Updating Forestry Road networks in Ontario Using Single Photon LiDAR and Deep Learning-enhanced algorithms

Hamid Ghanbari[1], Emmanuelle Baby-Bouchard[1], Catherine Chagnon[1], Alexandre Morin-Bernard[1], Alexis Achim[1]

[1]Department of Wood and Forest Sciences, Université Laval, Québec, Canada, hamid.ghanbari.1@ulaval.ca, emmanuelle.baby-bouchard.1@ulaval.ca, catherine.chagnon.2@ulaval.ca, alexandre.morin-bernard@sbf.ulaval.ca, alexis.achim@sbf.ulaval.ca

Keywords: Forest Management, Large-Scale Mapping, U-Net Models, Airborne Laser Scanning, 3D Point Cloud.

Abstract

Spatially accurate forestry road networks are essential for effective forestry operations, sustainable resource management, and conservation. Current forestry road databases in Ontario have significant location errors due to limitations and human errors associated with conventional road delineation approaches such as GPS-based field surveys and photointerpretation. A previously developed algorithm, which used airborne laser scanning (ALS) data, successfully corrected road locations in Quebec. However, its design limited its application in other landscapes, ALS instruments, and road construction and maintenance practices. This study advances that algorithm by integrating a deep learning component to improve its robustness and scalability for diverse forest conditions. A hybrid workflow combines the original friction-based conductivity surface with a road probability surface generated by an Attention Residual U-Net model trained on 11 LiDAR-derived features using road segments from five forest sites in Quebec. The enhanced workflow was applied to two forest management units in Ontario: Nipissing and Dryden. The results showed significant improvement in road alignment when compared to the existing provincial data and the outputs from the earlier automated approach. The deep learning-enhanced algorithm lowered mean positional error by 78% (from 9.36 m to 2.07 m) and increased the proportion of road centerline points within 3 m of the reference from 66.7% to 87.2%. These improved centerline accuracies will further support a scalable tool for rapid and accurate forestry road network mapping, which in turn will aid sustainable forest management and conservation planning at both provincial and national scales.

1. Introduction

In Canada, forestry road networks managed at the provincial level have been frequently characterized by significant positional inaccuracies and incomplete attributes (Morley et al., 2024). These inaccuracies largely originate from conventional road delineation methods that require photo interpreters to manually trace road alignments from aerial or high-resolution satellite imagery. These limitations are compounded by the gradual development of road mapping over many years, whereby some road segments may have been abandoned, partially delineated, or never recorded. Such methods are not only costly and labour-intensive, but also inconsistent, as they must contend with images of varying spatial resolution and quality, atmospheric interference, and the inherent difficulty of detecting roads in some regions such as beneath dense forest canopy using passive optical sensors (David et al., 2009). An alternative approach using GPS units mounted on vehicles is similarly limited, as coordinate accuracy and fixed recording frequency can produce rough edges and miss sharp turns. These challenges are particularly pronounced in the vast forests of Ontario, Canada, where positional discrepancies in the existing road network are significant enough to necessitate systematic correction (Morley et al., 2024). Addressing these inaccuracies is therefore critical, as reliable road networks support sustainable forest management, efficient forestry operations, and conservation activities across Canada's managed forest landscapes (Waga et al., 2020; Hoffmann et al., 2024).

Optical remote sensing imagery has demonstrated strong potential for road detection through its high spatial resolution and spectral richness. However, its effectiveness is significantly limited by cloud cover, shadows, interferences from spectrally similar background area, and the inability of passive sensors to penetrate dense forest canopy (Lu and Weng, 2025). On the other hand, Synthetic Aperture Radar (SAR) data offers a cloud-independent, canopy-penetrating alternative. However, its inherent speckle noise, complex backscatter response from forested terrain, relatively coarse spatial resolution compared to the narrow width of forestry roads, and sensitivity to surface roughness rather than spectral properties make accurate road delineation challenging in densely vegetated environments (Chen et al., 2022). Three-dimensional remotely sensed data, particularly airborne laser scanning (ALS) due to its large-scale capacity, offer a compelling alternative. By emitting active laser pulses and recording return signals, ALS generates dense 3D point clouds that capture both the Earth's surface and vegetation structure with high precision. A significant advantage of ALS is its ability to penetrate vegetation gaps and provide accurate terrain measurements regardless of illumination conditions, making it particularly valuable for forestry applications (Azizi et al., 2014).

A previous algorithm based on the linear-mode ALS has successfully been applied to correct forestry road positions and estimated road-level metrics across several forest sites in Quebec province, Canada (Roussel et al., 2022). The algorithm converts 3D point clouds into a set of 2D feature layers, each capturing a distinct characteristic of the scene. These features are then combined using a linear weighting scheme to generate a conductivity surface, which is subsequently used to derive the least-cost path and determine the corrected road alignment. Although it generated highly accurate vectorial road networks across Quebec and demonstrated transferability to selected Ontario forest management units (Morley et al., 2024), it revealed significant performance variability when applied to Ontario's broader forests. The limitations come from the algorithm's parametric design, which requires manual tuning of empirically determined weights for different LiDAR-derived metrics based on forest and road conditions. Additional challenges arise from

https://doi.org/10.5194/isprs-archives-XLIX-B1-2026-421-2026 | © Author(s) 2026. CC BY 4.0 License.

the transition to Single Photon LiDAR (SPL) instruments, which were used for data collection in Ontario and differ in point cloud characteristics from the linear-mode sensors on which the initial algorithm was developed in Quebec. Moreover, Ontario's heterogeneous landscapes and road construction standards, where roads may lack discernible shoulders or curbs, introduced new challenges.

The rapid progress of deep learning (DL) approaches has led to major improvements in automated road feature extraction and segmentation of remote sensing data, without requiring prior weighting of feature importance (Lu and Weng, 2025). DL models can capture non-linear interactions among input features, learn generalizable spatial patterns from training data, and adapt effectively to diverse landscape conditions without the need for manual parameter recalibration. In particular, Attention-based U-Net architectures have shown strong performance in automated road detection tasks (Botelho Jr et al., 2022; Lidberg, 2025; Oktay et al., 2018).

This study integrates a DL component into the existing ALS-based road correction framework to improve its robustness and scalability across Ontario's diverse forest environments. Building on previous work in Quebec, a hybrid workflow is developed that combines the original friction-based conductivity surface with a pixel-level road probability surface generated by a DL model trained on ALS-derived features from Quebec forestry road datasets. The enhanced model was then evaluated across two Forest Management Units (FMUs) in Ontario, encompassing diverse forest types, terrain conditions, and road management approaches. The open-source tool developed through this work aims to support rapid and reliable forestry road network delineation at a province-wide scale, with broad applications in sustainable forest management and ecological conservation.

2. Materials and Methods

2.1 Study Area

This study focuses on two FMUs in the province of Ontario, Canada including Nipissing Forest (NF; 11 395 km²) and Dryden Forest (DF; 3 071 km²). These regions encompass a wide range of forest types, terrain conditions, and road management practices, making them representative of the diverse environments encountered across Ontario's boreal and mixed-wood forests (Figure 1).

2.2 Airborne Single Photon LiDAR

The SPL data under leaf-on conditions were acquired using an airborne platform and the Leica SPL100 system during years 2019-2020 for NF and 2021-2022 for DF, and pre-processed following the methods of Gluckman (2016). SPL technology enables data acquisition at higher flight altitudes compared to conventional linear-mode LiDAR, resulting in greater area coverage per flight while maintaining comparable point cloud densities (White et al., 2021). The SPL point clouds were used to derive a suite of terrain and forest structure metrics at 1 m resolution, such as the multi-directional hillshade, Digital Surface Model (DSM), Canopy Height Model (CHM), and terrain slope.

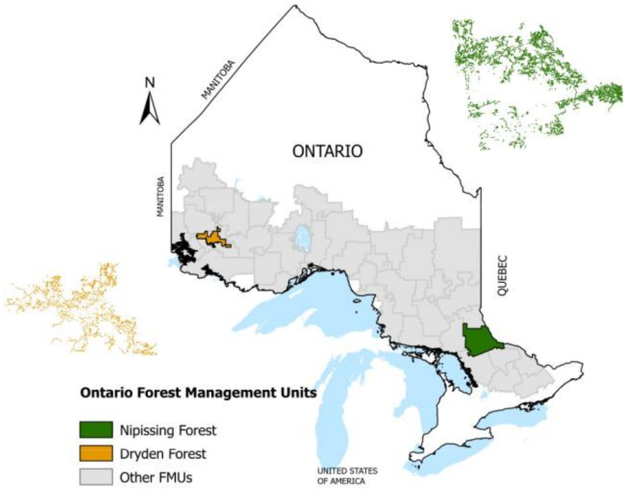

Figure 1. The location of the two FMUs examined in this study and their corresponding existing inaccurate road networks.

2.3 Existing Road Network Data

Existing forestry road network data for the two FMUs were obtained from provincial databases managed by the Ontario Ministry of Natural Resources and Forestry (MNRF). These vector datasets represent road centerlines that were originally delineated using manual photointerpretation of aerial and high-resolution satellite imagery. The positional accuracy of these data is variable and, in many areas, insufficient for operational planning, with deviations from true road positions that can reach up to 80 meters. These inaccurate road centerlines were used as prior spatial constraints in the road correction algorithm.

2.4 Training and Validating Data

A cross-regional training and validation design was implemented, whereby the model is trained on Quebec roads and evaluated on the structurally different landscapes and road networks of Ontario forests. This allowed us to assess the generalizability of the proposed algorithm beyond its training domain. A total of 890 manually corrected road segments from five forest sites in Quebec, originally compiled for the Roussel et al. (2022) framework, were used to train the deep learning model. For evaluation, an expert operator manually digitized a subset of road segments across both Ontario FMUs; 128 road segments from the DF and 46 road segments from the NF. To ensure geometric accuracy, the digitization process was guided by multiple LiDAR-derived features, including ground point density and the CHM, along with high-resolution aerial imagery of the respective regions. Validation road segments were carefully selected to represent the diverse landscape conditions present across the two FMUs, including variations in terrain complexity, forest cover density, and road management type.

2.5 Road Correction Workflow

This study builds upon the road correction framework introduced by Roussel et al. (2022), which derives a conductivity surface from LiDAR-derived predictor layers and applies a least-cost path algorithm to update road centerline positions. A key limitation of the original algorithm is the fixed weighting of predictor layers when generating the final conductivity surface. Because the importance of predictor raster layers varies across regions, this approach often led to inaccurate road mapping in heterogeneous landscapes. To address this, we combined the linear combination of predictors with the conductivity surface derived from an Attention Residual U-Net (Attention ResUNet) model trained to estimate road probability for each pixel (Figure 2). The model used 11 input features, including multi-directional hillshade, slope, DSM, terrain roughness, surface roughness, Sobel-filtered edge layer, laser return intensity, CHM, point count between 0.5–3 m, ground point density percentage, and a proximity layer indicating distance to the existing road centerline (see Figure 3 for a representative subset of these input features). The enhanced workflow retains the least-cost path correction of the original algorithm but replaces the fixed linear weighting with a data-driven conductivity surface predicted by a deep learning model. The final conductivity surface is computed as a weighted combination of the original parametric conductivity surface and the DL-predicted road probability raster, allowing the model to adaptively emphasize the most informative predictor layers across different landscape contexts (Figure 2A).

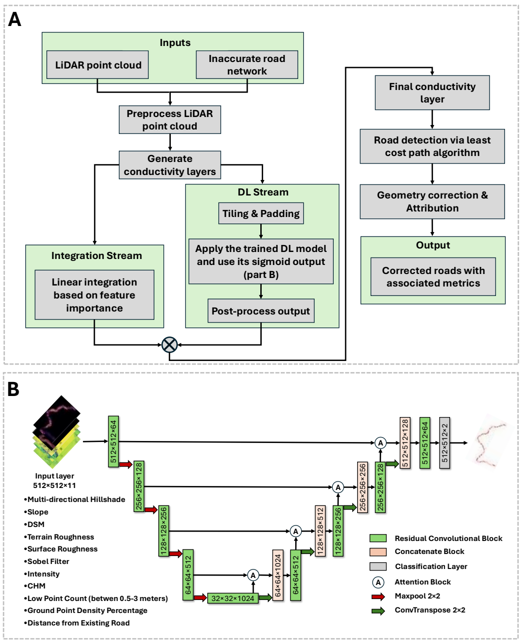

Figure 2. Road correction workflow: (A) overall algorithm structure; (B) Attention ResUNet architecture used in this study (adapted from Ghaznavi et al. (2022)).

2.5.1 Model Development: An Attention ResUNet was adopted as the DL backbone for pixel-wise road probability estimation. This design choice is motivated by the demonstrated advantages of attention-augmented architectures in capturing long-range dependencies and generalizing across diverse geographic contexts (Lu and Weng, 2025), which aligns with our cross-regional training and validation strategy. The attention mechanism in a ResUNet backbone enables the model to selectively weight spatially relevant information across the full extent of the multi-channel LiDAR-derived input. The architecture follows an encoder–decoder structure with skip connections, building on the standard U-Net design (Ronneberger et al., 2015). Residual blocks in the encoder enable the training of deeper networks by mitigating vanishing gradient effects, while attention gates incorporated into the skip connections allow the model to selectively weight spatial features, suppressing activations from non-road regions and focusing learning on road-relevant patterns (Ghaznavi et al., 2022).

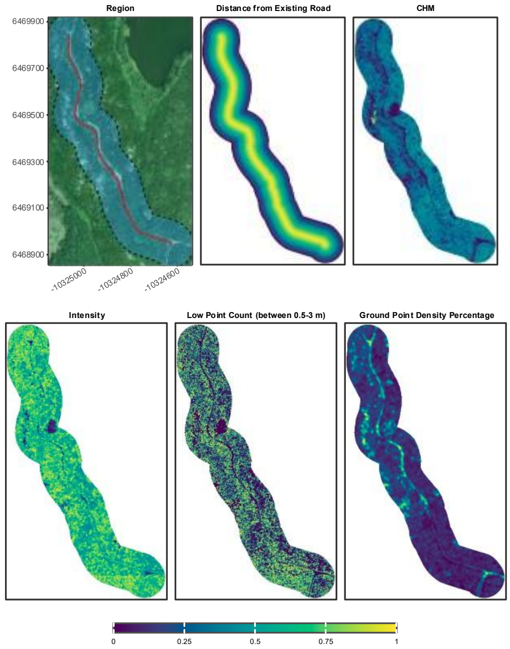

Figure 3. An example of LiDAR-derived features and their corresponding on-the-ground region. The red curve in the region panel indicates the existing inaccurate road centerline, and the blue shaded area represents an 80-metre buffer around the road. Spatial coordinates are in UTM Zone 15N (metres, GRS 1980 ellipsoid). All feature layers are normalized to [0, 1] using min–max normalization.

For each road segment, paired 512×512-pixel raster tiles of the 11 LiDAR-derived as input features and binary road mask as labels were extracted. The binary masks were generated by rasterizing the manually corrected road centerlines, where road pixels were assigned a value of 1 and all remaining pixels a value of 0. Tiles were centered on road centerlines with a fixed buffer of 80 m and resampled to a consistent spatial resolution of 1-m pixels. Data augmentation strategies including vertical and horizontal flipping, random rotation, and channel dropout (randomly zeroing individual input channels with a probability of 0.2, with a maximum threshold of 3 channels zeroed per data) were also applied during training to improve model generalization.

The model was trained using binary focal cross-entropy as the loss function, which addresses class imbalance between road and non-road pixels by down-weighting well-classified background samples and focusing training signal on difficult, ambiguous road pixels. However, since the road pixels barely cover the image area, the loss function was applied only on the buffered area and not the surroundings. The Adam optimizer was used for gradient-based parameter updates with an initial learning rate determined through hyperparameter optimization. Training was conducted for 200 epochs using early stopping with a patience of 30 epochs based on validation loss. To improve the robustness of validation, test-time augmentation was applied by averaging sigmoid predictions across six augmentations (i.e., original, horizontal flip, vertical flip, and three 90° rotations).

Hyperparameter optimization was performed using the Optuna framework (Ozaki et al., 2025), a black-box optimization library that employs Tree-structured Parzen estimator sampling to efficiently search the hyperparameter space. The following hyperparameters were optimized: learning rate (range: 1e-5 to 1e-

2), batch size (4, 8, 16, or 32), and focal loss parameters alpha (0.80 to 0.98) and gamma (1 to 4). Each trial was evaluated on a held-out validation subset using validation loss, and a total of 50 trials were conducted. The best-performing hyperparameter configuration was selected for final model training.

Model training was carried out on the Compute Canada (now Digital Research Alliance of Canada) high-performance computing infrastructure, using Narval cluster. Each training run was allocated 2× NVIDIA A100 40 GB, 16 CPU cores, and 64 GB of RAM. The model was implemented in Python using the PyTorch framework, with data loading handled via a custom rasterio-based DataLoader incorporating per-band min-max normalization. Training scripts and model weights will be made available as part of the open-source release.

2.6 Validation

Road alignment accuracy was assessed by comparing road centerline datasets (i.e., original parametric approach and DL-enhanced algorithm) against the manually digitized reference segments in Ontario. Positional errors were computed by sampling points at 5 m intervals along each corrected centerline and measuring the perpendicular distance from each sampled point to the nearest segment in the ground truth reference. Five metrics were derived from this error distribution: the median positional error (p50), the 95th percentile positional error (p95), the 99th percentile positional error (p99), the percentage of sampled points with an error below 3 m, and the percentage of sampled points with an error below 5 m. In addition, F1-score and IoU (Intersection over Union) were computed over buffered road polygons at a 1 m tolerance. Together, these metrics capture both the typical positioning accuracy and the tail behaviour of the error distribution, providing a comprehensive assessment of improvement across the two databases (i.e., corrected roads from the original parametric algorithm and the DL-enhanced algorithm) and across the diverse landscape conditions represented in the two Ontario FMUs.

3. Results

3.1 Hyperparameter Optimization

Hyperparameter optimization via Optuna converged after 50 trials. The optimal configuration identified for the Attention ResUNet included a learning rate of 1×10^{-3}, a batch size of 8, and focal loss parameters of alpha = 0.97 and gamma = 2.0, with a learning rate reduction factor of 0.5 applied when validation loss plateaued for 10 consecutive epochs and early stopping triggered after 30 epochs without improvement. The optimization process demonstrated that the learning rate was the most influential hyperparameter, followed by the focal loss function's alpha parameter, consistent with findings in other segmentation tasks on imbalanced remote sensing datasets.

3.2 Model Training Performance

The trained the Attention ResUNet was evaluated on the held-out Ontario validation set comprising ~349 tiles. The model achieved a precision (the proportion of predicted road pixels that are actually roads) of 0.24, recall (the proportion of actual road pixels correctly identified by the model) of 0.92, and an F1-score (harmonic mean of precision and recall) of 0.38 on the validation set, with a Matthews Correlation Coefficient (MCC; an overall measure of prediction quality that accounts for class imbalance) of 0.47, demonstrating strong sensitivity to road pixels while exhibiting a notably higher rate of false positives. Figure 4 presents three representative examples of CHM, point density, ground-truth mask, and DL-predicted sigmoid probability rasters, which help illustrate the sources of this low precision (high false positives). In all examples, the model correctly delineates the main road corridor but also activates strongly over an adjacent forest clearing, where low canopy height in the CHM and elevated point density mimic the structural signature of a road surface. Occasionally, false positives appear along harvest cutlines that were not labelled in the ground-truth mask, suggesting that the model detects real linear disturbances, or the sigmoid output assigns high probability to a neighboring region where the low-vegetation CHM pattern resembles a road, despite no corresponding ground-truth label. Taken together, these examples suggest that low precision stems from overlapping factors of structural ambiguity between roads and other low-canopy linear features in LiDAR-derived inputs and the model's tendency to over-generalize learned road signatures to visually similar non-road surfaces.

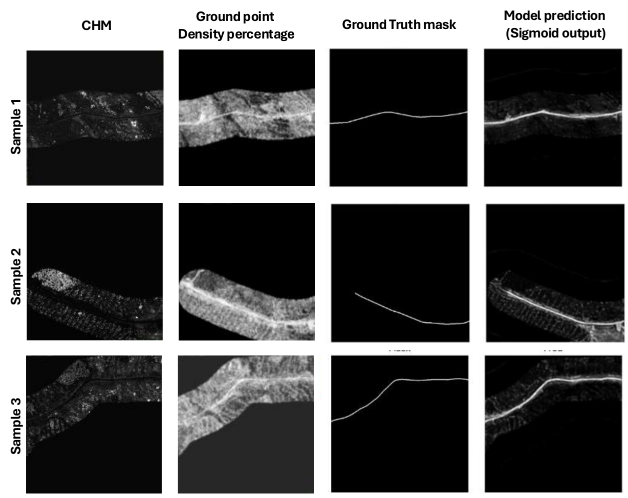

Figure 4. Representative examples of model inputs and predictions for three validation tiles from the Quebec hold-out set. Each row displays, from left to right: the CHM, ground point density percentage, ground-truth road mask, and DL-predicted road probability (sigmoid output).

3.3 Road Centerline Correction in Ontario

3.3.1 Qualitative Assessment: Visual comparison of the corrected road centerlines with high-resolution aerial imagery confirmed substantial improvement in positional accuracy relative to the original provincial road network across both Ontario FMUs. Figure 5 shows representative examples of road correction outcomes in the DF, where the original road centerlines exhibited positional offsets of several meters from the true road. The enhanced algorithm successfully repositioned centerlines within the road corridor in both open-terrain and canopy-covered road sections. Across the study area, the DL-enhanced conductivity surface proved particularly effective in resolving ambiguous road positions in areas where individual LiDAR metrics provided weak or conflicting signals, such as roads beneath partial canopy closure or roads in terrain with irregular surface features.

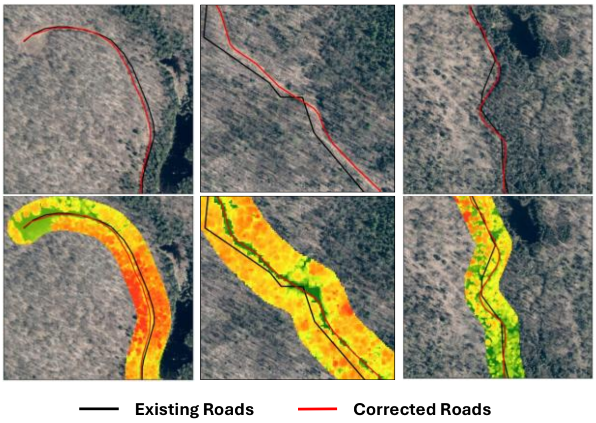

Figure 5. Qualitative examples of road centerline correction in DF using the DL-enhanced algorithm. Upper panels show existing and corrected centerlines overlaid on high-resolution aerial imagery and bottom panels present their respective CHM layers.

3.3.2 Quantitative Assessment: Quantitative evaluation was performed for DF and NF using the corrected reference road dataset comprising approximately ~110 km of road centerlines (Table 1). The DL-enhanced algorithm substantially outperformed the original baseline across all evaluated metrics. Compared to the original parametric algorithm, it reduced the mean positional error from 9.36 m to 2.07 m (78% reduction) and lowered the median error (p50) from 1.53 m to 1.00 m. The improvement is particularly evident in the tail of the error distribution: the 95th percentile positional error dropped from 51.43 m to 7.18 m, indicating that the DL-enhanced approach nearly eliminated the large systematic offsets that affected a subset of roads in the parametric output (Figure 6). This pattern was also reflected in the proportion of sampled points falling within 3 m of the reference which increased from 66.7% to 87.2%, and within 5 m from 72.9% to 92.7%. Improvements in geometric overlap metrics were consistent with the positional error results. The F1-score increased from 0.38 to 0.50 and IoU from 0.28 to 0.35, reflecting both better positional alignment and improved completeness of the corrected centerlines.

Table 1. Road alignment accuracy metrics for the original parametric algorithm and the DL-enhanced algorithm compared against manually digitized reference segments in Ontario.

Metric	Original Parametric Algorithm	DL-Enhanced Algorithm
Buffered Polygon Overlap (1 m tolerance)		
F1-Score	0.38	**0.50**
IoU	0.28	**0.35**
Positional Error (m)		
Median / p50	1.53	**1.00**
p95	51.43	**7.18**
p99	84.91	**23.90**
Mean	9.36	**2.07**
Points Within Threshold (%)		
Points < 3 m	66.7%	**87.2%**
Points < 5 m	72.9%	**92.7%**

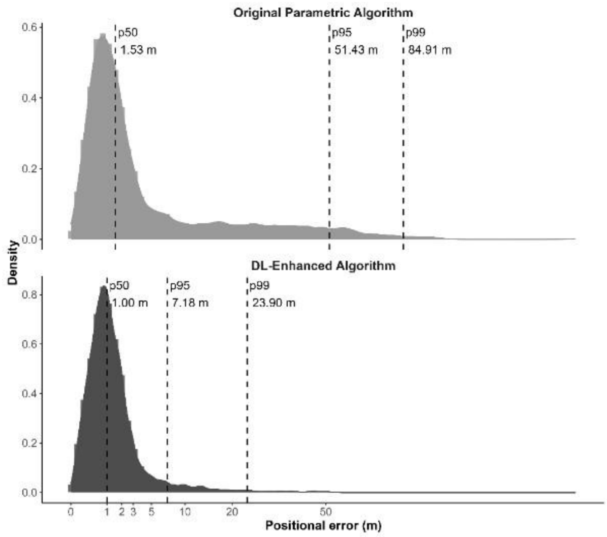

Figure 6. Positional error distributions for the original parametric algorithm and the DL-enhanced algorithm. The x-axis is on a square-root scale to improve visibility of the low-error range.

4. Discussion

4.1 Contribution of the Deep Learning Component

The integration of the Attention ResUNet into the road correction workflow substantially improved the generalizability of the algorithm across Ontario's heterogeneous forest landscapes. The primary limitation of the original parametric algorithm was its reliance on fixed, empirically calibrated weights for combining LiDAR-derived predictor layers, which assumed consistent relative feature importance across all forest contexts. In Ontario, variations in SPL point cloud density, terrain relief, forest composition, and road construction standards created conditions where this assumption was frequently violated, and in areas of complex terrain or ambiguous road signatures, the rule-based correction could deviate substantially from the true centerline, producing a small subset of segments with large systematic offsets that disproportionately inflated the high-percentile errors. This is reflected in the p95 and p99 of the original parametric output reaching 51.43 m and 84.91 m, respectively, compared to 7.18 m and 23.90 m for the DL-enhanced algorithm. The DL model's attention mechanism effectively addressed this by learning context-dependent features from the Quebec training data and generalizing these patterns across novel forest environments. The suppression of extreme positional deviations is consistent with the hypothesis that the model implicitly learns road geometry constraints during training, such as typical road curvature and continuity, which act as a soft regularizer against implausible offsets even in challenging landscape conditions where the parametric approach struggles.

The choice of binary focal cross-entropy as the loss function alongside limiting it to the buffered area rather than the whole receptive field, proved important for managing the strong class imbalance between road and non-road pixels, which is a persistent challenge in remote sensing road extraction tasks (Lu and Weng, 2025). By reducing the contribution of easily classified background pixels, the loss function directed model capacity toward learning the subtle signatures of narrow forestry roads, which might otherwise be overlooked when training over the entire receptive field.

4.2 Transferability from Quebec to Ontario

A key challenge in this study was training a model exclusively on Quebec data and deploying it across Ontario's forest conditions without fine-tuning. The effective transfer of learned features indicates that the model successfully captured generalizable patterns in the input data. Nevertheless, exposing the model to a broader range of road-landscape features during training would further strengthen its ability to generalize across diverse forest contexts. This aligns with findings in other remote sensing transfer learning studies, where models trained on ecologically diverse source domains have demonstrated effective cross-region generalization (Liu et al., 2024).

Certain roads in Ontario presented greater challenges for the algorithm. In particular, roads traversing recently harvested cutblocks, where the absence of canopy cover and the similar reflectance of bare soil made road surfaces difficult to distinguish, where the conductivity signal was too weak to reliably guide the pathfinder (Figure 4). These cases are likely to benefit from the inclusion of Ontario-specific training samples in future model iterations, which is facilitated by the ongoing production of a manually corrected Ontario road reference dataset.

The transition from linear-mode LiDAR to SPL also introduced differences in point cloud characteristics, including a higher proportion of noise returns at low incidence angles, differences in intensity calibration relative to linear-mode sensors, and a more diffuse return distribution in dense canopy conditions that slightly reduced the contrast between road and non-road ground-return densities. These differences have made the original algorithm fail; however, the 11 LiDAR-derived feature layers proved robust inputs to the model as they abstract raw point cloud properties into physically interpretable raster metrics that are comparable across sensor types.

4.3 Scalability and Computational Efficiency

The integration of the DL component introduced a five-fold increase in road alignment processing time relative to the original parametric approach. However, since the correction workflow is intended to be applied once per dataset rather than in real-time operational contexts, this additional computational cost is unlikely to pose a practical constraint. For training, the use of Compute Canada's high-performance computing infrastructure facilitated the training process, as processing the large SPL data required approximately 12 hours of GPU time for DL inference. For the deployment of the algorithm, the parallel processing of the workflow, which tiles the study area into overlapping processing units and reassembles outputs, allowed effective scaling without memory limitations. This design is directly transferable to province-wide deployment across Ontario's full managed forest, including more than 300 000 km of forestry roads.

The planned release as an open-source R package will make the workflow accessible to provincial forest managers and researchers without requiring specialized machine learning expertise. The modular design, separating DL inference, conductivity surface generation, and least-cost path correction, allows users to run individual components with locally available data or to integrate updated training data as Ontario-specific reference roads become available. This design aligns with the need for operational-grade, reproducible geospatial tools in forest management.

4.4 Limitations and Future Work

The deep learning model was trained exclusively on Quebec Road data, and while good generalization to Ontario was observed, the absence of Ontario-specific training samples may limit performance for road types that are underrepresented in the Quebec dataset. The production of a manually corrected Ontario reference dataset will enable fine-tuning of the model on target-domain data and quantitative benchmarking across additional FMUs.

The current validation is primarily qualitative and quantitative for two of the FMUs; however, a full quantitative assessment across more complex FMUs will be presented in subsequent work. Given the cost of manual road digitization, future work could explore semi-supervised approaches that leverage the large volume of uncorrected road segments as weak labels, or self-supervised pretraining on unlabeled LiDAR tiles to improve feature representations before fine-tuning on the labelled Quebec data.

The current workflow assumes the existence of a prior road network to guide the least-cost path correction, meaning it is fundamentally dependent on roads already being mapped. This represents a significant practical limitation, as a considerable portion of Ontario's forestry road network is believed to be unmapped entirely or partially. Extending the framework to support road detection in the absence of any prior network knowledge and transitioning from a correction task to a full detection and mapping task would substantially broaden its operational value.

Finally, the current study does not address the road metrics and classification scheme. The previous study has extracted road metrics and classification based on a threshold-wise approach, which are derived from the expert-defined system (Roussel et al., 2022). This approach may not fully reflect operational standards specific to Ontario's forest industry. Calibration of these thresholds using ground-truth field assessments from Ontario forest sites and integration of DL-driven insights would strengthen the operational applicability of the classification output.

5. Conclusions

This study presented a hybrid deep learning-enhanced workflow for updating forestry road networks in Ontario using Single Photon LiDAR data. By integrating an Attention ResUNet model into the existing parametric road correction framework, the algorithm overcame the key limitation of requiring pre-existing knowledge of forest and road conditions, enabling adaptive, context-dependent road delineation across diverse forest environments. Trained exclusively on Quebec data and deployed across two Ontario forest management units, the model demonstrated strong cross-regional generalizability, achieving a 78% reduction in mean positional offset compared to the original parametric approach. These results confirm that LiDAR-derived features sufficiently abstract sensor-specific differences to support transfer learning across instrument types and geographic domains. The intended open-source design of the workflow positions it as a scalable and practical tool for province-wide road network updating, with direct applications in sustainable forest management and conservation planning.

References

Azizi, Z., Najafi, A. and Sadeghian, S., 2014: Forest road detection using LiDAR data. *J. For. Res.*, 25(4), 975–980.

Botelho Jr, J., Costa, S.C., Ribeiro, J.G. and Souza Jr, C.M., 2022: Mapping roads in the Brazilian Amazon with artificial intelligence and Sentinel-2. *Remote Sens.*, 14, 3625.

Chen, Z., Deng, L., Luo, Y., Li, D., Junior, J.M., Gonçalves, W.N., Nurunnabi, A.A.M., Li, J., Wang, C. and Li, D., 2022: Road extraction in remote sensing data: a survey. *Int. J. Appl. Earth Obs. Geoinf.*, 112, 102833.

David, N., Mallet, C., Pons, T., Chauve, A. and Bretar, F., 2009: Pathway detection and geometrical description from ALS data in forested mountainous area. *Proc. Laserscanning 2009*, Paris, France.

Ghaznavi, A., Rychtáriková, R., Saberioon, M. and Štys, D., 2022: Cell segmentation from telecentric bright-field transmitted light microscopy images using a residual attention U-Net: a case study on HeLa line. *Comput. Biol. Med.*, 147, 105805.

Gluckman, J., 2016: Design of the processing chain for a high-altitude, airborne, single-photon lidar mapping instrument. *Laser Radar Technology and Applications XXI*, SPIE, Vol. 9832, pp. 20–28.

Hoffmann, M.T., Ostapowicz, K., Bartoń, K., Ibisch, P.L. and Selva, N., 2024: Mapping roadless areas in regions with contrasting human footprint. Sci. Rep., 14(1), 4722.

Lidberg, W., 2025: Deep learning-enhanced detection of road culverts in high-resolution digital elevation models: improving stream network accuracy in Sweden. *J. Hydrol. Reg. Stud.*, 57, 102148.

Liu, R., Wu, J., Lu, W., Miao, Q., Zhang, H., Liu, X., Lu, Z. and Li, L., 2024: A review of deep learning-based methods for road extraction from high-resolution remote sensing images. *Remote Sens.*, 16, 2056.

Lu, X. and Weng, Q., 2025: Deep learning-based road extraction from remote sensing imagery: progress, problems, and perspectives. *ISPRS J. Photogramm. Remote Sens.*, 228, 122–140.

Morley, I.D., Coops, N.C., Roussel, J.-R., Achim, A., Dech, J., Meecham, D., McCartney, G., Reid, D.E.B., McPherson, S., Quist, L. and McDonell, C., 2024: Updating forest road networks using single photon LiDAR in northern forest environments. *Forestry*, 97(1), 38–47. doi.org/10.1093/forestry/cpad021.

Oktay, O., Schlemper, J., Folgoc, L.L., Lee, M., Heinrich, M., Misawa, K., Mori, K., McDonagh, S., Hammerla, N.Y., Kainz, B., Glocker, B. and Rueckert, D., 2018: Attention U-Net: learning where to look for the pancreas. Preprint, arXiv:1804.03999. doi.org/10.48550/arXiv.1804.03999.

Ozaki, Y., Watanabe, S. and Yanase, T., 2025: OptunaHub: a platform for black-box optimization. Preprint, arXiv:2510.02798. doi.org/10.48550/arXiv.2510.02798.

Ronneberger, O., Fischer, P. and Brox, T., 2015: U-Net: convolutional networks for biomedical image segmentation. *In:* Navab, N., Hornegger, J., Wells, W.M. and Frangi, A.F. (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, Lecture Notes in Computer Science, Vol. 9351, Springer, Cham, pp. 234–241.

Roussel, J.-R., Bourdon, J.-F., Morley, I.D., Coops, N.C. and Achim, A., 2022: Correction, update, and enhancement of vectorial forestry road maps using ALS data, a pathfinder, and seven metrics. *Int. J. Appl. Earth Obs. Geoinf.*, 114, 103020.

Waga, K., Tompalski, P., Coops, N.C., White, J.C., Wulder, M.A., Malinen, J. and Tokola, T., 2020: Forest road status assessment using airborne laser scanning. *For. Sci.*, 66(4), 501–508.

White, J.C., Woods, M., Krahn, T., Papasodoro, C., Bélanger, D., Onafrychuk, C. and Sinclair, I., 2021: Evaluating the capacity of single photon lidar for terrain characterization under a range of forest conditions. *Remote Sens. Environ.*, 252, 112169.

This contribution has been peer-reviewed.
https://doi.org/10.5194/isprs-archives-XLIX-B1-2026-421-2026 | © Author(s) 2026. CC BY 4.0 License.