: 2025-10-07.
- [49] RoboSense, Ruby Plus Enhanced 128-Beam LiDAR User Guide, 2023. URL: <https://cdn.robotshop.com/media/U/UTR/RB-Utr-09/pdf/robosense-ruby-plus-enhanced-128-beam-lidar-for-l4-autonomous-vehicles-user-guide.pdf>, version 1.1, Accessed: 2025-10-07.

- [50] L. Li, K. N. Ismail, H. P. Shum, T. P. Breckon, DurLAR: A high-fidelity 128-channel LiDAR dataset with panoramic ambient and reflectivity imagery for multi-modal autonomous driving applications, in: 2021 International Conference on 3D Vision (3DV), IEEE, 2021, pp. 1227–1237.
- [51] Fundación Pública Gallega Centro Tecnológico de Supercomputación de Galicia, CESGA FinisTerra III, <https://www.cesga.es/en/infrastructures/computing/finisterrae-iii/>, 2022. Spanish Supercomputing Network (RES).
- [52] R. D. Hipp, SQLite, <https://www.sqlite.org/index.html>, 2025.
- [53] T. Wu, L. Pan, J. Zhang, T. Wang, Z. Liu, D. Lin, Density-aware chamfer distance as a comprehensive metric for point cloud completion, arXiv preprint arXiv:2111.12702 (2021).
- [54] J. Heo, C. Phillips, A. Gavrilovska, FLICR: A fast and lightweight LiDAR point cloud compression based on lossy RI, in: 2022 IEEE/ACM 7th Symposium on Edge Computing (SEC), IEEE, 2022, pp. 54–67.
- [55] F. Lin, Y. Yue, S. Hou, X. Yu, Y. Xu, K. D. Yamada, Z. Zhang, Hyperbolic chamfer distance for point cloud completion, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2023, pp. 14595–14606.
- [56] F. Williams, Point cloud utils, <https://github.com/fwilliams/point-cloud-utils>, 2022.