Transformer-based Reconstruction of Canopy Profiles from Large-Footprint Waveform LiDAR

Authors: Tahrir I. Siddiqui[1, *], Keith Krause[2], Jan A.N. van Aardt[1]

[1]Chester F. Carlson Center for Imaging Science, Rochester Institute of Technology, Rochester, NY 14623, USA

[2]Battelle, National Ecological Observatory Network, Boulder, CO 80301, USA

* Corresponding author.

E-mail addresses: ts5265@rit.edu (T.I. Siddiqui), kkrause@battelleecology.org (K. Krause), jvacis@rit.edu (J. van Aardt).

Note: This is a non-peer reviewed preprint submitted to EarthArXiv.

Abstract

Spaceborne laser scanning (SLS) presents a cost-effective means for frequent, global-scale monitoring of forest ecosystem parameters. Compared to airborne laser scanning (ALS), SLS offers substantially greater spatial coverage and revisit frequency, but at the cost of larger footprints, sparser sampling, and attenuated return signals. These constraints typically result in a loss of fine-scale vertical canopy structure in large-footprint waveform LiDAR, thereby limiting the retrieval of ecologically meaningful forest structural metrics. To address this challenge, we developed an encoder–decoder Transformer architecture to reconstruct high-resolution vertical canopy profiles from large-footprint waveform LiDAR observations. Using waveform data acquired by NASA's Land, Vegetation, and Ice Sensor (LVIS) – a high-altitude ALS instrument commonly used as a proxy for spaceborne missions – we trained the model to recover fine-scale canopy structure by leveraging overlapping ALS point clouds as reference data. The proposed Transformer leverages long-range vertical dependencies within waveform signals to infer canopy structural details that are degraded or unresolved in large-footprint, high-altitude observations. Results show that the proposed approach substantially improves the agreement between LVIS-derived and ALS-derived canopy structural complexity metrics, increasing correlations from R = 0.62 to 0.84 and from R = 0.76 to 0.90 for two representative metrics. This framework is readily transferable to current and future SLS missions, enabling the retrieval of super-resolved vertical canopy profiles and supporting large-area assessment of ecologically meaningful canopy structural metrics.

Keywords – Full-waveform LiDAR, Transformer, Canopy Structural Complexity, Generative Reconstruction.

1. Introduction

Light detection and ranging (LiDAR) is the leading technology for accurate mapping of three-dimensional (3D) canopy structure (Coops et al., 2021). In this context, aerial laser scanning (ALS) provides dense, high-resolution measurements and has been widely used to derive ecologically meaningful canopy structural metrics across large, forested landscapes (Coops et al., 2021; Siddiqui et al., 2025). However, ALS acquisitions are infrequent and spatially limited due to their high cost, thus restricting their utility for routine, global-scale monitoring and change detection (Hancock et al., 2021). Spaceborne laser scanning (SLS), on the other hand, offers continental to global coverage with frequent revisit intervals at substantially lower cost (Hancock et al., 2021), but at the expense of sparser sampling and reduced vertical resolution, which limits the retrieval of detailed canopy structural metrics. To overcome these limitations, prior studies have estimated canopy structural complexity (CSC) – a metric capturing both horizontal and vertical variability – by statistically linking SLS waveforms from NASA's Global Ecosystem Dynamics Investigation (GEDI) (Dubayah et al., 2020) to overlapping ALS-derived CSC and extrapolating these relationships to global scales (Liu et al., 2024; de Conto et al., 2024). Despite their scalability, such regression-based approaches have achieved only weak-to-moderate performance (R^2 = 0.58–0.68), thereby underscoring the need for methods that more directly recover fine-scale canopy structure from SLS observations.

The need for higher-resolution SLS observations extends beyond CSC estimation. Spaceborne LiDAR measurements with finer spatial sampling and vertical resolution would support a broad range of vegetation-focused applications, including precision agriculture, commercial forestry, wildfire fuel and hazard assessment, biodiversity and habitat mapping, and monitoring of deforestation and forest degradation (Zhao et al., 2018; Donellan et al., 2021). Upcoming full-waveform satellite LiDAR missions such as the Earth Dynamics Geodetic Explorer (EDGE) (UC San Diego, 2026), are poised to deliver denser sampling than all previous SLS missions combined (Harding et al., 2021; Yang et al., 2024), opening the door to canopy structural mapping at spatial scales previously attainable only with ALS. However, recovering fine-scale vertical canopy structure from SLS will remain challenging due to signal attenuation, background noise, and reduced sensitivity to weak returns at orbital altitudes (Chu, 2006; Lang et al., 2022), all of which can complicate the separation of canopy and ground signals within large-footprint waveforms (Lang et al., 2022). To address these limitations, we designed an encoder-decoder Transformer network to effectively enhance the vertical resolution of large-footprint waveforms acquired by NASA's Land Vegetation and Ice Sensor (LVIS) (NASA Goddard Space Flight Center, 2025a), a high-altitude ALS system commonly used as an analog for satellite LiDAR missions, including calibration and validation of GEDI products (NASA GSFC, 2025b; Hancock et al., 2019). LVIS operates at a nominal altitude of ~10 km and carries two instruments – LVIS Facility (LVIS-F), with footprint diameters of 7–10 m, and LVIS-Classic (LVIS-C), with larger footprints of 20–25 m (NASA GSFC, 2025a). We leveraged LVIS-F waveforms in this study because this instrument's footprint size closely matches the anticipated specifications for

upcoming missions (Yang et al., 2024). The high operating altitude also results in reduced vertical resolution relative to conventional low-altitude ALS, i.e., it closely emulates SLS conditions. Densely sampled ALS point clouds overlapping LVIS-F footprints are used as ground truth for training the proposed Transformer model.

Recent deep-learning approaches to reconstructing and super-resolving spaceborne LiDAR observations have emulated orbital acquisition by synthetically downsampling high-resolution airborne data (Ramirez-Jaime et al., 2024; Porras-Diaz et al., 2024; Ramirez-Jaime et al., 2025). Although these studies modeled selected aspects of the orbital sensing process, such as Gaussian beams and footprint aggregation, retrospective degradation of low-altitude, cross-track scanning ALS point clouds cannot fully reproduce the radiometric and geometric characteristics of real spaceborne acquisitions. The simulated inputs remain conditioned on the sampling geometry and return statistics of the source ALS data, whereas orbital waveforms are shaped by near-nadir viewing geometry, substantially longer propagation distances, atmospheric transmission, the transmitted pulse and receiver impulse response, detector noise, and the sensitivity-dependent loss of weak canopy and ground returns (Hancock et al., 2019; Lang et al., 2022). Moreover, because the models are trained and evaluated using observations generated by a prescribed synthetic degradation operator, they may learn to invert artifacts specific to that operator rather than the physical degradation encountered by an actual spaceborne sensor. Research on deep-learning based image super-resolution has shown that models trained using synthetic degradations can generalize poorly when those degradations differ from real acquisition conditions (Gu et al., 2019). By contrast, the LVIS waveforms used here are actual full-waveform

measurements acquired near the upper boundary of the mid-latitude troposphere. Their two-way optical paths therefore traverse most of the tropospheric column, which contains approximately 75–80% of atmospheric mass, nearly all water vapor and clouds, and most atmospheric aerosols (Wei et al., 2019; Dhaka & Kumar, 2023). The measured LVIS returns consequently incorporate much of the atmospheric path-dependent scattering and attenuation encountered by spaceborne observations, providing a more physically grounded analog of orbital acquisition than synthetically downsampled low-altitude ALS. Finally, previous studies evaluated reconstruction fidelity through generic similarity and error measures, canopy-height profiles, height percentiles, canopy height models (CHMs), and digital terrain models (DTMs) (Ramirez-Jaime et al., 2024; Porras-Diaz et al., 2024; Ramirez-Jaime et al., 2025). Such evaluations do not establish whether the reconstructed data preserve higher-order canopy structural metrics with demonstrated relevance to ecosystem processes and applications, thereby lacking an application-specific performance baseline.

We addressed these gaps in prior research by deriving two ecologically relevant CSC metrics from the reconstructed waveforms produced by the best-performing models and assessing their agreement with corresponding metrics derived from overlapping ALS data. This evaluation framework moves beyond generic similarity measures and directly tests whether the proposed reconstruction approach recovers fine-scale vertical canopy structure in a form that preserves ecologically meaningful structural information. The main contributions of this study are as follows:

1) We introduce a novel encoder–decoder Transformer architecture for reconstructing high-resolution vertical canopy profiles from degraded, high-altitude full-waveform

LiDAR measurements. When applied footprint by footprint to the dense, overlapping sampling patterns anticipated for upcoming spaceborne missions, the reconstructed profiles could be spatially integrated to generate contiguous three-dimensional representations with substantially finer structural detail than the original waveforms.

2) We evaluate reconstruction fidelity using two established CSC metrics derived from the reconstructed profiles and their ALS reference counterparts, thereby determining whether the recovered vertical structure preserves higher-order canopy characteristics relevant to ecosystem monitoring rather than merely reproducing the overall waveform shape.

3) We compare the proposed reconstruction framework with conventional waveform deconvolution methods, testing whether it more effectively resolves fine-scale vertical canopy features and improves the agreement of derived CSC metrics with those obtained from ALS reference data.

2. Materials and methods

2.1 Study site and dataset

LVIS and the National Ecological Observatory Network's (NEON) Airborne Observation Platform (AOP) (Kampe et al., 2010) acquired co-located LiDAR data over the Smithsonian Environmental Research Center (SERC), Maryland, USA, in August 2021, providing an ideal dataset for evaluating waveform-based canopy profile reconstruction against dense ALS reference data. SERC is particularly well suited for this study because it is a temperate, humid Mid-Atlantic site dominated by deciduous

forests, with a mean canopy height of approximately 38 m (NEON, n.d.). The site's tall, closed-canopy hardwood stands, including dominant species such as tulip poplar, American beech, and American sweetgum, provide vertically complex forest structure for testing whether fine-scale canopy information can be recovered from high-altitude waveform measurements (NEON, n.d.).

NASA's Gulfstream-V aircraft carried both the LVIS-C and (~25 m footprint) and LVIS-F (~10 m footprint) instruments over SERC on August 5, 2021, as part of a campaign to collect calibration and validation data for GEDI vegetation products over dense forested sites (NASA GSFC, 2025b). The aircraft collected 15 cm range-resolution full-waveform LiDAR data at a nominal altitude of 12 km, covering an approximately 2 km swath and producing an average footprint spacing of ~10 m along and across track (Cornejo et al., 2024). For each laser shot, the latitude and longitude of the corresponding footprint were geolocated with decimeter-level accuracy (NASA GSFC, 2025a) and the elevation at each sampling bin were referenced to the WGS-84 ellipsoid (Blair and Hofton, 2020a).

Small-footprint ALS data were acquired over SERC in August 2021 during NEON AOP flight campaigns using an Optech Galaxy Prime discrete-return LiDAR scanner (Teledyne Optech, n.d.). Flights were conducted at a nominal altitude of 1000 m above ground level (AGL) in flight lines with 37% overlap (product DP1.30003.001; NEON, 2026). The transmitted laser pulses had a beam divergence of 0.25 mrad ($1/e^2$), corresponding to an approximate nadir footprint diameter of 0.25 m at 1000 m AGL. The average pulse density was 38.5 pulses/m^2, and up to five discrete returns were recorded per pulse. The resulting point clouds had an average density of 56 points/m^2,

with each point geolocated with an absolute horizontal accuracy of ~10 cm (1 σ) (Teledyne Optech, n.d.) and an absolute vertical accuracy of less than 0.15 m root mean square error (*RMSE*) (NSF NEON, 2026). The discrete-return LiDAR data were distributed in LAS 1.3 format, with horizontal coordinates referenced to a UTM map projection and ITRF00 datum and elevations referenced to Geoid12A. The point clouds were processed and classified using LAStools v. 230901 (Isenburg, 2023) and provided as classified 1 × 1 km tiles.

2.1.1 LVIS data processing

LVIS-F shots were filtered to retain only those overlapping forested 30 m pixels in the National Land Cover Database (NLCD) 2021 land-cover product (Dewitz, 2023) in order to restrict the analysis to forested areas only. The raw received LVIS-F waveforms were preprocessed to remove noise and isolate the signal portion as follows:

1) Each waveform was smoothed using a Gaussian kernel to suppress high-frequency noise (Blair and Hofton, 2020b).
2) The mean noise level and standard deviation, provided within the LVIS-F L1B data product (Blair and Hofton, 2020a), were used to compute a signal threshold according to:

$$T_{signal} = \mu_{noise} + 4 * \sigma_{noise}$$

where T_{signal} is the signal threshold, μ_{noise} is the signal threshold, and σ_{noise} is the noise standard deviation.

3) The smoothed waveform was then clipped to retain the waveform segment exceeding T_{signal}, thereby isolating the canopy and ground return signal from

empty, noise-dominated bins (Fig. 1). These processed waveforms served as the input profiles for the proposed reconstruction framework.

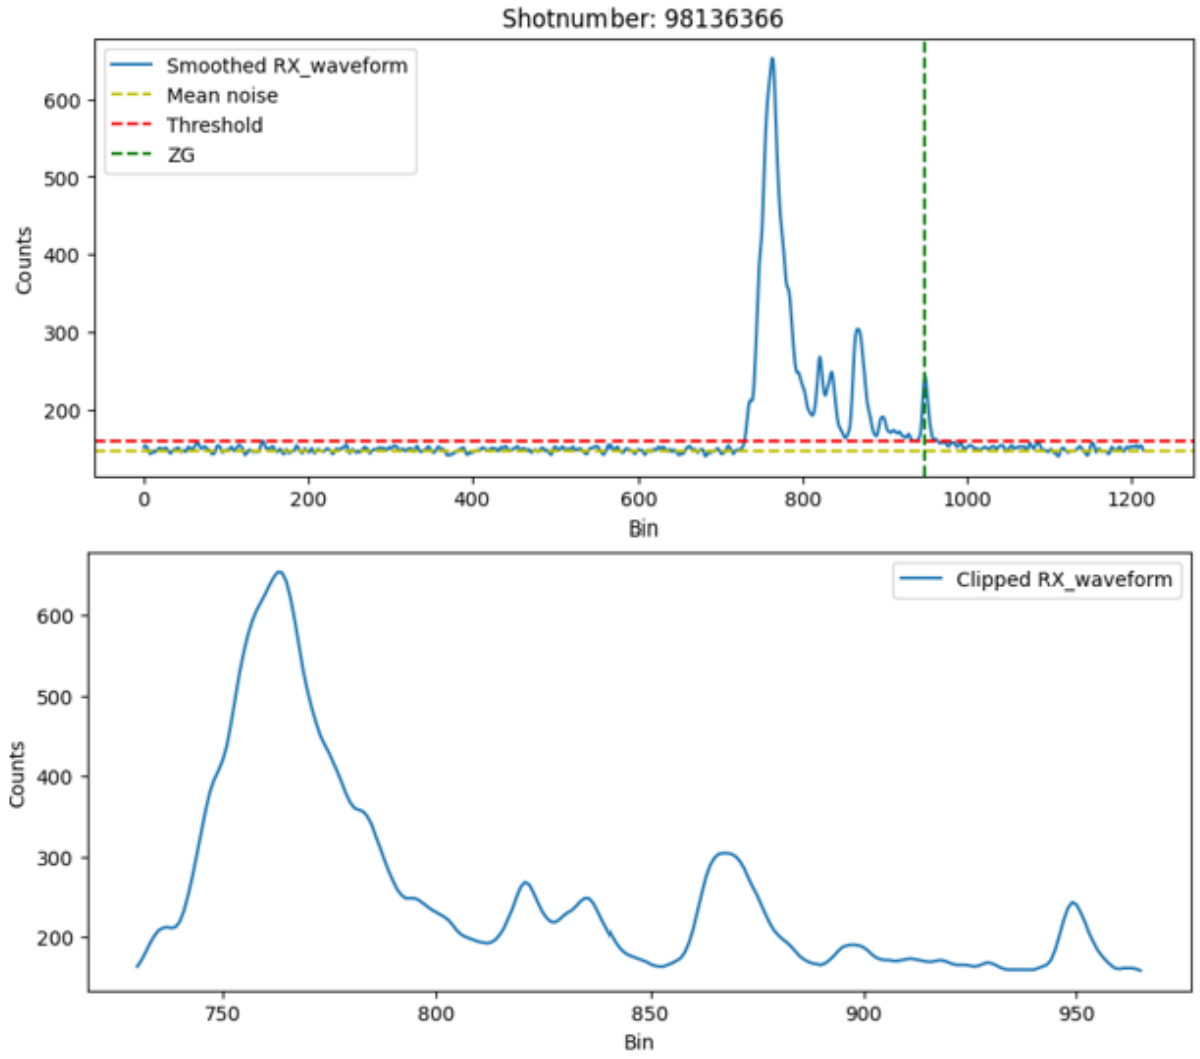

Fig. 1. Top: Smoothed LVIS-F received waveform. The mean noise level and assigned ground elevation (ZG) are provided with the data product. **Bottom:** Corresponding clipped waveform after applying the signal threshold.

2.1.2 ALS data processing

Classified ALS point clouds were clipped into square vertical columns spatially coincident with the filtered LVIS-F shots, with each column centered on the geolocated

ground coordinates of the corresponding LVIS-F footprint (Fig. 2). Square columns were used to provide a canopy structural representation consistent with prior work deriving ecologically meaningful canopy structural complexity metrics from LiDAR return columns (Atkins et al., 2018; Gough et al., 2019; LaRue et al., 2019; Siddiqui et al., 2025). Because LVIS-F footprints range from approximately 7–10 m in diameter, 7 × 7 m columns were extracted to prevent the inclusion of ALS returns outside the LVIS-F footprint extent. Clipping was performed using the clip_roi function in the lidR package (Roussel et al., 2020).

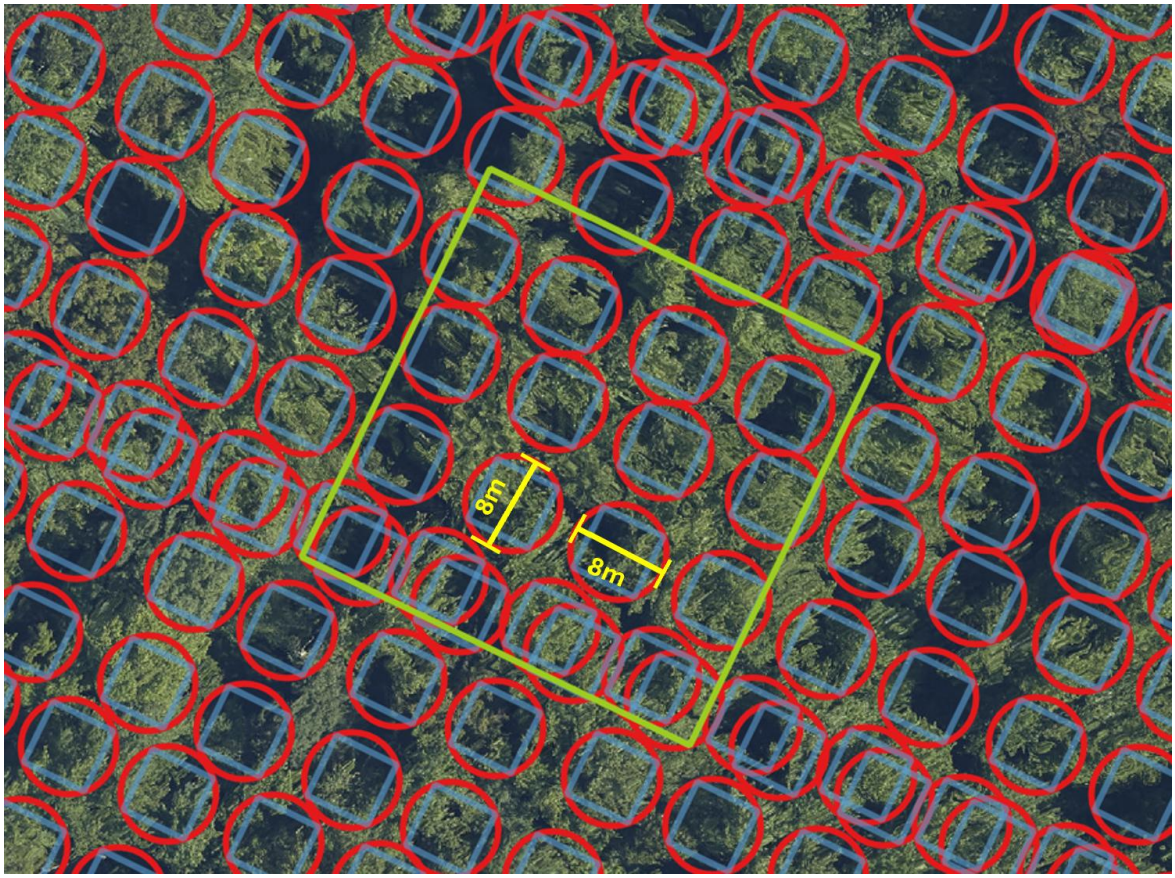

Fig. 2. LVIS-F shots (red circle) and overlapping AOP point cloud columns (blue square) at a forested area in SERC.

A series of filtering steps was then applied to the ALS columns to retain high-quality forest canopy samples:

1) Columns containing points outside the ground, vegetation, or noise classes were excluded.
2) Columns for which the combined number of medium- and high-vegetation returns did not exceed the number of low-vegetation points were removed to eliminate non-forest or sparsely vegetated samples.
3) Points classified as ground, low vegetation, or noise were removed, leaving only medium- and high-vegetation returns for canopy-profile reconstruction.

For applications requiring joint reconstruction of ground and canopy structure, this final point-level filtering step can be omitted. After these sequential filtering steps, 101,110 ALS columns with corresponding LVIS-F waveforms remained for analysis. The large number of co-located LVIS–ALS profile pairs is advantageous for the proposed Transformer-based approach, providing sufficient samples for model training, validation, and independent testing.

Finally, the ALS returns within each 7 × 7 m column were aggregated into 0.15 m vertical bins to match the vertical sampling interval of the LVIS-F waveforms. These vertically binned ALS return-count profiles served as the reference canopy profiles for training and evaluating the waveform reconstruction models. Figure 3 shows three representative LVIS–ALS waveform pairs from the final dataset.

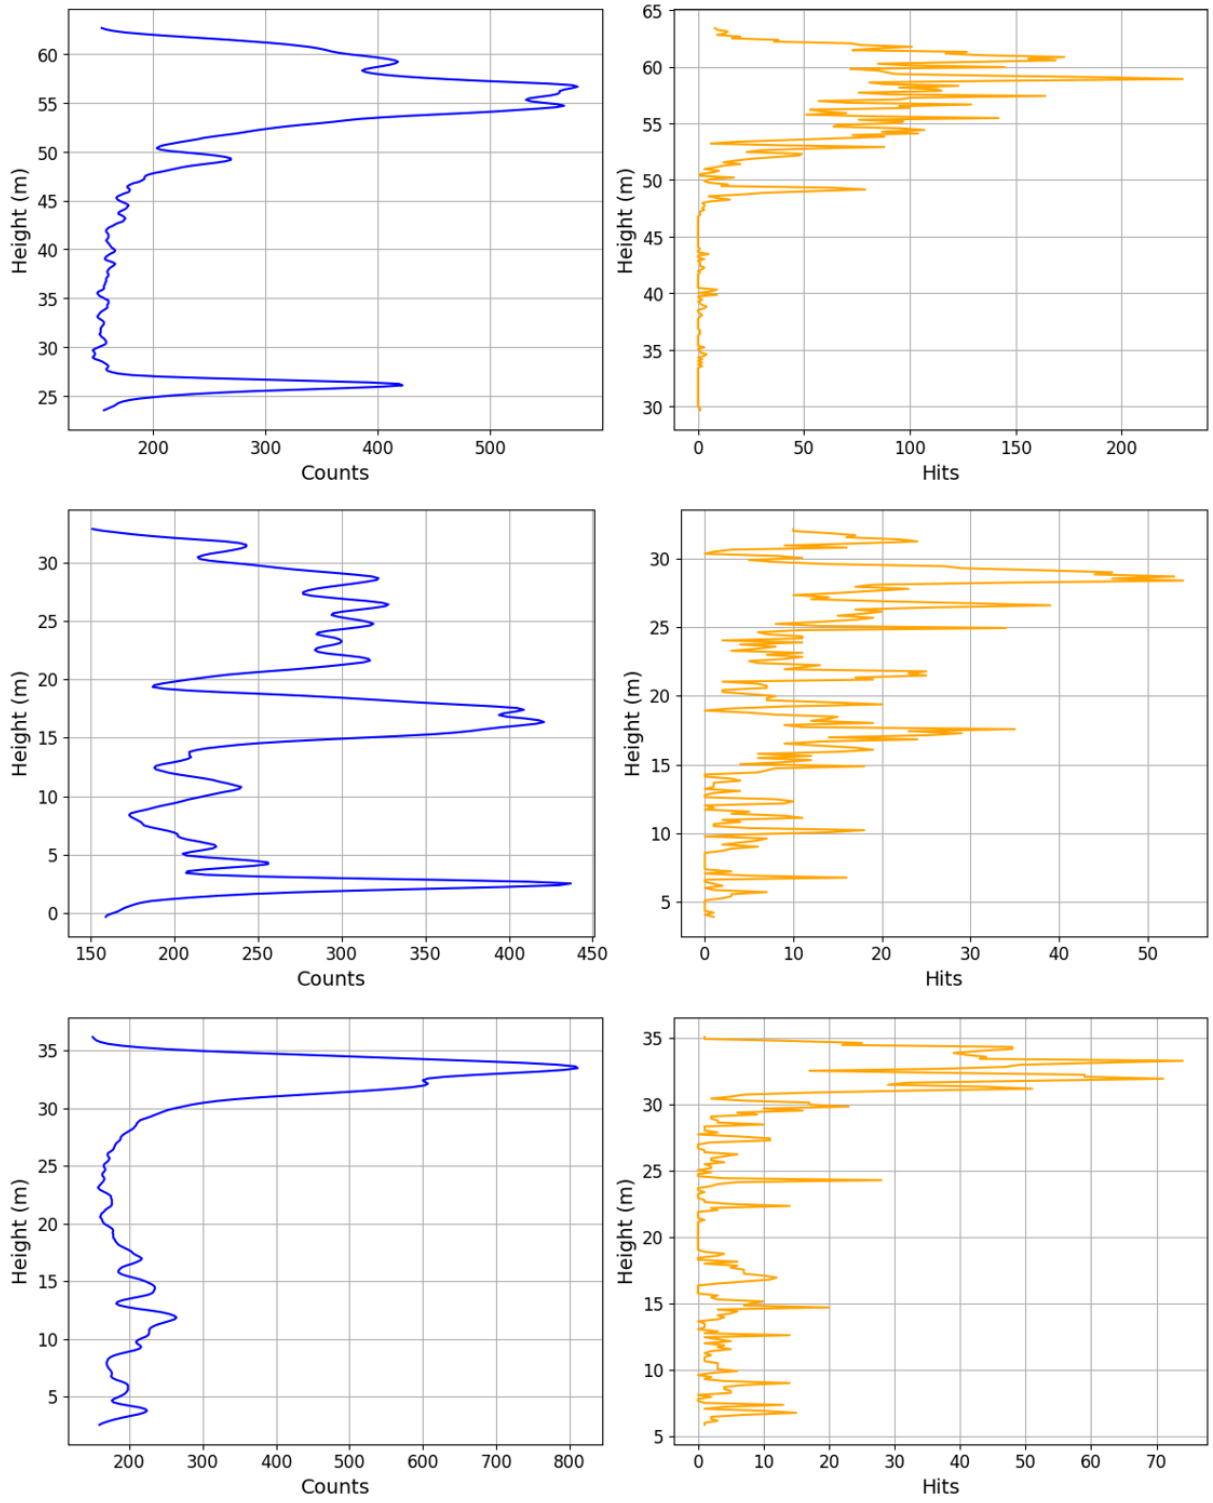

Fig. 3. Co-located LVIS-F waveforms (left) and corresponding ALS-derived canopy profiles (right). In the top pair, the LVIS-F waveform contains a pronounced ground-

return peak at approximately 26 m, which is absent from the ALS canopy profile because ground returns were removed during preprocessing. In the lower two pairs, the ALS profiles reveal finer within-canopy and subcanopy structure between approximately 5 and 25 m that is not resolved as clearly in the corresponding LVIS-F waveforms.

2.2 Transformer architecture for canopy profile reconstruction

We frame canopy profile reconstruction as a cross-sensor sequence-to-sequence learning problem in which a large-footprint LVIS-F waveform is mapped to a high-density ALS-derived canopy profile. The input consists of time-varying LVIS-F return amplitudes at 0.15 m intervals and includes canopy, near-ground, and ground signals. The target consists of canopy-only ALS return counts aggregated at the same vertical interval but containing substantially finer structural detail. Thus, although the input and target share a common vertical coordinate system, they represent measurements with markedly different sensing characteristics and information content. This formulation is analogous to one-dimensional signal super-resolution (Iversen et al., 2025) and sequence-transduction problems such as machine translation, in which an input sequence is transformed into a distinct output representation (Vaswani et al., 2017).

An encoder-decoder Transformer is well suited to this problem for three primary reasons. First, the encoder can accommodate waveforms of varying length through padding and attention masking, allowing the model to process differences in the vertical extent of the measured signal. Second, self-attention (Vaswani et al., 2017) enables the encoder to capture dependencies among vertically separated LVIS-F waveform bins, while decoder self-attention models relationships among output height positions within the reconstructed canopy profile. Together, these mechanisms allow the model to interpret each input and output position in the context of the full vertical sequence.

Third, decoder cross-attention (Vaswani et al., 2017) allows each output height query to condition its prediction on the most informative portions of the encoded LVIS-F waveform. This mechanism enables the model to infer canopy structure that may be attenuated, blended, or unresolved in the original large-footprint observation. The principal components of the proposed architecture are described in the following subsections.

2.2.1 Input embedding

The input embedding module maps each bin of a variable-length LVIS-F waveform to a fixed-dimensional token representation suitable for Transformer processing. Each token incorporates both the height of the bin and the corresponding waveform amplitude. Height values are encoded using learnable Fourier features, inspired by Tancik et al., 2020, with K = 32 frequency components. For a height value h (in meters), the encoding is:

$$\gamma(h) = \left[sin(\boldsymbol{\omega}h) \parallel cos(\boldsymbol{\omega}h) \parallel \tilde{h}\right]^{\top} \in \mathbb{R}^{2K+1}$$

where $\boldsymbol{\omega} \in \mathbb{R}^{K}$ contains log-spaced angular frequencies corresponding to wavelengths from 0.3 to 50 m, enabling the model to capture both fine-scale vertical variation and broader canopy-scale patterns, and $\tilde{h} = h/50$ is the normalized height. These features are projected through a multi-layer perceptron (MLP) to obtain the height embedding:

$$\mathbf{e}_h = \text{MLP}(\gamma(h)) \in \mathbb{R}^{d}.$$

The LVIS-F digital counts, which represent the recorded return-signal amplitude, were encoded using three complementary features. For a waveform with counts $\mathbf{c} = [c_1, \ldots, c_N]^{\top}$ at N height bins, the features for each bin i are:

$$\phi_i^{\text{shape}} = \frac{c_i}{\max_{j=1}^{N}(c_j)}, \qquad \phi_i^{\text{intensity}} = \frac{\log(c_i + 1)}{\log(C_{\max} + 1)}, \qquad \phi_i^{\text{energy}} = \frac{\log(\sum_{i=1}^{N} c_i + 1)}{\log(S_{\max} + 1)}$$

where $C_{\max}$ and $S_{\max}$ are the global maximum count and maximum waveform sum across the training dataset, respectively. The shape feature represents the relative waveform shape, the log-scaled intensity feature preserves absolute count magnitudes at each bin while compressing the wide dynamic range of LiDAR returns, and the energy feature represents the total waveform brightness. The first two features vary across height bins, whereas the waveform energy feature was shared across all bins in a given waveform. Additionally, two optional bin-wise features may be included: the local gradient, which helps identify peaks, slopes, and plateaus; and the distance from peak, which encodes the relative position of each bin with respect to the waveform's maximum. Descriptions and equations for these features are provided in the Appendix. These features are concatenated and projected through an MLP to produce the count embedding $\mathbf{e}_c^{(i)} \in \mathbb{R}^d$. The resulting count and height embeddings are then combined via element-wise addition, followed by layer normalization (Xiong et al., 2020), yielding the final input embedding:

$$\mathbf{z}_i = \text{LayerNorm}\left(\mathbf{e}_h^{(i)} + \mathbf{e}_c^{(i)}\right)$$

2.2.2 Encoder

The encoder processes variable-length LVIS full-waveform inputs and produces contextualized representations that capture both local signal characteristics and long-range structural patterns. It comprises four stacked encoder layers, each containing a multi-head self-attention module with four attention heads, followed by a position-wise

feed-forward network (FFN). The embedding dimension is $d = 128$, and the FFN hidden dimension is $d_{ff} = 512$. Pre-normalization is employed, wherein layer normalization is applied before each sub-layer, to improve optimization stability during training (Xiong et al., 2020). Dropout regularization is applied after the attention and feed-forward operations, and residual connections bypass each sub-layer. Padding masks ensured that padded bins introduced during batching did not contribute to the attention calculations. For a given input sequence $\mathbf{Z} = [\mathbf{z}_1, \ldots, \mathbf{z}_N]^\top \in \mathbb{R}^{N \times d}$, each encoder layer computes:

$$\mathbf{Z}' = \mathbf{Z} + \mathrm{Dropout}\big(\mathrm{SelfAttn}(\mathrm{LayerNorm}(\mathbf{Z}))\big)$$

$$\mathbf{Z}'' = \mathbf{Z}' + \mathrm{Dropout}\big(\mathrm{FFN}(\mathrm{LayerNorm}(\mathbf{Z}'))\big)$$

The self-attention mechanism allows every height bin to incorporate information from all other valid bins in the waveform, enabling the encoder to represent complex vertical dependencies across the full return signal. For example, interception and scattering by upper-canopy foliage influence the amount of laser energy reaching and returning from lower canopy strata. Self-attention may therefore use information from upper-canopy bins to contextualize attenuated returns at lower heights.

2.2.3 ALS Output Queries and Decoder

The decoder generates ALS return predictions through a set of learnable queries anchored to a fixed height grid spanning the ALS vertical range. For our dataset, the output grid extended from 1 to 80 m at 0.15 m intervals, yielding 526 output positions. Each output query combines two components: (1) a learnable content embedding that captures height-specific predictive characteristics, and (2) a height embedding produced

using the same Fourier-feature mapping applied to the LVIS-F input heights. Representing input and output positions within a common height-encoding space facilitates alignment between the LVIS-F waveform and the corresponding ALS canopy profile during cross-attention.

The decoder consists of four stacked pre-normalized Transformer layers, each with four attention heads. Each layer applies multi-head self-attention among the output queries, multi-head cross-attention to the encoded LVIS-F waveform, and a position-wise feed-forward network. Let $\mathbf{X}^{(l)} \in \mathbb{R}^{N_q \times d}$ denote the output-query representations entering decoder layer (l), where N_q is the number of output height positions, and let $\mathbf{E} \in \mathbb{R}^{N_s \times d}$ denote the encoder output for an LVIS-F waveform. The operations within each decoder layer are:

$$\mathbf{S}^{(l)} = \mathbf{X}^{(l)} + \mathrm{Dropout}\left(\mathrm{MHSA}\left(\mathrm{LN}_1\left(\mathbf{X}^{(l)}\right)\right)\right)$$

$$\mathbf{C}^{(l)} = \mathbf{S}^{(l)} + \mathrm{Dropout}\left(\mathrm{MHCA}\left(\mathrm{LN}_2\left(\mathbf{S}^{(l)}\right), \mathbf{E}, \mathbf{E}; \mathbf{M}\right)\right)$$

$$\mathbf{X}^{(l+1)} = \mathbf{C}^{(l)} + \mathrm{FFN}\left(\mathrm{LN}_3\left(\mathbf{C}^{(l)}\right)\right)$$

where $MHSA$ denotes multi-head self-attention, $MHCA$ denotes multi-head cross-attention, LN denotes layer normalization, and $\mathbf{M}$ is the encoder padding mask. In the self-attention operation, the normalized output-query representations serve as the queries, keys, and values, allowing information to be exchanged among all output heights. During cross-attention, the normalized decoder representations serve as queries, whereas the encoded LVIS-F representations serve as keys and values. This allows each output height query to selectively combine information from the observed waveform while preventing padded LVIS-F positions from contributing to the

reconstruction. After the final decoder layer, an additional layer-normalization operation produces the decoder output, $\mathbf{D} = \mathrm{LN}_{\mathrm{final}}\left(\mathbf{X}^{(L)}\right)$.

Because reconstruction at a given height can depend on canopy structure both above and below that position, non-causal (bidirectional) attention was used in the encoder and decoder. The final decoder representations are subsequently passed to the prediction head, which maps the contextualized embedding at each height position to the parameters used to estimate the ALS return distribution.

2.2.4 Count prediction head

The decoder output representations are mapped to discrete ALS return count predictions using a two-layer feed-forward network. Instead of directly regressing counts, we model the output as a Negative Binomial (NB) (Ross & Preece, 1985) distribution to account for the overdispersion characteristic of ALS return counts, where variance frequently exceeds the mean due to the sparse, peaky nature of discrete-return waveforms. For a decoder output vector $h \in \mathbb{R}^d$ at a given height bin, a hidden representation is first computed through two fully connected layers with GELU activations and dropout:

$$\mathbf{f}_1 = \mathrm{Dropout}\left(\mathrm{GELU}(\mathbf{W}_1\mathbf{h} + \mathbf{b}_1)\right)$$

$$\mathbf{f}_2 = \mathrm{Dropout}\left(\mathrm{GELU}(\mathbf{W}_2\mathbf{f}_1 + \mathbf{b}_2)\right)$$

The NB distribution parameters, mean μ and dispersion r, are then predicted as:

$$\mu = \mathrm{softplus}\left(\mathbf{w}_\mu^\top \mathbf{f} + b_\mu\right), \qquad r = \mathrm{softplus}(\mathbf{w}_r^\top \mathbf{f} + b_r)$$

Here, $\mathbf{W}_1 \in \mathbb{R}^{(d\,x\,d/2)}$, $\mathbf{W}_2 \in \mathbb{R}^{(d/2\,x\,d/2)}$, and $\mathbf{w}_\mu$, $\mathbf{w}_r \in \mathbb{R}^{(d/2)}$ are learnable parameters, , where $d_h = d/2$ is the hidden dimension. GELU denotes the Gaussian Error Linear Unit activation (Hendrycks & Gimpel, 2016, and the $\mathrm{softplus}$ function, $\mathrm{softplus}(x) = \log(1 + e^x)$, ensures positivity of both parameters. Under this parameterization, the predicted variance is $Var(y) = \mu + \mu^2/r$, allowing the model to capture count distributions ranging from near-Poisson (large r) to highly overdispersed (small r). The predicted ALS return count at each height bin is given by μ.

2.2.5 Loss functions

Training is guided by a composite objective designed to optimize three complementary aspects of canopy-profile reconstruction: return count fidelity, similarity of the overall profile shape, and suppression of spurious predictions outside the height range of the reference ALS profile. For each sample, the valid region is defined as the set of output-grid positions whose heights fell within the vertical range of the corresponding reference ALS profile. The zero region comprised all remaining positions in the fixed output grid outside this range. The three loss functions are:

1. **Negative Binomial Negative Log-Likelihood**: The primary loss term maximizes the likelihood of observed ALS counts under the predicted NB distribution. For target count y with predicted parameters μ and r, this loss is:

$$\mathcal{L}_{\mathrm{NB}} = \log\Gamma(r) + \log\Gamma(y+1) - \log\Gamma(y+r) + r\log\left(\frac{r+\mu}{r}\right) + y\log\left(\frac{r+\mu}{\mu}\right)$$

where $\log\Gamma(\cdot)$ denotes the natural logarithm of the gamma function. The loss is averaged over all bins in the valid regions of the training batch.

2. **Shape Similarity Loss**: To encourage recovery of the relative vertical distribution of ALS returns independently of their absolute magnitude, cosine distance is computed between the predicted mean profile and the reference profile:

$$\mathcal{L}_{\text{shape}} = 1 - \frac{\mu^{\top} y}{\| \mu \| \| y \|}$$

All output positions contributed to this loss, but positions within the valid profile range were assigned a weight of 1.0, whereas positions outside this range were assigned a reduced weight of 0.1. This weighting allows predictions in the empty regions to influence the profile shape loss without dominating it, while the separate zero-region penalty provides the primary mechanism for enforcing near-zero predictions outside the valid range. The loss was calculated independently for each profile and then averaged across the batch.

3. **Zero-Region Penalty**: To suppress spurious predictions above and below the height range of the reference ALS profile, an $\ell 1$ penalty is applied to the predicted mean counts in the zero region:

$$\mathcal{L}_{\text{zero}} = \frac{1}{| \mathcal{Z} |} \sum_{i \in \mathcal{Z}} | \mu_i |$$

where $\mathcal{Z}$ denotes the set of output positions outside the valid target region. Separating the losses over valid and zero regions allows the model to optimize count fidelity and profile shape within the reference canopy profile range while independently penalizing unsupported predictions elsewhere in the output grid.

The total training objective is given by the weighted sum of these losses:

$$\mathcal{L}_{\text{total}} = \lambda_1 \mathcal{L}_{\text{NB}} + \lambda_2 \mathcal{L}_{\text{shape}} + \lambda_3 \mathcal{L}_{\text{zero}}$$

where λ_1, λ_2, and λ_3 control the relative contributions of the three loss terms. Figure 4 presents a schematic overview of the complete Transformer architecture.

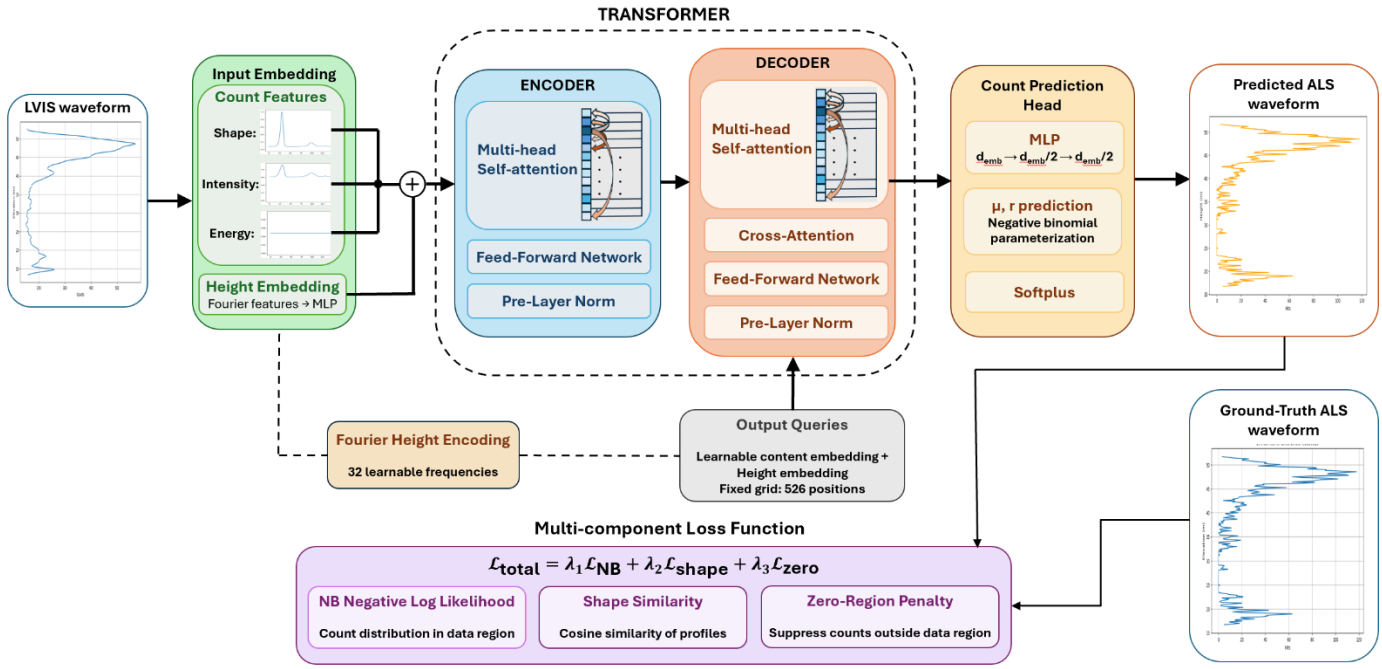

Fig. 4. Transformer-based canopy profile reconstruction architecture. Example input LVIS and ground-truth ALS waveforms are shown from a real paired observation. The predicted ALS waveform is illustrated by recoloring the ground truth for visualization purposes.

2.3 Training setup

The paired LVIS–ALS dataset was partitioned randomly into training, validation, and test sets using an 80/10/10 split (n = 80,888 / 10,111 / 10,111 waveform pairs). The model architecture employed an embedding dimension of 128, four encoder and decoder layers, four attention heads per layer, and a feed-forward hidden dimension of 512. Loss weights were set to λ_1 = 1.2 (NB loss), λ_2 = 0.4 (shape similarity), λ_3 = 1.0 (zero-region penalty) in the initial training run and were tuned in subsequent runs based on validation set performance. Models were trained for a maximum of 100 epochs with

a batch size of 32, and early stopping was triggered after 15 consecutive epochs without improvement in the validation correlation between the predicted and reference ALS profiles. Optimization was performed using Adam with decoupled weight decay regularization (Loshchilov and Hutter, 2017). The initial learning rate was set to 1×10^{-4}, and a cosine-annealing-with-warm-restarts schedule was used to periodically reduce and restart the learning rate during training (Loshchilov and Hutter, 2017). This schedule was selected to promote stable convergence while allowing the optimizer to move away from suboptimal local solutions following each restart. Gradient clipping with a maximum norm of 1.0 was applied to ensure training stability.

The following dataset-level statistics are required to construct the input height and waveform amplitude features and create the output height grid: LVIS-F input height range, ALS reference height range, maximum LVIS-F waveform amplitude, and maximum LVIS-F cumulative waveform energy. Training was conducted on an NVIDIA A100 GPU with 40 GB memory using computing resources provided by the Rochester Institute of Technology (Rochester Institute of Technology, 2026). A typical training run required approximately four hours. The complete codebase, including the model architecture, training scripts, and data preprocessing pipelines, is available in the repository provided in the data availability statement to facilitate replication and future extension of this work.

2.4 Evaluation

Reconstruction performance was evaluated using two complementary pooled correlation measures. Correspondence in absolute return count distributions was quantified by concatenating all predicted canopy profiles within an evaluation subset

into a single vector and similarly concatenating the corresponding ALS reference profiles. Pearson's correlation coefficient (R) was then calculated between the pooled vectors. Profile shape correspondence was assessed using max-normalized correlation (R_n). Each predicted profile and its corresponding ALS reference profile were independently divided by their respective maximum values before concatenation, after which Pearson's correlation was calculated between the resulting normalized vectors. Thus, R reflects correlation between the unnormalized vertical return distributions and remains sensitive to variation in return magnitude among samples, whereas R_n emphasizes similarity in profile shape independently of absolute return counts.

During model development, validation set R was used as the sole criterion for hyperparameter tuning and checkpoint selection. Within each training run, the checkpoint attaining the highest validation set R was retained. The hyperparameter configuration achieving the highest overall validation set R was designated as the best-performing model and applied to the held-out test subset for final evaluation. The pooled R, R_n, and the canopy structural complexity metrics described below were then calculated on the test predictions.

2.4.1 Evaluation of CSC metrics

To assess whether the best-performing model preserved ecologically meaningful canopy structure, foliage height diversity (FHD) (Clawges et al., 2008) and vertical canopy rugosity (VCR) (Siddiqui et al., 2025) were independently derived from the predicted and ALS reference profiles for every sample in the test subset. The same metrics were also derived from the corresponding LVIS-F input waveforms to provide a baseline performance without reconstruction. For each metric, agreement with the ALS-

derived values was quantified using Pearson's *R*, and prediction error was measured using *RMSE*. Performance of the reconstructed profiles was then compared with the LVIS-F baseline to determine whether the proposed model improved recovery of ALS-derived canopy structural complexity.

This evaluation extends beyond pooled profile similarity by directly testing whether the reconstructed vertical return distributions preserve higher-order structural attributes relevant to ecological applications. Previous research has shown that the spatial variability of FHD and VCR metrics is strongly predictive of plot-level primary production in temperate forests (Siddiqui et al., 2025). FHD has also been widely used to characterize habitat complexity and biodiversity (Clawges et al., 2008; Simonson et al., 2014) and to assess changes in vertical forest structure during succession (van Ewijk et al., 2011).

2.4.2 Metric-aware training using FHD and VCR losses

After the best-performing model had been selected using validation set *R*, additional experiments were conducted to determine whether directly incorporating FHD and VCR into the training objective could improve the reconstruction of these metrics. Differentiable FHD and VCR loss terms were implemented following their established formulations (Clawges et al., 2008; Siddiqui et al., 2025) (see Appendix).

The architecture and training hyperparameters of the best-performing model were held constant, while the inclusion and weights of the metric-specific loss terms were varied. Three metric-aware objectives were investigated: (1) the baseline objective supplemented with an FHD loss, (2) the baseline objective supplemented with a VCR

loss, and (3) the baseline objective supplemented with both FHD and VCR losses. Within each training run, the checkpoint achieving the highest validation set *R* was retained. The metric-aware configuration achieving the highest validation set *R* was then evaluated on test subset using pooled *R*, R_n, together with Pearson's *R* and *RMSE* for the predicted FHD and VCR values. These results were compared with those of the baseline model to determine whether metric-aware optimization improved the recovery of ecologically relevant canopy structure without degrading overall profile reconstruction.

2.4.3 Comparison with waveform deconvolution algorithms

Deconvolution techniques have been widely applied to LiDAR waveforms to recover fine-scale vertical features by reducing system and environmental contributions that broaden or obscure the received signal (Wu et al., 2011). Among these techniques, Richardson–Lucy (RL) and Gold have demonstrated strong performance for recovering vegetation cross sections, detecting unresolved echoes, and deriving forest structural parameters relative to alternative methods such as Wiener filtering and nonnegative least squares (Wu et al., 2011; Zhou et al., 2017). Gold and RL were therefore selected as representative conventional deconvolution benchmarks for comparison with the proposed reconstruction framework. Their implementation was adapted from Bhatta et al. (2026), who optimized both algorithms for LVIS waveforms, and the corresponding code is available through the repository provided in the data availability statement.

Because LVIS-F waveforms and ALS reference profiles represent different sensing domains, profile-level performance was compared using the max-normalized pooled correlation, R_n, defined above, together with the corresponding pooled root mean

square error ($RMSE_n$). Each unprocessed (input) LVIS-F waveform, Transformer-reconstructed profile, deconvolved waveform, and ALS reference profile was independently divided by its maximum value before concatenation across the held-out test subset. The following comparisons were then conducted: (1) unprocessed LVIS-F waveforms versus ALS reference profiles, which served as the input baseline; (2) Transformer-reconstructed profiles versus ALS reference profiles; (3) Gold-deconvolved waveforms versus ALS reference profiles; and (4) RL-deconvolved waveforms versus ALS reference profiles. This shape-based comparison reduces the influence of differences in absolute magnitude between LVIS waveform amplitudes and ALS return counts and is consistent with the FHD and VCR analyses, which derive structural information from the normalized vertical distribution of returns.

Recovery of canopy structural complexity was evaluated using the same four profile sources. FHD and VCR were independently derived from the unprocessed LVIS-F waveforms, Transformer predictions, Gold-deconvolved waveforms, and RL-deconvolved waveforms for each test sample. The resulting values were compared with the corresponding ALS-derived FHD and VCR values using Pearson's *R* and *RMSE*.

An important distinction among the methods is that the ALS reference profiles used in this study contain only canopy returns, whereas both the original and deconvolved LVIS-F waveforms retain canopy and ground components. Conventional deconvolution can sharpen blurred waveform features but, without an additional ground-identification and removal procedure, does not transform a canopy-plus-ground waveform into a canopy-only profile. In contrast, the proposed supervised model is explicitly trained to perform this cross-domain transformation and is therefore expected to correspond more

closely with the canopy-only ALS references. Consequently, the deconvolution results should be interpreted as conventional waveform-processing baselines rather than strictly equivalent canopy-profile reconstruction methods. Applications requiring both canopy and ground information could repeat this comparison using ALS reference profiles that retain ground returns.

3. Results

The best-performing model achieved pooled correlations of R = 0.67 and R_n = 0.80 on the test set. The normalized profile shape correlation therefore increased by 0.07 relative to the unprocessed LVIS-F baseline (R_n = 0.73). The reconstructed profiles generally recovered the dominant vertical canopy return patterns observed in the ALS references while suppressing the prominent low-elevation ground returns present in the LVIS-F inputs. Several alternative hyperparameter configurations produced comparable validation performance, indicating that similar reconstruction accuracy was attainable across multiple model settings. The complete hyperparameter configuration of the selected model is provided in the Appendix.

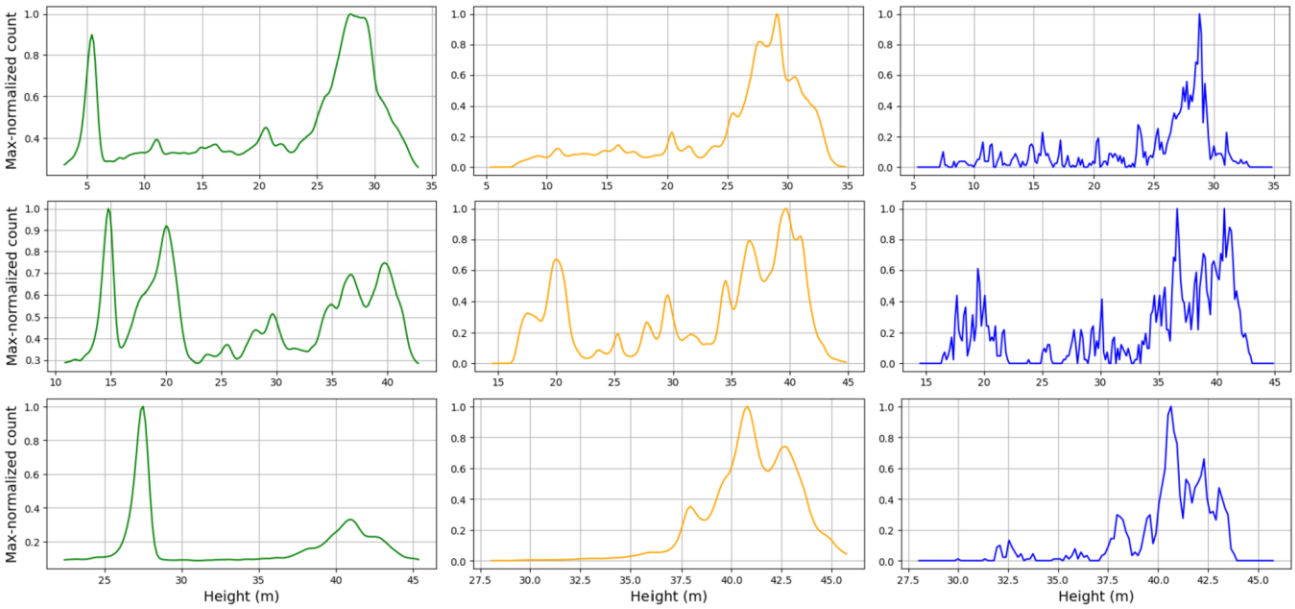

Fig. 5. LVIS-F input waveform (left), Transformer-reconstructed canopy profiles (center), and corresponding ALS reference profiles (right) for three random test samples, displayed by row. Each profile was independently normalized by its maximum value.

3.1 CSC metric derivation performance

CSC metrics derived from the Transformer-reconstructed profiles showed strong agreement with those derived from the ALS reference profiles. FHD values achieved R = 0.90 and $RMSE$ = 0.20, while VCR values achieved R = 0.84 and $RMSE$ = 15.58 m. These results represented substantial improvements over metrics derived directly from the unprocessed LVIS-F waveforms, which achieved R = 0.74 and $RMSE$ = 0.65 for FHD and R = 0.62 and $RMSE$ = 49.85 m for VCR. Thus, reconstruction increased the correlations with ALS-derived VCR and FHD by 0.22 and 0.16, respectively, while reducing the corresponding RMSE values by approximately 69%.

3.2 Metric-aware training results

Table 1 summarizes the test-set performance of the metric-aware training experiments. Adding only the FHD loss produced results similar to those of the baseline model. The max-normalized pooled correlation decreased slightly from R_n = 0.80 to R_n = 0.79, while FHD recovery improved marginally from R = 0.90 and *RMSE* = 0.20 to R=0.91 and RMSE = 0.16. VCR performance remained nearly unchanged.

In contrast, objectives containing the VCR loss produced markedly lower profile-level performance. The highest R_n values were 0.51 for the VCR-only objective and 0.56 for the combined FHD-and-VCR objective, compared with 0.80 for the baseline. However, direct optimization of VCR produced modest improvements in VCR recovery. The VCR-only and combined objectives both achieved R = 0.86, with *RMSE* values of 14.65 m and 14.60 m, respectively, compared with R = 0.84 and *RMSE* = 15.58 m for the baseline. These improvements were accompanied by lower FHD accuracy and substantial degradation in overall profile correspondence. The results therefore indicate that directly optimizing VCR improved recovery of that metric slightly, but at the cost of broader canopy-profile reconstruction fidelity.

Table 1. Test set performance of the metric-aware training objectives. For each objective, results are reported for the configuration attaining the highest validation set R.

Training objective	Pooled *R; RMSE* (counts)	Pooled R_n; $RMSE_n$ (%)	FHD *R; RMSE*	VCR *R; RMSE* (m)
Baseline	0.67; 19.66	0.80; 0.18	0.90; 0.20	0.84; 15.58
FHD only	0.66; 19.83	0.79; 0.18	0.91; 0.16	0.84; 15.78
VCR only	0.43; 27.49	0.51; 0.38	0.83; 0.29	0.86; 14.65
FHD + VCR	0.47; 27.47	0.56; 0.36	0.84; 0.28	0.86; 14.60

3.3 Comparison with waveform deconvolution algorithms

Table 2 compares the proposed Transformer reconstruction with the unprocessed LVIS-F baseline and the Gold- and Richardson–Lucy-deconvolved waveforms. The Transformer-reconstructed profiles achieved the highest max-normalized pooled correlation (R_n = 0.80), exceeding the unprocessed LVIS-F (R_n =0.73), Gold (R_n = 0.68), and RL (R_n = 0.72) results. The Transformer also produced the strongest agreement with ALS-derived CSC metrics, achieving R = 0.90 for FHD and R = 0.84 for VCR.

Neither deconvolution algorithm improved CSC metric correlations relative to the unprocessed LVIS-F baseline. Gold produced FHD and VCR correlations of 0.73 and 0.62, respectively, while RL produced corresponding correlations of 0.74 and 0.62. These values were nearly identical to, or slightly lower than, those obtained from the unprocessed LVIS-F waveforms. Although Gold and RL achieved lower $RMSE_n$ values than the Transformer, their lower R_n values indicate weaker correspondence in the pooled variation of profile shape across the test subset. Overall, the Transformer provided the clearest improvement in both normalized profile correspondence and recovery of ALS-derived canopy structural complexity.

Table 2. Correlation and CSC metric derivation results between ALS reference profiles and unprocessed LVIS-F waveforms, Transformer-reconstructed profiles (baseline model), and deconvolved waveforms on the test set.

Profile source	Pooled R_n; $RMSE_n$ (%)	FHD *R; RMSE*	VCR *R; RMSE* (m)

Unprocessed LVIS-F waveforms	0.73; 0.18	0.74; 0.65	0.62; 49.85
Transformer-reconstructed profiles	0.80; 0.18	0.90; 0.20	0.84; 15.58
Gold-deconvolved waveforms	0.68; 0.12	0.73; 0.57	0.62; 50.66
RL-deconvolved waveforms	0.72; 0.14	0.74; 0.58	0.62; 49.25

4. Discussion

This study demonstrates that a Transformer-based reconstruction framework can substantially improve footprint-level canopy-profile reconstruction from high-altitude full-waveform LiDAR. Compared with unprocessed LVIS-F waveforms and conventional deconvolution baselines, the proposed model produced stronger correspondence with ALS reference profiles and markedly improved the derivation of CSC metrics. The improvement was especially clear for FHD and VCR, for which correlations with ALS-derived values increased substantially relative to both the unprocessed LVIS-F baseline and the Gold- and RL-deconvolved waveforms. These results indicate that the proposed architecture recovers ecologically meaningful vertical structure that is not preserved by direct waveform use or by conventional deconvolution alone.

4.1 Comparison with waveform deconvolution algorithms

The comparison with deconvolution methods is particularly important because Gold and Richardson–Lucy represent strong conventional waveform-processing baselines for vegetation applications. Deconvolution can sharpen waveform features by reducing the influence of the transmitted pulse, receiver impulse response, and other convolutional effects, but it does not, by itself, transform a large-footprint canopy-plus-ground

620 waveform into a fine-scale canopy-only profile. In contrast, the proposed supervised
621 framework learns this cross-domain mapping directly from co-located LVIS-F and ALS
622 observations. This distinction explains why the Transformer-reconstructed profiles
623 showed stronger agreement with canopy-only ALS references and better preservation of
624 ALS-derived CSC metrics than the deconvolved LVIS-F waveforms. The results
625 therefore suggest that learning-based reconstruction provides an important complement
626 to physically motivated waveform-processing approaches when the objective is not
627 merely signal sharpening, but recovery of application-relevant canopy structure.

628 ### 4.2 Metric-aware training performance

629 The metric-aware training experiments suggest that FHD and VCR losses should not
630 be included in the final model objective in their current form. Adding FHD loss produced
631 only negligible improvement, while objectives containing VCR loss slightly improved
632 VCR recovery but substantially degraded overall profile reconstruction. This behavior is
633 consistent with known challenges in multi-objective and multi-task learning. Weighted
634 combinations of losses can induce trade-offs among competing objectives, and
635 improvements in one objective may come at the expense of another when gradients are
636 poorly aligned (Sener and Koltun, 2018; Liu et al., 2021). In this study, the baseline
637 objective already recovered FHD and VCR with strong accuracy, leaving limited room
638 for metric-specific losses to improve performance. The additional metric-aware terms
639 therefore likely introduced optimization conflict and training instability rather than useful
640 complementary supervision. Future work could revisit metric-aware optimization using
641 gradient-surgery methods (Yu et al., 2020) or Pareto-based optimization (Sener and

Koltun, 2018), but the present results support using the baseline reconstruction objective as the preferred training strategy.

4.3 Implications for future spaceborne missions

The strong agreement between CSC metrics derived from reconstructed profiles and those derived from dense ALS reference data highlights the potential of this approach for future spaceborne laser scanning (SLS) missions. Current and past SLS systems, such as GEDI and ICESat-2, provide valuable global observations but remain limited by discontinuous sampling and sparse along-track or multi-beam coverage. Future missions and technology demonstrations, including the Concurrent Artificially Intelligent Spectrometry and Adaptive Lidar System (CASALS) under development at NASA Goddard, are expected to provide denser sampling through adaptive scanning and overlapping footprints (Harding et al., 2021; Yang et al., 2023). If applied to each waveform in such dense sampling configurations, the proposed reconstruction framework could enable ALS-like gridded or voxelized three-dimensional representations of forest structure. These products would support derivation of higher-order structural metrics from 3D forest plots rather than isolated footprints alone.

A key advantage of the proposed framework is its potential scalability. By training Transformers on co-located SLS waveforms and ALS data from representative calibration regions, the learned relationship between high-altitude waveform structure and fine-scale canopy profiles could be extrapolated to larger areas sampled only by SLS. In this way, strategically collected ALS data could serve as high-fidelity reference information for extending ALS-level structural characterization across broader spatial domains. Such transfer will require careful sampling of calibration sites across canopy

types, biomes, terrain conditions, phenological states, and sensor geometries, since model performance may degrade when applied outside the structural and radiometric conditions represented in the training data. Nevertheless, the approach provides a practical pathway for combining the spatial coverage of SLS with the structural detail traditionally associated with ALS.

4.4 Limitations and future directions

Several limitations should be noted. First, although the proposed model improved normalized profile correspondence and CSC metric recovery, the unnormalized pooled correlation remained moderate and did not improve as strongly as the normalized correlation. This indicates that the model is more effective at recovering relative profile shape than absolute return-count magnitude. Consequently, the present framework may be less appropriate for applications requiring raw count fidelity. This limitation is less problematic for many ecosystem applications, including the CSC metrics evaluated here, because such metrics are commonly derived from normalized return distributions rather than raw return counts. Nevertheless, future work should investigate whether losses that explicitly constrain total predicted return counts or separate modules that estimate a sample-specific count-scaling factor can improve absolute count reconstruction without degrading profile shape fidelity.

A second limitation is that the LVIS-F data used here predominantly consist of non-overlapping or weakly overlapping footprints (see Fig. 2) relative to the dense overlapping patterns planned for CASALS. As a result, we could evaluate footprint-level canopy-profile reconstruction and footprint-level CSC metric recovery, but we could not test whether aggregating reconstructed overlapping waveforms produces accurate plot-

level 3D structural metrics relative to ALS point clouds. Once CASALS or comparable dense SLS waveform data become available, future work should evaluate whether 3D representations generated from aggregating reconstructed overlapping waveforms yield accurate plot-level metrics relative to ALS point clouds.

Future work should also link reconstructed canopy profiles directly to field-measured ecosystem variables. In this study, evaluation focused on ALS-derived structural metrics because ALS provides dense, spatially coincident reference data at the footprint scale. The next step is to determine whether reconstructed profiles improve prediction of field-measured forest attributes such as productivity, biomass, canopy fuels, and biodiversity indicators. The framework is also flexible with respect to the target representation. Although the present study trained the model to predict canopy-only profiles, applications requiring ground and canopy information could retain ground returns in the ALS reference profiles and train the Transformer to reconstruct joint canopy-ground structure. Such an extension would support derivation of plot-level canopy cover, canopy base height, fuel load metrics, and terrain-related variables that depend on both vegetation and ground returns.

5. Conclusions

This study developed and evaluated an encoder–decoder Transformer framework for reconstructing high-resolution canopy profiles from high-altitude LVIS-F full-waveform LiDAR. Using co-located ALS canopy profiles as reference data, the proposed model improved footprint-level profile reconstruction relative to unprocessed LVIS-F waveforms and conventional waveform deconvolution baselines. The best-

710 performing model achieved a pooled correlation of R = 0.67 and a max-normalized
711 pooled correlation of R_n = 0.80 on the test subset, improving upon the normalized LVIS-
712 F baseline of R_n = 0.73.

713 The reconstructed profiles also substantially improved the derivation of ecologically
714 relevant canopy structural complexity metrics. FHD derived from reconstructed profiles
715 achieved R = 0.90 and *RMSE* = 0.20 relative to ALS-derived FHD, compared with R =
716 0.74 and *RMSE* = 0.65 for metrics derived directly from unprocessed LVIS-F
717 waveforms. Similarly, VCR derived from reconstructed profiles achieved R = 0.84 and
718 *RMSE* = 15.58 m, compared with R = 0.62 and *RMSE* = 49.85 m for the LVIS-F
719 baseline. These results demonstrate that the proposed approach recovers vertical
720 canopy structure in a form that better preserves higher-order structural metrics relevant
721 to ecosystem monitoring.

722 The Transformer framework also outperformed Gold- and Richardson–Lucy-
723 deconvolved waveforms in normalized profile correspondence and CSC metric
724 recovery. Gold and RL deconvolution achieved R_n = 0.68 and R_n = 0.72, respectively,
725 and neither approach improved FHD or VCR correlations relative to the unprocessed
726 LVIS-F baseline. These findings indicate that conventional deconvolution can sharpen
727 waveform features but is not sufficient, by itself, to transform large-footprint canopy-
728 plus-ground waveforms into fine-scale canopy-only profiles.

729 Overall, the proposed Transformer-based reconstruction framework provides a
730 promising pathway for recovering ecologically meaningful canopy structure from high-
731 altitude and future spaceborne waveform LiDAR observations. When paired with dense
732 sampling configurations from upcoming missions and representative ALS calibration

data, this approach could help extend ALS-like structural characterization to broader spatial extents sampled by spaceborne laser scanning.

Data and code availability

The complete codebase for training and implementing the proposed encoder–decoder Transformer, including the best-performing model checkpoint, is available at https://github.com/tahriribraq/Transformer-waveform-reconstruction. The preprocessed LVIS waveforms and corresponding ALS reference profiles used in the study are available at https://doi.org/10.5281/zenodo.21154804.

Acknowledgements

This work was supported by the National Aeronautics and Space Administration's (NASA) Decadal Survey Incubation (DSI) program [grant 80NSSC22K1097].

References

1. Coops, N. C., Tompalski, P., Goodbody, T. R., Queinnec, M., Luther, J. E., Bolton, D. K., ... & Hermosilla, T. (2021). Modelling lidar-derived estimates of forest attributes over space and time: A review of approaches and future trends. *Remote Sensing of Environment*, *260*, 112477. https://doi.org/10.1016/j.rse.2021.112477.
2. Siddiqui, T., Alveshere, B., Gough, C., van Aardt, J., & Krause, K. (2025). Modeling Primary Production in Temperate Forests Using Three-Dimensional Canopy Structural Complexity Metrics Derived from Airborne LiDAR Data. *Remote Sensing*, *17*(16), 2817. https://doi.org/10.3390/rs17162817.
3. Hancock, S., McGrath, C., Lowe, C., Davenport, I., & Woodhouse, I. (2021). Requirements for a global lidar system: spaceborne lidar with wall-to-wall coverage. *Royal Society open science*, *8*(12), 211166. https://doi.org/10.1098/rsos.211166.

4. Dubayah, R., Blair, J. B., Goetz, S., Fatoyinbo, L., Hansen, M., Healey, S., ... & Silva, C. (2020). The Global Ecosystem Dynamics Investigation: High-resolution laser ranging of the Earth's forests and topography. *Science of remote sensing*, *1*, 100002. https://doi.org/10.1016/j.srs.2020.100002.
5. Liu, X., Feng, Y., Hu, T., Luo, Y., Zhao, X., Wu, J., ... & Su, Y. (2024). Enhancing ecosystem productivity and stability with increasing canopy structural complexity in global forests. *Science Advances*, *10*(20), eadl1947. https://doi.org/10.1126/sciadv.adl1947.
6. de Conto, T., Armston, J., & Dubayah, R. (2024). Characterizing the structural complexity of the Earth's forests with spaceborne lidar. *Nature Communications*, *15*(1), 8116. https://doi.org/10.1038/s41467-024-52468-2.
7. Donnellan, A., Harding, D., Lundgren, P., Wessels, K., Gardner, A., Simard, M., ... & Treuhaft, R. (2021). Observing Earth's changing surface topography and vegetation structure: A framework for the decade. *NASA surface topography and vegetation incubation study*, *210*.
8. Zhao, K., Suarez, J. C., Garcia, M., Hu, T., Wang, C., & Londo, A. (2018). Utility of multitemporal lidar for forest and carbon monitoring: Tree growth, biomass dynamics, and carbon flux. Remote Sensing of Environment, 204, 883-897. https://doi.org/10.1016/j.rse.2017.09.007.
9. Donnellan, A., D. Harding, P. Lundgren, K. Wessels, A. Gardner, M. Simard, C. Parrish, C. Jones, Y. Lou, J. Stoker, K.J. Ranson, B. Osmanoglu, M. Lavalle, S. Luthcke, S. Saatchi, R. Treuhaft, 2021. Observing Earth's Changing Surface Topography and Vegetation Structure: A Framework for the Decade, NASA Surface Topography and Vegetation Incubation Study, 210 pp.
10. UC San Diego, Scripps Institution of Oceanography, 2026. EDGE | Earth Dynamics Geodetic Explorer, https://edge.ucsd.edu/. (Accessed 16 June 2026).
11. Harding, D., Yang, G., Chen, J., Stephen, M., Sun, X., Durachka, D., Li, H., Lu, W., Mackinnon, J., Wise, T., Ranson, J., & Dabney, P. (2021). CASALS: An adaptive lidar and spectrometry SmallSat for a NASA Explorer Mission. In M. Hollaus & N. Pfeifer (Eds.), *Proceedings of the SilviLaser Conference 2021* (pp. 320–322). Technische Universität Wien. https://doi.org/10.34726/wim.2023.

12. Yang, G., Chen, J. R., Medley, B., Harding, D. J., Mazarico, E., Stephen, M. A., ... & Hasselbrack, W. (2024, July). Development of concurrent artificially-intelligent spectrometry and adaptive lidar system (casals) for swath mapping from space. In *IGARSS 2024-2024 IEEE International Geoscience and Remote Sensing Symposium* (pp. 926-929). IEEE. 10.1109/IGARSS53475.2024.10642000.
13. Chu, X., 2006. Lecture 11. LIDAR Equation. University of Colorado Boulder. https://superlidar.colorado.edu/Classes/Lidar2006/Lecture11.pdf (Accessed 16 June 2026).
14. Lang, N., Kalischek, N., Armston, J., Schindler, K., Dubayah, R., & Wegner, J. D. (2022). Global canopy height regression and uncertainty estimation from GEDI LIDAR waveforms with deep ensembles. *Remote sensing of environment*, *268*, 112760. https://doi.org/10.1016/j.rse.2021.112760.
15. NASA Goddard Space Flight Center, 2025. Land, Vegetation, and Ice Sensor (LVIS), https://lvis.gsfc.nasa.gov/Home/index.html. (Accessed 16 June 2026).
16. NASA Goddard Space Flight Center, 2025. Land, Vegetation, and Ice Sensor (LVIS) - GEDI 2021, https://lvis.gsfc.nasa.gov/Data/Maps/GEDI2021Map.html. (Accessed 16 June 2026)
17. Hancock, S., Armston, J., Hofton, M., Sun, X., Tang, H., Duncanson, L. I., ... & Dubayah, R. (2019). The GEDI simulator: A large-footprint waveform lidar simulator for calibration and validation of spaceborne missions. *Earth and Space Science*, *6*(2), 294-310. https://doi.org/10.1029/2018EA000506.
18. Ramirez-Jaime, A., Pena-Pena, K., Arce, G. R., Harding, D., Stephen, M., & MacKinnon, J. (2024). Hyperheight lidar compressive sampling and machine learning reconstruction of forested landscapes. *IEEE Transactions on Geoscience and Remote Sensing*, *62*, 1-16. https://doi.org/10.1109/TGRS.2024.3356389.
19. Porras-Diaz, N., Ramirez-Jaime, A., Arce, G. R., Pena-Pena, K., Harding, D., Stephen, M., ... & Vargas, R. (2024). Transformer end-to-end optimization of compressive LiDARs using imaging spectroscopy side information. *IEEE Transactions on Geoscience and Remote Sensing*, *62*, 1-17. https://doi.org/10.1109/TGRS.2024.3401614.

20. Ramirez-Jaime, A., Porras-Diaz, N., Arce, G. R., & Stephen, M. (2025). Super-resolved 3d satellite lidar imaging of earth via generative diffusion models. *IEEE Transactions on Geoscience and Remote Sensing*. https://doi.org/10.1109/TGRS.2025.3543670.
21. Gu, J., Lu, H., Zuo, W., & Dong, C. (2019). Blind super-resolution with iterative kernel correction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 1604-1613).
22. Wei, P. S., Chiu, H. H., Hsieh, Y. C., Yen, D. L., Lee, C., Tsai, Y. C., & Ting, T. C. (2019). Absorption coefficient of water vapor across atmospheric troposphere layer. *Heliyon*, *5*(1). https://doi.org/10.1016/j.heliyon.2019.e01145.
23. Dhaka, S. K., & Kumar, V. (2023). Composition and thermal structure of the earth's atmosphere. In *Atmospheric remote sensing* (pp. 1-18). Elsevier. https://doi.org/10.1016/B978-0-323-99262-6.00023-7.
24. Kampe, T. U., Johnson, B. R., Kuester, M. A., & Keller, M. (2010). NEON: the first continental-scale ecological observatory with airborne remote sensing of vegetation canopy biochemistry and structure. *Journal of Applied Remote Sensing*, *4*(1), 043510. https://doi.org/10.1117/1.3361375.
25. National Ecological Observatory Network. (n.d.). *Smithsonian Environmental Research Center NEON / SERC*. NSF NEON | Open Data to Understand our Ecosystems. https://www.neonscience.org/field-sites/serc. (Accessed 18 June 2026)
26. Cornejo, H., Hofton, M., & Blair, J. B. (2024, July). *NASA's Land, Vegetation, and Ice Sensor (LVIS) Facility technical reference document* (Version 1.4). NASA Goddard Space Flight Center.
27. Blair, J. B. & Hofton, M. (2020). LVIS Facility L1B Geolocated Return Energy Waveforms. (LVISF1B, Version 1). [Data Set]. Boulder, Colorado USA. NASA National Snow and Ice Data Center Distributed Active Archive Center. https://doi.org/10.5067/XQJ8PN8FTIDG. Date Accessed 06-18-2026.
28. Teledyne Optech. (n.d.). *Galaxy PRIME airborne lidar terrain mapper* [Brochure]. DirectIndustry. https://pdf.directindustry.com/viewerCatalog/optech/galaxy-prime/25132-951631.html.
29. NSF NEON, 2026. Lidar, https://www.neonscience.org/data-collection/lidar. (Accessed 20 June, 2026)

30. NEON (National Ecological Observatory Network). Discrete return LiDAR point cloud (DP1.30003.001), RELEASE-2026. https://doi.org/10.48443/5ts2-rc92. Dataset accessed from https://data.neonscience.org/data-products/DP1.30003.001/RELEASE-2026 on June 18, 2026.
31. Isenburg, M. (2023). *LAStools - efficient tools for LiDAR processing* (Version 230901) [Computer software]. rapidlasso GmbH. https://rapidlasso.de/lastools/
32. Dewitz, J., 2023, National Land Cover Database (NLCD) 2021 Products: U.S. Geological Survey data release, https://doi.org/10.5066/P9JZ7AO3.
33. Blair, J. B. & Hofton, M. (2020). *LVIS Facility L2 Geolocated Surface Elevation and Canopy Height Product*. (LVISF2, Version 1). [Data Set]. Boulder, Colorado USA. NASA National Snow and Ice Data Center Distributed Active Archive Center. https://doi.org/10.5067/VP7J20HJQISD. Date Accessed 06-18-2026.
34. Atkins, J. W., Bohrer, G., Fahey, R. T., Hardiman, B. S., Morin, T. H., Stovall, A. E., ... & Gough, C. M. (2018). Quantifying vegetation and canopy structural complexity from terrestrial Li DAR data using the forestr r package. *Methods in Ecology and Evolution*, *9*(10), 2057-2066. https://doi.org/10.1111/2041-210X.13061.
35. Gough, C. M., Atkins, J. W., Fahey, R. T., & Hardiman, B. S. (2019). High rates of primary production in structurally complex forests. https://doi.org/10.1002/ecy.2864.
36. LaRue, E. A., Hardiman, B. S., Elliott, J. M., & Fei, S. (2019). Structural diversity as a predictor of ecosystem function. *Environmental Research Letters*, *14*(11), 114011. https://doi.org/10.1088/1748-9326/ab49bb.
37. Siddiqui, T., Alveshere, B., Gough, C., van Aardt, J., & Krause, K. (2025). Modeling Primary Production in Temperate Forests Using Three-Dimensional Canopy Structural Complexity Metrics Derived from Airborne LiDAR Data. *Remote Sensing*, *17*(16), 2817. https://doi.org/10.3390/rs17162817.
38. Roussel, J. R., Auty, D., Coops, N. C., Tompalski, P., Goodbody, T. R., Meador, A. S., ... & Achim, A. (2020). lidR: An R package for analysis of Airborne Laser Scanning (ALS) data. *Remote sensing of environment*, *251*, 112061. https://doi.org/10.1016/j.rse.2020.112061.

39. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, *30*.
40. Iversen, M., Khan, M., Miraki, A., & Arghandeh, R. (2025). T2SR: Super-Resolution in Smart Meter Data Using a Transformer-Based Framework. *IET Smart Grid*, *8*(1), e70010. https://doi.org/10.1049/stg2.70010.
41. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., ... & Ng, R. (2020). Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in neural information processing systems*, *33*, 7537-7547.
42. Xiong, R., Yang, Y., He, D., Zheng, K., Zheng, S., Xing, C., ... & Liu, T. (2020, November). On layer normalization in the transformer architecture. In *International conference on machine learning* (pp. 10524-10533). PMLR.
43. Ross, G. J. S., & Preece, D. A. (1985). The negative binomial distribution. *Journal of the Royal Statistical Society: Series D (The Statistician)*, *34*(3), 323-335.
44. Hendrycks, D., & Gimpel, K. (2016). Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*. https://doi.org/10.48550/arXiv.1606.08415.
45. Loshchilov, I., & Hutter, F. (2017). Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*. https://doi.org/10.48550/arXiv.1711.05101.
46. Rochester Institute of Technology. Research Computing Services. Rochester Institute of Technology. https://doi.org/10.34788/0S3G-QD15.
47. Clawges, R., Vierling, K., Vierling, L., & Rowell, E. (2008). The use of airborne lidar to assess avian species diversity, density, and occurrence in a pine/aspen forest. *Remote sensing of environment*, *112*(5), 2064-2073. https://doi.org/10.1016/j.rse.2007.08.023.
48. Simonson, W. D., Allen, H. D., & Coomes, D. A. (2014). Applications of airborne lidar for the assessment of animal species diversity. *Methods in Ecology and Evolution*, *5*(8), 719-729. https://doi.org/10.1111/2041-210X.12219.
49. van Ewijk, K. Y., Treitz, P. M., & Scott, N. A. (2011). Characterizing forest succession in Central Ontario using LiDAR-derived indices. *Photogrammetric Engineering & Remote Sensing*, *77*(3), 261-269. https://doi.org/10.14358/PERS.77.3.261.

50. Wu, J., Van Aardt, J. A. N., & Asner, G. P. (2011). A comparison of signal deconvolution algorithms based on small-footprint LiDAR waveform simulation. *IEEE Transactions on Geoscience and Remote Sensing*, *49*(6), 2402-2414. https://doi.org/10.1109/TGRS.2010.2103080.
51. Zhou, T., Popescu, S. C., Krause, K., Sheridan, R. D., & Putman, E. (2017). Gold–A novel deconvolution algorithm with optimization for waveform LiDAR processing. *ISPRS Journal of Photogrammetry and Remote Sensing*, *129*, 131-150. https://doi.org/10.1016/j.isprsjprs.2017.04.021.
52. Bhatta, R. (2026). *Waveform-LiDAR-Analysis: Novel-LAI-estimation* (Version main) [Computer software]. GitHub. https://github.com/Bhatta6190/Waveform-LiDAR-Analysis.
53. Sener, O., & Koltun, V. (2018). Multi-task learning as multi-objective optimization. *Advances in neural information processing systems*, *31*. https://doi.org/10.48550/arXiv.1810.04650.
54. Liu, B., Liu, X., Jin, X., Stone, P., & Liu, Q. (2021). Conflict-averse gradient descent for multi-task learning. *Advances in neural information processing systems*, *34*, 18878-18890. https://doi.org/10.48550/arXiv.2110.14048.
55. Yu, T., Kumar, S., Gupta, A., Levine, S., Hausman, K., & Finn, C. (2020). Gradient surgery for multi-task learning. *Advances in neural information processing systems*, *33*, 5824-5836.

Appendix

Auxiliary count features

1. **Local gradient:** Captures the rate of change in counts between adjacent height bins, helping identify structural features such as peaks (where gradient crosses zero), slopes (positive or negative gradient), and plateaus (near-zero gradient). For bin i, the gradient is computed as the first-order difference normalized by the waveform maximum:

$$\phi_i^{\text{gradient}} = \frac{c_i - c_{i-1}}{\max_{j=1}^{N}(c_j)}, i = 2, \ldots, N \qquad \text{Eq. (A.1)}$$

916

917 with $\phi_1^{\text{gradient}} = 0$ for the first bin and N is the sequence length.

918 **2. Distance from peak:** Encodes each bin's relative position with respect to the
919 waveform's maximum count, providing structural context within the vertical
920 profile. For bin i, the distance is computed as:

$$\phi_i^{dist} = \frac{i - i^*}{N} \quad \text{Eq. (A.2)}$$

921

922 where $i^* = \text{argmax}_{j=1}^{N}(c_j)$ is the index of the bin with maximum count. This
923 feature is negative for bins below the peak and positive for bins above, with
924 values normalized to approximately $[-1, 1]$.

925 **FHD and VCR loss implementations**

926 **Foliage Height Diversity (FHD) loss.**

927 The FHD loss directly optimizes the Shannon diversity index of the vertical return
928 distribution. Waveform counts are first aggregated into 1-meter height bins, then
929 normalized to obtain proportions. For a waveform with counts aggregated into M height
930 bins, let p_m denote the proportion of total counts in bin m:

931 $$p_m = \frac{\sum_{i \in \text{bin}_m} c_i}{\sum_{i=1}^{N} c_i}$$

932 The FHD is computed as the Shannon entropy:

933 $$\text{FHD} = -\sum_{m=1}^{M} p_m \log(p_m), \text{for } p_m > 0$$

The FHD loss is the mean squared error between predicted and target FHD values:

$$\mathcal{L}_{\text{FHD}} = \frac{1}{B}\sum_{b=1}^{B}\left(\text{FHD}_{\text{pred}}^{(b)} - \text{FHD}_{\text{target}}^{(b)}\right)^2$$

where B is the batch size.

Vertical Canopy Rugosity (VCR) loss.

The VCR loss optimizes the weighted variance of heights, which measures vertical structural complexity. For a waveform with counts $\mathbf{c} = [c_1, \ldots, c_N]^\top$ at heights $\mathbf{h} = [h_1, \ldots, h_N]^\top$, the normalized counts are:

$$p_i = \frac{c_i}{\sum_{j=1}^{N} c_j}$$

The weighted mean height is:

$$\bar{h} = \sum_{i=1}^{N} p_i \cdot h_i$$

The VCR is computed as the weighted variance:

$$\text{VCR} = \sum_{i=1}^{N} p_i \cdot (h_i - \bar{h})^2$$

The VCR loss is the mean squared error between predicted and target VCR values:

$$\mathcal{L}_{\text{VCR}} = \frac{1}{B}\sum_{b=1}^{B}\left(\text{VCR}_{\text{pred}}^{(b)} - \text{VCR}_{\text{target}}^{(b)}\right)^2$$

Hyperparameter configuration of best-performing model

The hyperparameter configuration of the best-performing model is provided below using the exact parameter names defined in the source code. Hyperparameters not listed here were set to their default values in the released codebase. Readers are referred to the source code and accompanying configuration files for definitions and implementation-specific behavior associated with each parameter.

```
# Model architecture
embed_dim = 128
use_auxiliary_features = True
use_cls_token = False
als_query_fusion = "Add"
encoder_layers = 4
decoder_layers = 4
num_heads = 4
ffn_dim = 512
dropout = 0.2

# Data statistics
lvis_height_range = (-15.0, 82.0)
als_height_range = (1.0, 80.0)
height_resolution = 0.15
global_max_count = 4093.0
global_max_sum = 80684.0
```

```
    # Loss weights
    lambda_count = 0.6
    lambda_zero_penalty = 0.4
    lambda_shape = 1.5
    lambda_fhd = 0.0
    lambda_vcr = 0.0

    # Training parameters
    batch_size = 32
    learning_rate = 1e-4
    weight_decay = 0.05
    num_epochs = 100
    scheduler = "cosine"
    early_stopping_patience = 15

    # Data configuration
    train_split = 0.8
    val_split = 0.1
    augment_train = False
    boundary_buffer = 2.0
```