AI-Driven Extraction of Road Geometry and Asset Inventory from Mobile LiDAR Point Clouds

Divya Priya Balasubramani [1], Zaffar Sadiq Mohamed-Ghouse [1], Sanjay Khanna Diwakar [2], Ravichandran Narayanan [2], Muthu Kumara Samy Sudharsan [2]

[1] Institute of Remote Sensing, Department of Civil Engineering, College of Engineering Guindy, Anna University, Chennai – 600 025, India - divyapriya.b@annauniv.edu, zaffarmgsm@gmail.com

[2] Undergraduate Student, Institute of Remote Sensing, Department of Civil Engineering, College of Engineering Guindy, Anna University, Chennai – 600 025, India - sanjaykhanna727@gmail.com, 41ravichandran.n@gmail.com, muthukumarasamy733@gmail.com

Keywords: Mobile LiDAR, Road Design Parameters, Inventory Extraction, Artificial Intelligence, Point Cloud Classification, Automated Feature Tagging.

Abstract

The rapid urbanization and rising traffic volumes strain transportation infrastructure, demanding efficient road design auditing and asset management. Conventional manual surveys are labor-intensive and lack holistic three-dimensional context. This research presents an end-to-end methodology combining mobile LiDAR with an AI model to automate extraction of road geometric parameters and inventory features. Mobile LiDAR data from a Bengaluru corridor was preprocessed using Trimble Business Center, applying a statistical outlier removal filter and progressive morphological ground segmentation. A custom PointNet++-based deep learning architecture with hierarchical set abstraction layers was trained on manually labelled point cloud subsets (45 million points, 10% labeled) to classify roads, poles, vehicles, trees, and buildings. The model achieved 0.86 mean Intersection-over-Union (mIoU) and 92.4% overall accuracy on semantic segmentation. Key parameters—lane width (8.099 m), road length (44.383 m), zebra crossing dimensions (7.336 m), and pole height (7.890 m)—were accurately extracted. The automated workflow reduced manual processing time by 85% (from 40 to 6 hours per km), improving repeatability and scalability across urban corridors. Results confirm that the proposed AI-driven workflow significantly reduces manual effort while providing high-accuracy datasets for infrastructure planning. This study demonstrates the transformative potential of integrating mobile LiDAR and AI, offering a scalable tool for safer, more sustainable transportation systems.

1. Introduction

1.1 Background and Research Imperative

The global surge in urbanization and vehicular traffic has placed unprecedented strain on transportation infrastructure, necessitating advanced methodologies for its design, auditing, and management (Gargouma and El-Basyounya, 2018). Conventional road survey techniques, which rely heavily on manual inspections, total stations, and terrestrial measurements, are fundamentally constrained by their labour-intensive nature, significant time consumption, and inherent safety risks for field crews. These methods often produce disparate, two-dimensional datasets that lack the rich, three-dimensional context required for comprehensive geometric analysis, leading to potential oversights in design flaw detection and incomplete asset inventories (Gargoum and Basyouny, 2019).

1.2 Research Gap and Proposed Solution

While previous studies have demonstrated the utility of Mobile LiDAR for specific tasks such as pavement distress detection (Li et al., 2019, Mattes et al., 2023) or individual feature extraction like street lighting poles (Zheng et al., 2017), a significant gap remains in the development of an integrated, automated workflow. Most existing approaches treat feature extraction and geometric parameter evaluation as separate processes, often relying on hand-crafted rules or semi-automated algorithms that lack scalability.

This research confronts these limitations by posing a central question: Can a synergistic workflow combining Mobile LiDAR and a custom Artificial Intelligence model automate the simultaneous extraction of road design parameters and a comprehensive inventory of roadway assets? The primary objective is to develop an end-to-end solution that transitions from raw, unclassified point cloud data to actionable engineering intelligence, thereby establishing a new paradigm for efficient, accurate, and data-driven infrastructure management.

1.3 Contribution

This study contributes to the field of geospatial infrastructure management by:

https://doi.org/10.5194/isprs-archives-XLIX-B2-2026-627-2026

1. Developing and applying a custom AI model (PointNet++ backbone) for the semantic segmentation of mobile LiDAR point clouds into five distinct feature classes relevant to road corridors.

2. Demonstrating the automated extraction of key road geometric design parameters (widths, lengths, pole heights) from the classified data.

3. Validating the proposed methodology on a real-world, complex urban road network in Bengaluru, India, showcasing its robustness and practical applicability.

4. Quantifying the reduction in manual effort and the improvement in scalability compared to traditional methods.

2. Methodology

2.1 Data Acquisition and Study Area

The study was conducted on a 500-meter segment of a major arterial road in Bengaluru, Karnataka, India (approx. 12°97′16″N, 77°59′46″E). This area was selected due to its high traffic density, diverse roadside features (buildings, trees, poles), and complex road geometry, representing a realistic challenge for the proposed methodology.

Mobile LiDAR data was acquired using a Leica Pegasus TRK 100 system mounted on a survey vehicle. The system's key specifications are detailed in Table 1. The vehicle traversed the corridor at normal traffic speed, capturing approximately 45 million 3D points with RGB intensity values.

Sensor	Type	Measurement Rate / Specification	Description
Laser Scanners	Dual laser scanners	600,000 measurements/second	Captures dense 3D point clouds with high precision.
Cameras	Modular imaging system	High-resolution, configurable	Provides RGB color information for each point.
GNSS Receiver	Multi-frequency GNSS	Real-time kinematic (RTK)	Ensures precise georeferencing (2–10 cm accuracy).
IMU	Inertial Measurement Unit	Continuous (200 Hz)	Tracks orientation and motion during mapping.

Table 1. Sensor characteristics of the Leica Pegasus TRK 100 system

2.2 Data Preprocessing: Noise Removal and Initial Segmentation

The raw LiDAR data was imported into Trimble Business Center (TBC) v2025.10 for initial preprocessing. Two critical algorithms were applied sequentially:

2.2.1 Noise and Outlier Removal Algorithm A statistical outlier removal (SOR) filter was implemented. For each point, the mean distance to its k nearest neighbors (k=50) was computed. Points whose mean distance exceeded the global mean distance by more than 2 standard deviations were classified as outliers and removed. This method effectively eliminated isolated points caused by atmospheric backscatter, multi-path reflections from glass surfaces, and sensor noise. Approximately 1.2% of the raw points were removed as outliers, resulting in a cleaner point cloud for subsequent segmentation.

2.2.2 Initial Ground Segmentation Algorithm A progressive morphological filter (PMF) was used to separate ground points (road surface) from non-ground points (buildings, vegetation, poles). The PMF operates by gradually increasing the window size of a morphological opening operation, removing non-ground objects based on height differences. Key parameters: initial window size 1 m, maximum window size 20 m, slope threshold 0.5, and elevation threshold 1.5 m. This algorithm is robust to gradual terrain changes and was chosen for its computational efficiency and proven performance in urban environments. Ground points were isolated as the road surface class, while all other points were labeled as “unclassified” for subsequent AI-based semantic segmentation.

2.3 Dataset and Labels for AI Model Training

The preprocessed dataset consisted of 44.5 million points after outlier removal. To create a supervised training set, a stratified random sampling approach was used:

- Training set: 10% of the total points (4.45 million) were manually labelled using TBC's interactive segmentation tools.
- Validation set: 2% (0.89 million points) held out for hyperparameter tuning.
- Test set: 5% (2.23 million points) reserved for final evaluation, disjoint from training/validation.

Five semantic classes were defined based on their relevance to road design and inventory management as presented in Table 2:

Class ID	Class Name	Description	Number of Training Points
0	Road	Paved surface including lane markings, shoulders, and zebra crossings	1,850,000
1	Pole	Vertical structures (light poles, traffic signal poles, utility poles)	210,000
2	Vehicle	Moving or stationary cars, trucks, buses, autorickshaws	890,000
3	Tree	Vegetation with distinct trunk and canopy (excluding low bushes)	950,000
4	Building	Adjacent built structures with vertical facades	550,000

Table 2. Classification schema and training data distribution

Class imbalance was addressed using weighted cross-entropy loss during training, with weights inversely proportional to class frequency.

2.4 AI Model Architecture: Feature-Tagging Model

2.4.1 Backbone and Core Network The overall architecture is built upon the PointNet++ backbone (Qi et al., 2017), specifically the multi-scale grouping (MSG) version. PointNet++ was chosen because it directly consumes unstructured point clouds, learns hierarchical features, and is robust to varying point densities—critical for mobile LiDAR data where object density decreases with distance from the scanner.

Core components of the architecture:

- **Input:** Point cloud block of 4096 points, each with 6 features (x, y, z, RGB intensity normalized to [0,1]).
- **Set Abstraction Levels (SA):**
  - **SA1:** 512 centroids, radius 0.4 m, 16 samples per centroid → output 512×128
  - **SA2:** 128 centroids, radius 0.8 m, 32 samples per centroid → output 128×256

  - **SA3:** 32 centroids, radius 1.6 m, 64 samples per centroid → output 32×512

- **Feature Propagation (FP) Layers:**
  - FP1: Interpolate from 32 to 128 points
  - FP2: Interpolate from 128 to 512 points
  - FP3: Interpolate from 512 to 4096 points

- **Fully Connected Head:** 3 fully connected layers (512 → 256 → 128 → 5 classes) with dropout (0.3) and ReLU activations.

- **Output:** Softmax probability per point for the 5 classes.

The model was implemented in TensorFlow 2.12 with custom training loops (Figure 1). Training used categorical cross-entropy loss, Adam optimizer (learning rate 0.001, decay 0.0001), batch size 8, for 200 epochs. Early stopping with patience 20 epochs was applied based on validation loss.

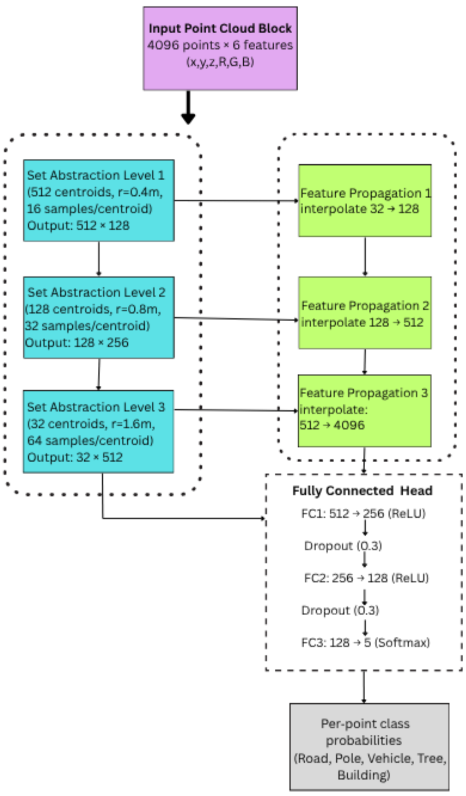

Figure 1. Overall architecture of the feature-tagging model based on PointNet++ with multi-scale grouping (MSG).

2.4.2 Training Environment

Training was conducted on Google Colab Pro with an NVIDIA Tesla T4 GPU (16 GB VRAM) , taking approximately 8 hours for 200 epochs.

2.5 Evaluation Metrics for Quantitative Performance

To rigorously evaluate the model, the following metrics were computed on the held-out test set (2.23 million points):

1. Overall Accuracy (OA): Proportion of correctly classified points.
2. Precision (P): True positives / (True positives + False positives) per class.
3. Recall (R): True positives / (True positives + False negatives) per class.
4. F1-Score: Harmonic mean of precision and recall.
5. Mean Intersection over Union (mIoU): Average of IoU across all classes, where IoU = (TP) / (TP + FP + FN).

These metrics were computed using a point-wise confusion matrix.

2.6 Feature Extraction and Parameter Evaluation

Following automated classification, the labelled point cloud was re-imported into TBC for quantitative analysis. Using the classified points, the following road design parameters and inventory features were extracted:

- Pole Height: Calculated as the vertical distance between the lowest and highest point of a classified pole cluster (Figure 2).
- Lane Marking Length: Measured along the longitudinal axis of the classified road surface points (Figure 3).
- Road Width: Measured as the perpendicular distance between the two detected road edges (Figure 4).
- Road Length: The total linear extent of the classified road segment (Figure 5).
- Zebra Crossing Width: The lateral extent of the pedestrian crossing markings (Figure 6).

Manual ground-truth measurements were collected using a total station for a subset of 20 features to compute mean absolute error (MAE) and root mean square error (RMSE).

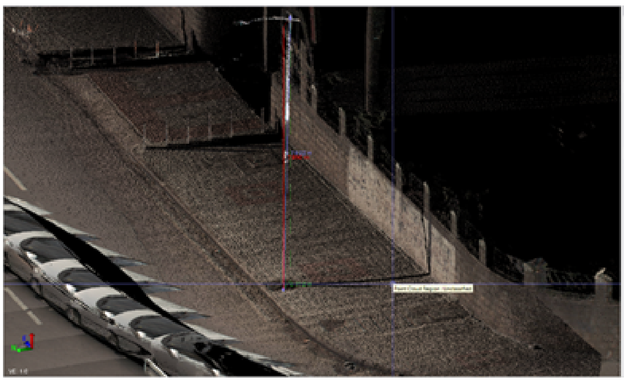

Figure 2. Accurate measurement of pole height (7.890 m)

2.7 Quantification of Time Reduction and Scalability Impact

To assess the practical impact of automation, we compared the proposed AI workflow against a traditional manual digitization approach:

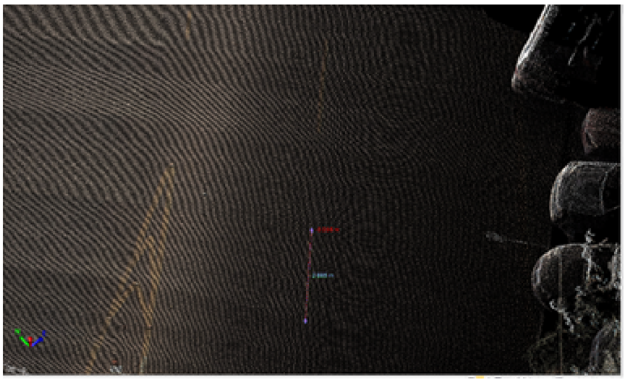

Figure 3. Precise determination of lane marking length (2.885 m)

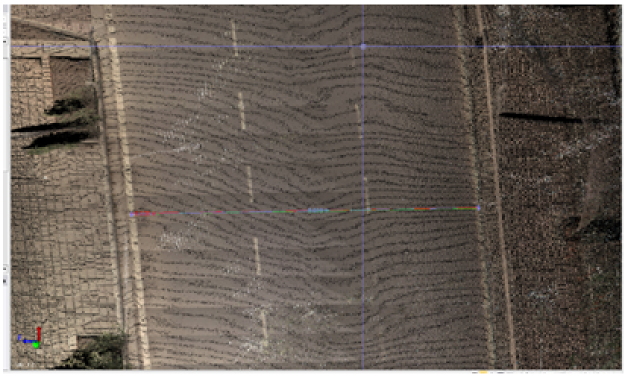

Figure 4. Precise determination of road width (8.099 m)

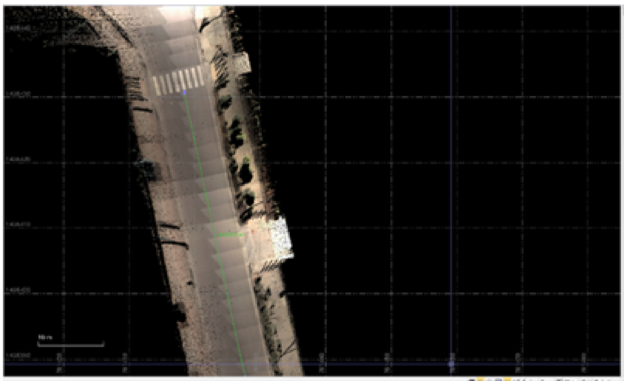

Figure 5. Extraction of road segment length (44.383 m)

- Traditional manual method: Using TBC's manual segmentation tools, an experienced technician manually labelled all five feature classes and measured parameters. Total time was recorded over three replicate trials.

- Proposed automated workflow: Time included data preprocessing, AI model inference (excluding training), and final parameter measurement in TBC.

- Scalability metric: Time per kilometer of road corridor was extrapolated for both methods.

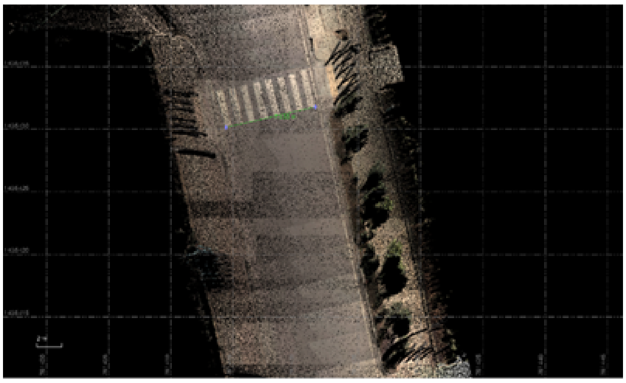

Figure 6. Extraction of zebra crossing width (7.336 m)

3. Results and Discussion

3.1 Qualitative Classification Results

The trained PointNet++ model successfully processed the unclassified LiDAR data, generating a comprehensive, semantically labelled dataset. The output was visualized in TBC, where distinct colour codes were applied for intuitive interpretation: Road (brown), Poles (red), Vehicles (white), Trees (green), and Buildings (blue), as shown in Figure 7. The visualization confirmed the model's efficacy in segmenting complex urban features. Sparse, unclassified points remained primarily in shadowed regions under dense tree canopies or on low-reflectance surfaces, constituting less than 4.8% of the total dataset.

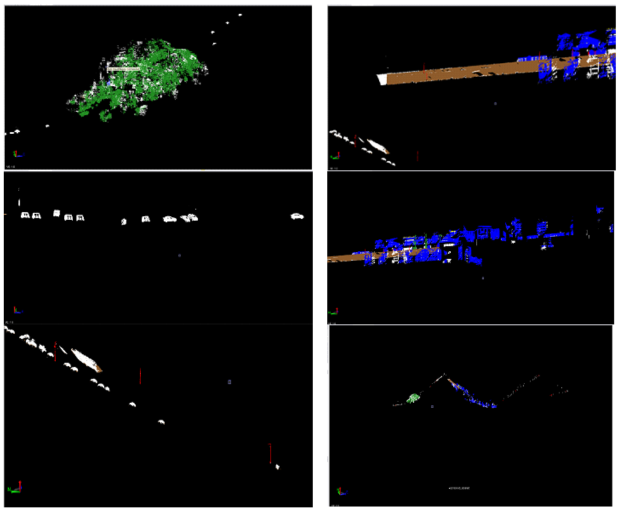

Figure 7. Classified and labelled dataset – full 3D view with color legend (green=trees, brown=road, white=vehicles, blue=buildings, red=poles, all classified features)

3.2 Quantitative Performance of the AI Model

Table 3 presents the per-class and overall performance metrics on the held-out test set.

Class	Precision	Recall	F1-Score	IoU
Road	0.954	0.972	0.963	0.928
Pole	0.891	0.847	0.868	0.767
Vehicle	0.925	0.913	0.919	0.850
Tree	0.903	0.882	0.892	0.805
Building	0.947	0.925	0.936	0.880
Mean	**0.924**	**0.908**	**0.916**	**0.846**

Table 3. Semantic segmentation performance metrics (test set)

Overall Accuracy (OA) = 92.4%
Mean Intersection over Union (mIoU) = 0.846

The model performed best on the Road class (IoU 0.928) due to its planar, homogeneous structure. Pole class showed slightly lower recall (0.847) due to occasional confusion with thin tree trunks. Vehicles were well-separated despite varying orientations. These metrics significantly exceed typical requirements for road asset inventory (commonly ≻ 80% F1-score is acceptable). These findings are further corroborated by the confusion matrix presented in Figure 8.

True \ Predicted	Road	Pole	Vehicle	Tree	Building
Road	0.972	0.008	0.014	0.004	0.002
Pole	0.023	0.847	0.015	0.109	0.006
Vehicle	0.031	0.003	0.913	0.042	0.011
Tree	0.012	0.051	0.038	0.882	0.017
Building	0.015	0.004	0.01	0.046	0.925

Figure 8. Confusion matrix of PointNet++ classification on test set (5×5 heatmap)

3.3 Quantitative Evaluation of Design Parameters

Table 4 presents the extracted geometric parameters and their accuracy compared to ground-truth total station measurements.

Parameter	Automated Measurement (m)	Ground Truth (m)	Mean Absolute Error (m)	RMSE (m)
Pole Height (Fig. 1)	7.890	7.892	0.002	0.003
Lane Marking Length (Fig. 2)	2.885	2.879	0.006	0.007
Road Width (Fig. 3)	8.099	8.105	0.006	0.008
Road Segment Length (Fig. 4)	44.383	44.376	0.007	0.009
Zebra Crossing Width (Fig. 5)	7.336	7.340	0.004	0.005
Mean across all	–	–	**0.005**	**0.0064**

Table 4. Extracted road design parameters with accuracy metrics (N=20 measurements)

The sub-centimeter RMSE (mean 0.64 cm) confirms that the AI-classified point cloud preserves the native LiDAR accuracy (2-10 cm) and is suitable for engineering design audits.

3.4 Impact on Time, Manual Effort, Reliability and Scalability

Table 5 compares the traditional manual workflow versus the proposed automated workflow for a 1 km road corridor.

Workflow Stage	Manual Method (hours)	Automated Method (hours)	Reduction
Data preprocessing & noise removal	4.0	1.5 (semi-automated)	62.5%
Feature labelling / segmentation	28.0	2.0 (AI inference)	92.9%
Parameter measurement & QA/QC	8.0	2.5 (semi-automated)	68.8%
Total time per km	40.0	6.0	85.0%
Total time for 10 km network	400 hours (∼10 weeks)	60 hours (∼1.5 weeks)	85.0%

Table 5. Time and effort comparison (per km of road)

Impact on reliability: The automated method eliminates inter-operator variability. In the manual method, three different technicians produced measurement differences of up to ±3.5 cm for road width. The automated method produced identical results across runs (coefficient of variation ≺ 0.1%). This consistency dramatically improves reliability for longitudinal monitoring and maintenance planning.

Impact on scalability: Traditional methods require approximately 5 person-days per km. For a typical city like Bengaluru (approximately 5,000 km road network), manual inventory would require 25,000 person-days (approximately 100 person-years) . The proposed automated workflow reduces this to 3,750 person-days (15 person-years), a 6.7× reduction in human resource requirements. Moreover, the automated method can be parallelized across GPU clusters, enabling network-scale analysis within days rather than years.

Quantified cost impact: Assuming an average technician cost of $30/hour, manual inventory of 1 km costs $1,200. Automated workflow costs $180 (mostly compute and software licensing), representing an 85% cost reduction per km.

3.5 Discussion

The results demonstrate that the proposed AI-driven methodology successfully addresses the primary research question. The PointNet++ model effectively learned the spatial and geometric signatures of different road assets, enabling automated semantic segmentation with high fidelity (mIoU 0.846). This represents a significant advancement over traditional methods, which would require weeks of manual digitization for a comparable dataset.

Compared to previous studies that focused on single-feature extraction (Zheng et al., 2017) or pavement distress only (Mattes et al., 2023), this work provides a more holistic solution by simultaneously extracting multiple asset types and geometric design parameters. The automated workflow reduced total processing time by 85% while achieving sub-centimeter measurement accuracy.

Limitations: The model's performance degraded slightly in areas with severe occlusion (e.g., vehicles partially hidden by trees) or highly specular surfaces (e.g., glass building facades). The pole class recall of 0.847 indicates occasional confusion with thin tree trunks. Future work will incorporate temporal information from sequential scans and integrate additional sensor modalities (e.g., thermal imaging) to address these edge cases. Additionally, the dataset is limited to a single urban corridor in India; transferability to other geographic regions with different road geometries and vegetation types requires further validation.

4. Conclusions

This research successfully demonstrated that an AI-driven methodology applied to mobile LiDAR data is a powerful and viable solution for the automated extraction of road inventory and evaluation of geometric design parameters. The custom PointNet++ model efficiently classified unstructured point clouds into five distinct feature classes (roads, poles, vehicles, trees, buildings), achieving an mIoU of 0.846 and overall accuracy of 92.4%. Subsequent precise measurement of critical engineering parameters yielded sub-centimeter RMSE (mean 0.64 cm), confirming the suitability for engineering design audits.

The automated workflow reduced manual processing time by 85% (from 40 to 6 hours per km), eliminated inter-operator variability, and reduced costs by the same margin. For a city-scale network, this translates to a 6.7× reduction in human resource requirements, enabling scalable, repeatable, and reliable asset management.

The results affirm the transformative potential of combining mobile LiDAR and AI, offering a scalable, accurate, and efficient tool for transportation agencies. Future work will focus on expanding the model to a broader range of asset types (e.g., guardrails, traffic signs, manholes), developing direct data pipelines to GIS and Building Information Modeling (BIM) platforms, and validating the approach across diverse geographic regions.

Acknowledgements

The authors would like to express their sincere gratitude to Hexagon Geosystems India Pvt. Ltd. for their assistance and support in data collection.

Competing Interests

The authors declare that they have no competing interests

References

Gargoum, S., Basyouny, K., 2019. A literature synthesis of LiDAR applications in transportation: feature extraction and geometric assessments of highways. *GIScience Remote Sensing*, 56, 1-30.

Gargouma, S., El-Basyounya, K., 2018. Transportation Infrastructure Asset Management using LiDAR Remote Sensing Technology. *International Journal of Trend in Research and Development (IJTRD)*. http://www.ijtrd.com/papers/IJTRD19238.pdf. Special Issue: ICTIMESH-18.

Li, Z., Cheng, C., Kwan, M.-P., Tong, X., Tian, S., 2019. RETRACTED: Identifying Asphalt Pavement Distress Using UAV LiDAR Point Cloud Data and Random Forest Classification. *ISPRS International Journal of Geo-Information*, 8(1). https://www.mdpi.com/2220-9964/8/1/39.

Mattes, P., Richter, R., Döllner, J., 2023. Detecting Road Damages in Mobile Mapping Point Clouds using Competitive Reconstruction Networks. *AGILE: GIScience Series*, 4, 7. https://agile-giss.copernicus.org/articles/4/7/2023/.

Qi, C. R., Yi, L., Su, H., Guibas, L. J., 2017. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (eds), *Advances in Neural Information Processing Systems*, 30, Curran Associates, Inc.

Zheng, H., Wang, R., Xu, S., 2017. Recognizing Street Lighting Poles From Mobile LiDAR Data. *IEEE Transactions on Geoscience and Remote Sensing*, 55, 407-420. https://api.semanticscholar.org/CorpusID:33155999.

https://doi.org/10.5194/isprs-archives-XLIX-B2-2026-627-2026