Low-Cost LiDAR Mapping on Bicycles for Urban Road and Sidewalk Detection

Titiksha Bhatia[1], Salil Goel[1], Aditya Medury[1]

[1]Department of Civil Engineering, Indian Institute of Technology Kanpur, Kanpur, Uttar Pradesh, India

(titikshab, sgoel, amedury)@iitk.ac.in

Keywords: Autonomous vehicle navigation, Bicycle-mounted, Cost-effective, Ground point filtering, LiDAR point cloud, Road extraction, Sidewalk detection, Velodyne VLP-16.

Abstract

Extracting road information from Lidar point cloud data is crucial for autonomous vehicle navigation, urban planning, and infrastructure management applications. Lidar technology provides detailed 3D representations of environments, making it an effective tool for capturing road and terrain features. Traditional setups, where Lidar sensors are mounted on vehicles or drones, can be limited in complex environments like narrow streets or areas with dense vegetation. This research introduces a novel approach by mounting a Velodyne VLP-16 Lidar sensor on a bicycle, offering increased manoeuvrability in restricted areas and enabling data collection in places inaccessible to vehicles or drones. This bicycle-mounted setup is also cost-effective, providing high-resolution data without expensive equipment. The study presents a methodology that begins with ground point extraction, filtering out non-ground elements to isolate the road surface. Further, specialised algorithms were developed to accurately identify and extract road and sidewalk points from the filtered data, accommodating the varying elevations and textures typical of urban environments. The Lidar data was supplemented with RGB images collected simultaneously during data acquisition, providing additional context for validation. Comparison with ground truth data yielded an 85% to 90% accuracy rate, demonstrating the reliability of the approach in identifying road and sidewalk features. The results of this study have broad applications, particularly in urban planning and autonomous navigation systems.

1. INTRODUCTION

Lidar (Light Detection and Ranging) technology has become a cornerstone in spatial data collection, offering exceptional precision in creating high-resolution 3D models of various environments. Its applications span across fields such as urban planning, infrastructure management, autonomous vehicle navigation, and environmental monitoring. Accurately mapping and detecting road surfaces and sidewalks is critical for developing safe, efficient, and sustainable transportation systems. Recent advancements in Lidar technology have significantly enhanced the capacity to generate detailed 3D representations of both urban and non-urban environments. Traditionally, vehicle or drone-mounted Lidar systems have been the primary methods for large-scale data acquisition, especially in open terrains and metropolitan areas. However, these methods often encounter limitations in constrained or complex environments, such as narrow city streets or dense forests, where access is more challenging.

Traditional vehicle-mounted LiDAR systems offer high accuracy and extensive coverage for urban road mapping but are often constrained by traffic regulations, limited accessibility in narrow streets, and high operational costs. Similarly, drone-based LiDAR provides excellent aerial perspectives for large-scale mapping but faces challenges such as short battery life, restricted flight zones, and high acquisition costs. Bicycle-mounted LiDAR presents a unique advantage by navigating narrow urban roads, pedestrian paths, and areas inaccessible to larger vehicles and drones. Unlike vehicle-mounted systems, which struggle in pedestrian-heavy zones and complex terrains, bicycles can collect detailed point clouds at lower altitudes with minimal obstructions. Additionally, this method is significantly more cost-effective due to its lower equipment and maintenance expenses than drones, which require frequent battery replacements and regulatory approvals. However, a key limitation of bicycle-mounted systems is their slower data acquisition rate compared to vehicle-based LiDAR. Unlike drones, which capture larger areas in a single flight, bicycles must traverse the entire study region, making data collection more time-intensive. Nevertheless, for applications requiring fine-grained urban road and sidewalk detection, this approach provides a compelling alternative with unique strengths in accessibility and affordability.

2. LITERATURE REVIEW

In recent years, Lidar technology has advanced beyond its traditional uses in topography and forestry, becoming a vital tool in urban planning, infrastructure monitoring, and environmental assessment. As noted by Shan and Toth (2018), Lidar's capability to penetrate dense vegetation and provide highly accurate ground measurements has been transformative, especially in urban environments. This precision in capturing 3D spatial data has revolutionised urban mapping, producing detailed models crucial for modern urban development. However, extracting road networks from Lidar data remains challenging in diverse environments. Urban areas often feature complex structures and dense vegetation that obscure road detection. Premebida et al. (2014) demonstrated the benefits of combining Lidar with RGB data to improve road detection in urban settings, but this approach still faces limitations in regions with dense canopy cover or narrow pathways. Haichi et al. (2022) highlighted the importance of advanced ground-filtering algorithms in effectively distinguishing ground from non-ground points, particularly in challenging terrains.

Bicycle-mounted Lidar systems offer a novel solution for data collection, especially in environments where vehicle-mounted or drone-based systems may not be practical. These systems combine manoeuvrability with precision, making them ideal for capturing data in narrow streets, forest trails, and other complex terrains. Tokorodani (2019) demonstrated how a bicycle equipped with multilayer Lidar could generate high-resolution 3D point clouds at a low cost. While this approach is cost-effective, the dynamic movement of the bicycle presents challenges related to data stability and quality. Zai et al. (2019) further refined these techniques by integrating sensors like RGB cameras and GPS to enhance accuracy and mitigate these challenges.

Recent developments in Lidar-based road and sidewalk detection have led to significant improvements in urban infrastructure mapping, particularly for autonomous vehicles and urban planning. Huang et al. (2021) developed a real-time curb and lane detection system using Lidar point clouds, achieving high accuracy by leveraging geometric features, making it suitable for autonomous navigation. Similarly, Horváth et al. (2021) introduced a real-time system for efficiently segmenting road and sidewalk surfaces, improving vehicle navigation and safety. Ma et al. (2021) utilised PointNet++ with a two-step post-processing approach to extract urban road footprints from airborne Lidar, providing a scalable solution for city-wide infrastructure monitoring. Ground filtering, essential in Lidar data processing, has also seen advancements, such as Zhang et al.’s (2016) Cloth Simulation Filter (CSF), which effectively distinguishes ground from non-ground points, and the experimental results yield an average total error of 4.58%, which is comparable with most of the state-of-the-art filtering algorithms. Hui et al. (2016) employed morphological and interpolation filters. Paigwar et al. (2021) introduced GndNet, a deep learning-based method for rapid ground classification, suitable for real-time applications like autonomous driving. These studies illustrate Lidar's growing role in precise, large-scale urban mapping.

The increasing demand for flexible and adaptive Lidar data acquisition methods has led to the exploration of alternative systems. Bicycle-mounted Lidar offers a unique solution, combining the manoeuvrability of bicycles with the precision of Lidar technology to collect data in challenging environments such as narrow streets and forest trails. This paper aims to address the limitations of traditional Lidar data acquisition methods by introducing an innovative approach using a Velodyne VLP-16 Lidar sensor mounted on a bicycle. By leveraging the accessibility of bicycles in complex environments, this method offers a cost-effective and versatile alternative for Lidar data collection. The primary objective of this research is to accurately extract road and sidewalk features from Lidar data, with applications aimed at improving autonomous navigation, urban planning, and infrastructure management.

3. EXPERIMENTAL SETUP AND DATA COLLECTION

Figure 1 illustrates a specialised data collection setup mounted on a bicycle. This setup enables the capture of Lidar point cloud data and synchronised RGB video footage. This setup was designed to navigate the intricate road networks within the IIT Kanpur campus, providing a flexible and low-cost solution for urban road detection, especially in environments where conventional vehicle-based methods may be less practical.

Figure 1. A specialised data collection setup mounted on a bicycle.

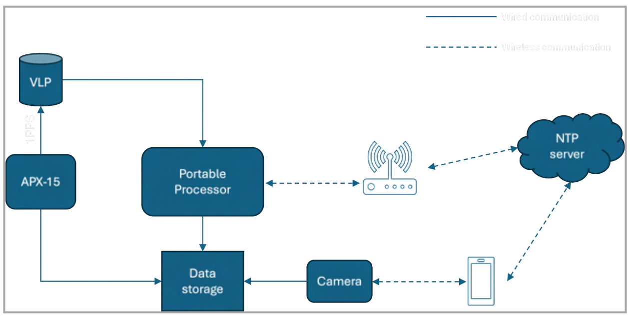

Figure 2. The flow chart of the equipment setup showcases how the different components are integrated to enable seamless data collection during the study.

The central component of this setup was the Velodyne VLP-16 Lidar sensor. The VLP-16, a compact and lightweight Lidar unit, can produce high-resolution 3D point clouds with a 360-degree horizontal field of view. The VLP-16 was directly connected to an APX-15 positioning and orientation system, which provided precise GNSS and inertial measurements. The APX-15 system integrates GNSS data with inertial measurements to deliver accurate position and orientation information, which is crucial for georeferencing the Lidar data. The integration of these two systems ensured that each point in the Lidar cloud was accurately tied to a real-world coordinate, even as the bicycle moved through the campus environment.

To complement the Lidar data, an RGB camera was mounted alongside the VLP-16 sensor. The camera was configured to capture high-definition video footage, providing visual context that could be used to enhance the interpretation of the Lidar data. This setup ensured that each frame of video footage could be precisely matched with the corresponding Lidar data points. The camera settings, including exposure and focus, were adjusted using a mobile application, allowing for real-time optimization of image quality as the lighting conditions varied during the data collection process. The entire data collection process was managed by a portable processor, which was mounted on the bicycle alongside the other equipment. The portable processor handled the continuous data streams from the Lidar sensor, the APX-15 system, and the RGB camera. This processor ensured that all incoming data were synchronised and correctly stored, maintaining the integrity of the dataset. Before initiating the data collection, the portable processor was connected to the internet to log into a Network Time Protocol (NTP) server. This step was critical for synchronising the system clock with global time standards, ensuring that all data points were accurately time-stamped. This synchronisation was essential for post-processing, where precise timing is required to align the different data streams. Once the time synchronisation was completed, the internet connection was disabled, and the bicycle-mounted system began recording data.

Figure 3 illustrates the route traversed by the bicycle within IIT Kanpur. The journey commenced from a marked red point on one side of the road, with the rider heading in the left direction and eventually returning to the starting point. The VLP-16 LiDAR sensor continuously captured 3D point cloud data throughout the journey, while an RGB camera recorded video footage. All collected data were stored on a connected storage drive, resulting in a comprehensive dataset that integrated spatial, visual, and positional information.

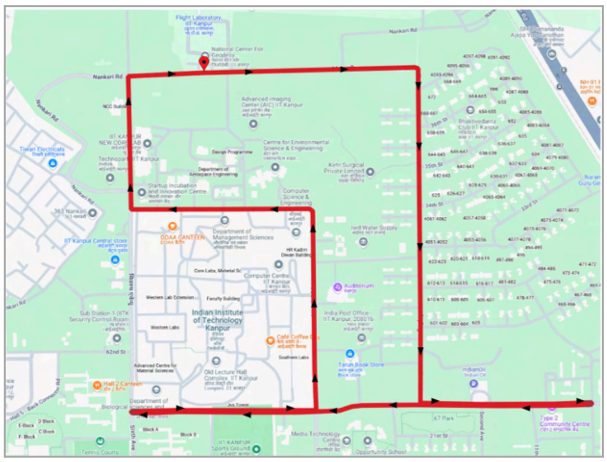

Figure 3. Route traversed by the bicycle within the IIT Kanpur campus.

The data collection process lasted approximately 25 minutes, covering a 5 km stretch of the IIT Kanpur campus. The VLP-16 Lidar sensor captured around 12,780 frames during this ride, covering various road types and conditions to ensure a diverse and comprehensive dataset. The high frame rate of the VLP-16 allowed for detailed 3D point cloud generation, capturing subtle variations in the road surface and surrounding environment. Simultaneously, the RGB camera recorded video, providing complementary visual data that enriched the dataset and enabled more robust analysis and feature extraction in subsequent processing stages.

Figure 4. RGB image illustrating sidewalk on one side (image taken by the camera onboard the bicycle).

Figure 5. RGB image illustrating sidewalk on both sides (image taken by the camera onboard the bicycle).

Figure 4. and Figure 5. illustrate the RGB images of two different road types captured during data collection. In the first case, shown in Figure 4, a sidewalk is on one side of the road, specifically on the left. The RGB image captured during the video validates the key features analysed in this study: the sidewalk on the left is marked with a green line, the road in the centre is black, and the roadside area on the right is indicated in yellow. Blue markings represent areas where changes in angle were detected, helping to differentiate these key elements within the scene. In the second case, as shown in Figure 5, sidewalks are on both sides of the road, marked in green.

(a) RGB image of a frame taken by the camera onboard the bicycle.

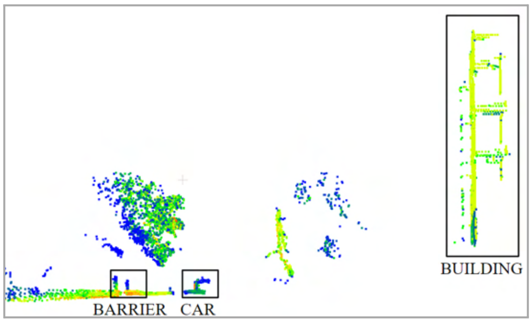

(b) Front view of the VLP16 data frame with visible car features, barriers, and building features.

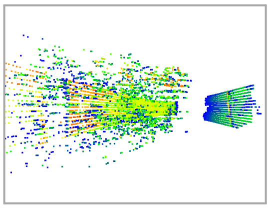

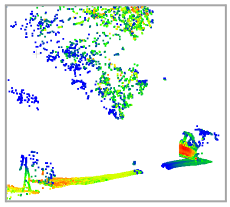

(c) Zoomed-in Top view (d) Zoomed-in side view

Figure 6. A detailed analysis of a single frame from the collected data. Panel (a) shows the RGB image captured by the camera, providing the visual context of the environment. Panel (b) presents a front view of the LiDAR data frame, highlighting features such as a car, barriers, and a building. Panel (c) offers a top view of the LiDAR point cloud, illustrating the spatial distribution and structure of the environment. Panel (d) provides a side zoomed-in view, allowing for a closer examination of specific elements within the point cloud, enhancing the understanding of the data's depth and detail.

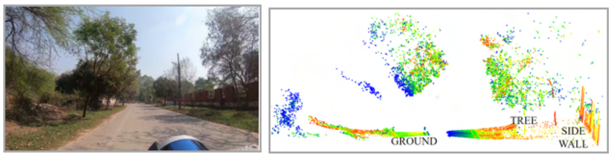

(a) (b)

Figure 7. Sample dataset displaying (a) RGB images on the left and (b) the corresponding LiDAR point clouds on the right. A comparison between the LiDAR and RGB views highlights the identification of key features, including the ground surface, trees, and sidewalls.

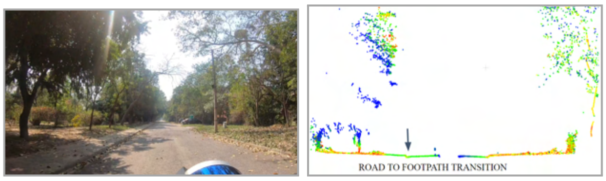

(a) (b)

Figure 8. (a) RGB images on the left and (b) the corresponding LiDAR point clouds on the right, illustrating the transition from road to footpath as highlighted in the LiDAR point cloud.

4. METHODOLOGY

The workflow of this study is depicted in Figure 9. Road surface points are extracted from the raw LiDAR data using the Cloth Simulation Filtering (CSF) algorithm, with careful tuning and manipulation of its parameters. For road and sidewalk detection, Principal Component Analysis (PCA) is applied at various stages, utilising the point cloud's x, y, and z coordinates. A threshold selection method is then employed to accurately identify the transition point between the road and the sidewalk.

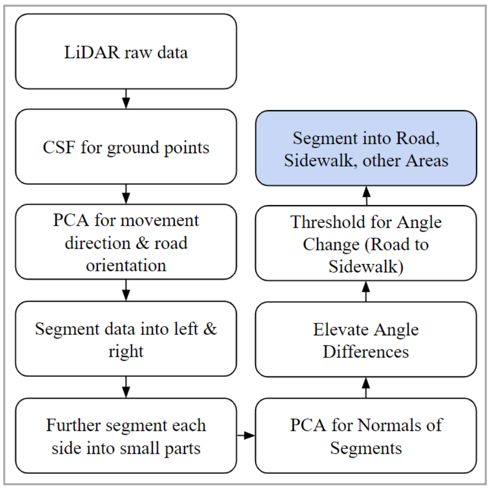

Figure 9. The workflow of road and sidewalk detection algorithm.

4.1 Ground Filtering Method

The Cloth Simulation Filter (CSF) is a ground filtering method used in LiDAR data processing. It is a simulation-based technique that employs a physics-based cloth simulation to classify LiDAR points as either ground or non-ground. The overview of the CSF algorithm is illustrated in Figure 10 as explained by Zhang et al. (2016). The Cloth Simulation Filter (CSF) method begins by inverting the original point cloud data, effectively turning the terrain upside down. Following this inversion, a virtual cloth is simulated as it "drops" onto the inverted surface from above. As the cloth descends, its nodes interact with the corresponding Lidar points. These interactions are influenced by factors such as the cloth's rigidity, resolution, and the time step of the simulation. As the cloth conforms to the contours of the inverted surface, it gradually settles into a shape that mirrors the underlying terrain. The final configuration of the cloth is then used as a reference to classify the original Lidar points. Points that lie on or near the surface of the cloth are categorised as ground points, while those that do not are classified as non-ground points. This method effectively distinguishes between terrain and above-ground objects, such as buildings or vegetation, enabling accurate ground surface extraction even in complex environments. The process is iterative, with multiple passes often required to refine the cloth's shape and achieve optimal classification accuracy.

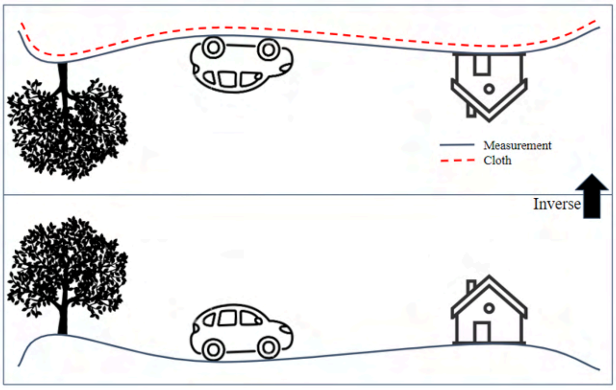

Figure 10. Overview of the CSF algorithm (Zhang et al. 2016).

CSF has gained popularity for its adaptability to different terrains and its effectiveness in classifying ground points in various environmental conditions. In this section, we will delve into the details of the CSF method, explaining how it works and the key parameters involved.

Cloth Resolution (CR) represents the horizontal distance between two neighbouring particles in the cloth grid, controlling the level of detail and accuracy in ground surface modelling. Mathematically, it can be expressed as

$$CR = \frac{1}{N}\sum_{i=1}^{N} d_i \tag{1}$$

where N is the total number of grid points and d_i is the distance between neighbouring particles.

Time Step (Δt) controls the displacement of particles from gravity during each iteration, influencing how the cloth interacts with the inverted terrain surface. The iterative process to minimise the cloth's energy over time can be modelled as

$$X(t + \Delta t) = X(t) + V(t) \cdot \Delta t \tag{2}$$

where $X(t)$ is the position of a cloth particle at time (t), $V(t)$ is the velocity of the cloth particle at time (t) and Δt is the time step (Δt).

Rigidness (R) controls the resistance of the cloth to deformation, affecting its ability to conform to the terrain surface—higher rigidness results in less deformation. There is no direct formula, but it impacts the elasticity and flexibility of the cloth during the simulation. It is an integer that varies from 1 to 3; 1 stands for very soft (to fit rugged terrain), 2 stands for medium, and 3 stands for hard cloth (for flat terrain). The steep slope fit factor (ST**)** is an optional parameter indicating whether post-processing is required to handle steep slopes. If activated, it ensures that the cloth accurately fits steep terrains. There is no direct formula, but it is a binary setting that triggers additional processing steps if set. When steep slopes exist, set this parameter to TRUE to reduce errors during post-processing.

The Classification Threshold (CT) defines the distance used to differentiate ground points from non-ground points. It represents the scalar distance from the simulated cloth surface, with values ranging from 0 to 1. Through experimentation, the value between 0.1 and 1 was adjusted to obtain optimal classification results. The equation is as follows:

$$\text{Ground Classification} = \begin{cases} \text{Ground Point} & \text{if } d \leq CT \\ \text{Non-Ground Point} & \text{if } d > CT \end{cases} \tag{3}$$

where d is the distance from a LiDAR point to the cloth and CT is the classification threshold.

The iteration process involves repeating the simulation multiple times to model the cloth's behavior until convergence is achieved. The default iteration count is set to 500. The equation governing this iterative process is as follows:

$$X_{i+1} = X_i + \Delta X \tag{4}$$

where X_i represents the state of the cloth at iteration i and ΔX is the change in the cloth's position during each iteration.

The filtering equation is used to classify points from a point cloud P as ground points based on their proximity to a simulated cloth surface in the Cloth Simulation Filter (CSF) method. The overall equation is

$$\text{Ground Points} = \{p \in P \mid \min(|p - X(t)|) \leq CT\} \tag{5}$$

where $p \in P$ represents an individual point in the point cloud P, $X(t)$ represents the position of the simulated cloth at time t. The cloth's position changes over time as it drapes over the inverted point cloud to simulate ground detection, $\min(|p-X(t)|)$: The minimum distance between the point p and the cloth $X(t)$ over all possible positions of the cloth, CT (Classification Threshold) is a threshold value used to determine whether a point p is classified as a ground point.

In the CSF methodology, the LiDAR dataset consists of a point cloud containing XYZ coordinates, where each point represents a snapshot of the environment, including roads, sidewalks, and other urban infrastructure. By applying the Cloth Simulation Filtering (CSF) to this raw dataset, the method identifies points that are classified as ground by measuring the minimum distance between each point p and the simulated cloth $X(t)$. If the distance is less than or equal to the classification threshold (CT), the point is labelled as a ground point. This process filters the point cloud, removing all above-ground features such as buildings and trees, and leaving behind only the ground-level information, as illustrated in Figure 11. This filtered data provides a clean representation of the ground surface, crucial for road and sidewalk extraction in the further process.

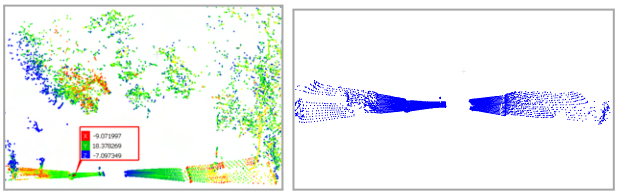

(a) (b)

Figure 11. (a) A raw data frame with environmental information, (b) Only ground point cloud after CSF.

4.2 Road and Sidewalk Split Algorithm

This study introduces a comprehensive methodology for extracting road and sidewalk features from LiDAR point cloud data by combining Principal Component Analysis (PCA) and angle-based thresholding. PCA is used for dimensionality reduction and calculating normal vectors, providing insights into the orientation of surfaces. Angle-based thresholding is then applied to detect transitions between surfaces, such as roads and sidewalks. This approach allows for precise segmentation and classification of roadways and adjacent structures, offering a detailed understanding of urban infrastructure. The method is highly applicable to autonomous navigation, urban planning, and infrastructure monitoring, where accurate road and sidewalk extraction is crucial.

The input data consists of point cloud datasets containing only ground-level information after applying the Cloth Simulation Filtering (CSF), which removes all above-ground environmental details. Each file captures the 3D spatial coordinates collected by a LiDAR sensor, with every point comprising three components: x, y, and z.

Let the point cloud data from each file be represented as:

$$D = \{p_1, p_2, \ldots, p_n\}, \quad p_i = (x_i, y_i, z_i) \tag{6}$$

where D is the set of n points, and p_i is a point in 3D space with coordinates x_i, y_i, z_i.

Principal component analysis (PCA) is applied to the point cloud data to determine the orientation of the road surface relative to the data collection direction. PCA is employed to reduce the dimensionality of the data while preserving the most significant variance. The covariance matrix Σ of the dataset is calculated:

$$\Sigma = \frac{1}{n-1} \sum_{i=1}^{n} (p_i - \mu)(p_i - \mu)^T \tag{7}$$

where p_i is the individual points and μ is the mean of the points.

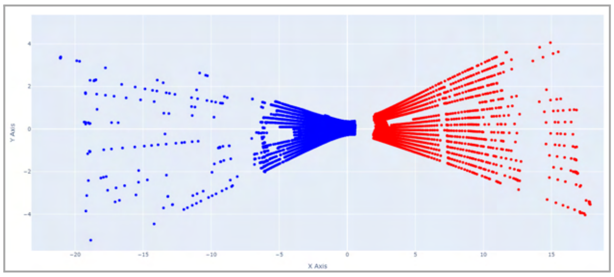

Figure 12. Presents a 2D plot of the LiDAR point cloud, illustrating the dataset divided into two segments—blue on the left and red on the right. The split occurs along a designated axis aligned with the road's primary direction.

Each side of the point cloud (positive and negative) is further segmented into smaller patches, based on the transformed x-axis coordinates. The segmentation is performed with a pre-defined segment size Δ*s*, creating a sequence of non-overlapping patches:

$$\text{Segment range} = [x_i, x_i + \Delta s] \tag{9}$$

These segments enable localised analysis of the point cloud data to account for variations in road surfaces and adjacent features. This granularity is necessary for detailed surface analysis and subsequent normal calculation. For each segment, PCA is applied again to calculate the normal vectors. The covariance matrix of the points in each segment is recalculated, and the normal vector n is determined by solving for the eigenvector corresponding to the smallest eigenvalue:

$$\Sigma_{\text{segment}} n = \lambda_{\min} n \tag{10}$$

where λ_{min} represents the smallest eigenvalue, and n is the normal vector. This normal vector describes the orientation of the surface in 3D space, which is crucial for identifying road versus non-road surfaces.

The normal vectors for each segment are compared to determine changes in surface orientation. The angle θ between two adjacent normal vectors n_1 and n_2 is calculated using the dot product formula:

$$\theta = \arccos\left(\frac{n_1 \cdot n_2}{\|n_1\| \|n_2\|}\right) \tag{11}$$

where $n_1 \cdot n_2$ is the dot product of the normals, and $\|n_1\|$ and $\|n_2\|$ are the magnitudes of the vectors. A significant angle difference between adjacent normals typically indicates a transition between surfaces (e.g., road to sidewalk).

An angle threshold detects transitions between road surfaces and other adjacent areas (such as sidewalks). If the calculated angle θ between two segments exceeds the specified threshold $\theta_{threshold}$, a surface transition is identified:

$$\theta > \theta_{\text{threshold}} \implies \text{surface transition (e.g., road to sidewalk)} \tag{12}$$

This thresholding approach enables the classification of point cloud segments based on surface characteristics.

Based on the identified transitions, the point cloud data is classified into distinct segments representing roads, sidewalks, and other features. The segments on either side of the identified transition points are labeled and saved accordingly.

7. RESULT AND DISCUSSION

The results of this study demonstrate the effective use of a low-cost, bicycle-mounted LiDAR system in accurately identifying and classifying road and sidewalk features in an urban setting. The analysis covered approximately 5 kilometers of campus roads in one direction, with 1 kilometer featuring a sidewalk on only one side, while the remaining sections had sidewalks on both sides, mimicking a typical urban environment suitable for modeling future city planning. A total of 12,780 frames were captured along this stretch, providing substantial data for analysis. The RGB images, LiDAR point clouds, and filtered ground points from Figures 9 and 10 illustrate the success of the proposed algorithm in distinguishing and segmenting key features within these varied urban road layouts, offering promising results for urban infrastructure mapping.

(a)

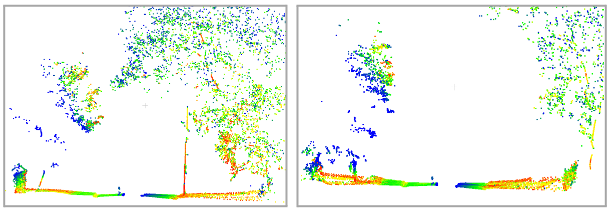

(b)

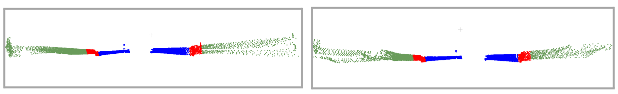

(c)

Figure 13. Illustrates the Sidewalk only on the Left.

(a)

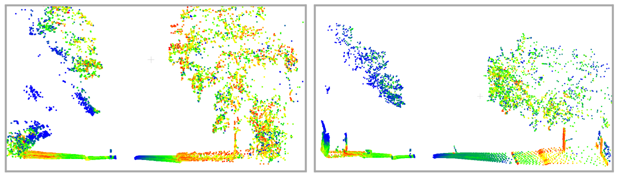

(b)

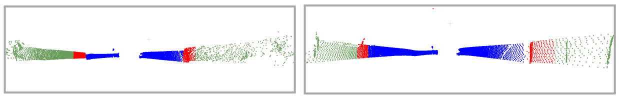

(c)

Figure 14. Shows the example frames with Sidewalks on both ends.

Figure 13 and Figure 14 illustrate the effectiveness of the proposed algorithm in segmenting road and sidewalk features. In each figure, (a) shows the RGB images, (b) displays the corresponding complete raw LiDAR point cloud, and (c) highlights the processed ground point cloud with road points in blue, sidewalk points in red, and remaining data in green. The segmentation results align well with the visual information from the RGB images, providing a qualitative validation of the algorithm's accuracy.

From the results, the confusion matrix for the near-side (left) detection is explained in Table 1:

	Predicted Road	Predicted Sidewalk	Precision	Recall
Actual Road	5775 (TP)	268 (FN)	0.857	0.956
Actual Sidewalk	962 (FP)	5775 (TN)		

Table 1: For the detection of road and sidewalk features on the near side (left) of the road, the results show that 5,775 frames are correctly identified as road (true positives), and 5,775 frames are accurately classified as sidewalk (true negatives). However, in 962 frames, the point clouds are incorrectly classified as road (false positives), and in 268 frames, the point clouds are misclassified as sidewalk (false negatives).

The confusion matrix for the far-side (left) detection is explained in Table 2:

	Predicted Road	Predicted Sidewalk	Precision	Recall
Actual Road	5492 (TP)	556 (FN)	0.816	0.908
Actual Sidewalk	1240 (FP)	5492 (TN)		

Table 2: The algorithm yields comparable results for the detection on the right side. 5,492 frames are correctly identified as roads (true positives), and 5,492 frames are accurately classified as sidewalks (true negatives). However, 1,240 frames are incorrectly classified as roads (false positives), and 556 frames are misclassified as sidewalks (false negatives).

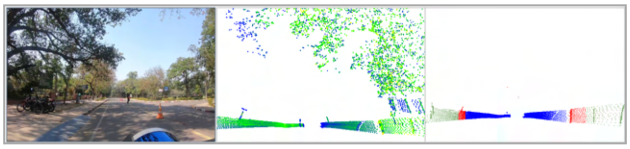

(a)

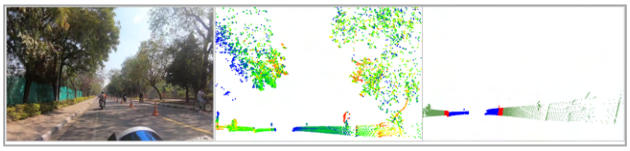

(b)

Figure 15. (a) and (b) Illustrates the example result of False Positive and False Negative respectively. **False positives**, where sidewalk points were misclassified as roads, were primarily due to the smooth transition between road and sidewalk surfaces. Conversely, **false negatives**, where road points were incorrectly classified as sidewalks, were often caused by objects like cone dividers or passing pedestrians. These elements introduced additional point clouds that were not entirely filtered out by the CSF (Cloth Simulation Filtering) algorithm, leading to incorrect classifications. These factors significantly impacted the precision and recall rates, elaborated through the confusion matrix analysis above.

These results highlight the robustness of the proposed algorithm in accurately identifying and segmenting road and sidewalk features using a lightweight, portable data collection system. Integrating LiDAR and RGB imagery enhances the system's capability to interpret and classify spatial data. RGB imagery provides valuable visual validation alongside the precise 3D measurements from the LiDAR point clouds. Despite certain challenges, such as smooth transitions between road and sidewalk or the presence of temporary obstacles like cone dividers and pedestrians, the algorithm’s ability to handle these environments demonstrates its practical value. The methodology effectively reorients the coordinate system on the fly. It determines road boundaries and sidewalks using passive data collection, underscoring its potential for applications in autonomous vehicle navigation, urban planning, and smart city infrastructure development.

5. CONCLUSION AND FUTURE WORK

This research successfully demonstrates the potential of a low-cost, bicycle-mounted Lidar system for accurately detecting and classifying road and sidewalk features in urban environments. Combining Lidar point clouds with RGB imagery offers a practical and adaptable solution to urban mapping, overcoming the limitations of traditional vehicle- or drone-mounted sensors in constrained or densely vegetated areas. The Velodyne VLP-16 sensor, mounted on a bicycle, provides a cost-effective method to capture high-resolution 3D data. The developed algorithm, tested over a 5-kilometer stretch of campus roads, accurately segmented roads and sidewalks, achieving more than 90% accuracy on the left side and around 85% on the right side, when validated with RGB imagery. This highlights its effectiveness for various urban applications, from autonomous navigation to infrastructure management.

While the system performed with high accuracy, challenges such as false positives and false negatives were encountered. These issues were primarily due to smooth transitions between road and sidewalk surfaces and obstacles like pedestrians and traffic cones, which introduced noise that the CSF algorithm could not fully eliminate. Despite these challenges, the algorithm maintained strong precision and recall, showcasing its robustness for urban infrastructure mapping. Although the RGB imagery was used solely for validating the Lidar-based predictions, the integration of multimodal sensing proved essential in enhancing classification accuracy. Future work could explore sensor fusion strategies for road asset identification, moving beyond validation to improve detection capabilities in applications like autonomous vehicle navigation, smart city planning, and urban infrastructure management.

Future work will focus on refining the ground filtering and segmentation algorithms to enhance road and sidewalk boundary detection. Additionally, we plan to extend data collection to narrow roads and other challenging urban settings to evaluate the system’s adaptability in constrained environments. Another key aspect for future research is a more in-depth discussion of the applicability and limitations of the developed algorithms, particularly their generalizability across diverse urban landscapes. Furthermore, a comprehensive cost-effectiveness analysis will be conducted, addressing not only the initial setup expenses but also the long-term operational and maintenance costs associated with a bicycle-mounted LiDAR system. Expanding the system for city-wide deployments could enable continuous, real-time road and sidewalk data collection, supporting dynamic traffic management, pedestrian safety, and efficient urban planning—positioning it as a valuable tool for next-generation urban infrastructure solutions.

REFERENCES

Fritsch, J., Kühnl, T., & Geiger, A. (2019). A new performance measure and evaluation benchmark for road detection algorithms. in Proceedings of the IEEE Intelligent Transportation Systems Conference (ITSC), 27–30.

Han, X., Huan, W., Lu, J., Zhao, C., (2017). Road detection based on the fusion of Lidar and image data. *International Journal of Advanced Robotic Systems*. 14(6).

Horváth, E., Barsi, Á., Szabó, G., & Lovas, T. (2021). Real-time LIDAR-based urban road and sidewalk detection for autonomous vehicles. *Sensors*, 22(1), 194.

Hu, H., Hu, Z., Wang, X., & Lu, J. (2021). Lidar-based real-time 3D mapping for smart city applications. *Remote Sensing*, 13(11), 2239.

Huang, J., Yang, G., Zhang, Y., Zhao, Q., & Zhou, F. (2021). Real-time road curb and lane detection for autonomous driving using LiDAR point clouds. *IEEE Transactions on Intelligent Vehicles*, 6(1), 134-146.

Jung, J., Kim, D., Lee, K., & Heo, J. (2020). Multi-sensor data fusion for the detection and mapping of urban elements. *Sensors*, 20(14), 3891.

Levinson, J., Askeland, J., Becker, J., Dolson, J., Held, D., Kammel, S., Kolter, J. Z., Langer, D., Pink, O., Pratt, V., & Thrun, S. (2018). Towards fully autonomous driving: Systems and algorithms. *IEEE Transactions on Intelligent Vehicles*, 3(4), 53-68.

Ma, H., Ma, H., Zhang, L., Liu, K., & Luo, W. (2021). Extracting urban road footprints from airborne LiDAR point clouds with PointNet++ and two-step post-processing. *Remote Sensing*, 14(3), 789.

Ma, H., Ma, H., Zhang, L., Liu, K., & Luo, W. (2021). Extracting urban road footprints from airborne LiDAR point clouds with PointNet++ and two-step post-processing. *Remote Sensing*, 13(2), 212.

Narksri, P., Xia, S., Yan, W., Wu, J., Li, X., & Yokoya, N. (2020). Two-stage ground filtering of airborne LiDAR data using RANSAC. *ISPRS International Journal of Geo-Information*, 9(4), 223.

Paigwar, A., Rangesh, A., & Trivedi, M. M. (2021). GndNet: Fast ground plane estimation and point cloud segmentation for autonomous vehicles. *IEEE Robotics and Automation Letters*, 6(2), 3498-3505.

Premebida, C., Carreira, J., Batista, J., & Nunes, U. (2014). Pedestrian detection combining RGB and dense LIDAR data. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE.

Shan, J., & Toth, C. K. (2018). *Topographic Laser Ranging and Scanning: Principles and Processing*. CRC Press.

Shen, Q., Lu, T., Luo, Z., & Wang, J. (2019). Efficient point cloud segmentation on rough terrains using Jump-Convolution-Process. *IEEE Geoscience and Remote Sensing Letters*, 17(5), 813-817.

Tokorodani, K., Hashimoto, M., Aihara, Y., & Takahashi, K. (2019). Point-cloud mapping using lidar mounted on two-wheeled vehicle based on NDT scan matching. *IEEE International Conference on Robotics and Automation (ICRA)*.

Velodyne Lidar. (2021). *Velodyne VLP-16 User Manual*. Velodyne Lidar, Inc.

Wang, G. J., Wu, J., He, R., & Tian, B. (2021). Speed and accuracy tradeoff for LiDAR data-based road boundary detection. *IEEE/CAA Journal of Automatica Sinica*, 8(6), 1210-1220.

Winiwarter, L., Pfeifer, N., Mandlburger, G., & Briese, C. (2017). Improving GPS accuracy using LiDAR data for urban road extraction. *Remote Sensing*, 9(6), 567.

Wu, B., Wan, A., Yue, X., & Keutzer, K. (2017). SqueezeSeg: Convolutional neural nets with recurrent CRF for real-time road-object segmentation from 3D LiDAR point cloud, *Proc. IEEE International Conference on Robotics and Automation*. 1887–1893.

Zai, D., Tang, Y., Wei, H., & Song, M. (2019). Integration of LiDAR, RGB imagery, and GPS for precise urban mapping. *Photogrammetric Engineering & Remote Sensing*, 85(10), 689-698.

Zhang, W., Qi, J., Wan, P., Wang, H., Xie, D., Wang, X., & Yan, G. (2016). An easy-to-use airborne LiDAR data filtering method based on cloth simulation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 8(6), 501.

https://doi.org/10.5194/isprs-archives-XLVIII-G-2025-197-2025