

Computational LiDAR from Space

*Toward the Next Generation of
Satellite LiDAR Systems*

From sparse photon returns to dense 3D Earth structure

Gonzalo Arce


Department of Electrical and Computer Engineering
University of Delaware

Satellite sensing • Computational imaging • 3D reconstruction



LiDAR Imaging

Unmatched sensing tools for planetary and Earth science



High-resolution
3D point clouds



From meters
to 100's of km



Rapid, wide-area
data acquisition

Revealing detail
others can't see



Forested landscapes

- ▶ Topography beneath vegetation
- ▶ Canopy structure
- ▶ Height characterization



Ocean and glaciers

- ▶ Ice-sheet elevation & thickness
- ▶ Bathymetry



Landform evolution

- ▶ Hydrologic & flood modeling
- ▶ Surface earthquakes



Urban growth and sustainability

NASA Spaceborne LiDAR Missions: GEDI and ICESat-2

Footprints, sampling geometry, and applications

GED

Global Ecosystem
Dynamics Investigation

- ISS-mounted full-waveform lidar
- Focus: forests, canopy structure, biomass, topography
- Nominal footprint ~25 m



ICESat-2

Ice, Cloud and land
Elevation Satellite-2

- Photon-counting ATLAS lidar
- Focus: ice, elevation, water, vegetation
- Nominal footprint ~11 m



Platform

GED

ICESat-2



Measurement

ISS

Polar-orbiting satellite



Footprint

Full-waveform lidar

Photon-counting lidar



Sampling geometry

~25 m

8 tracks; 60 m along-track; 600 m across-track

~11 m

6 beams in 3 pairs; 0.7 m along-track;
90 m within pair; ~3.3 km between pairs




Primary applications

Forest structure, biomass, carbon, habitat

Ice-sheet change, sea ice, water height, canopy

GED1: Forest Structure from the ISS

Full-waveform lidar for canopy height, biomass, and ecosystem structure




MISSION SNAPSHOT

- Installed on the International Space Station in 2018
- Measures land between 51.6° N and 51.6° S
- Three lasers generate 8 ground tracks
- Primary targets: forest canopy, topography, and biomass



SAMPLING GEOMETRY

- Footprint diameter: ~25 m
- ⋮ Along-track spacing: 60 m
- ⋈ Across-track spacing: 600 m
- ↔ Across-track width: ~4.2 km



KEY APPLICATIONS




Canopy height & vertical structure



Aboveground biomass & carbon stocks




Habitat & biodiversity



Ground elevation / terrain under canopy

ICESat-2: Photon-Counting Precision

Dense along-track elevation sampling for ice, water, and vegetation



MISSION SNAPSHOT

- Launched September 15, 2018
- ATLAS photon-counting laser altimeter
- Six beams arranged in 3 pairs
- Measures ice sheets, sea ice, water, canopy, and terrain




SAMPLING GEOMETRY

- Footprint diameter: ~11 m
- ↔ Along-track spacing: 0.7 m
- ↔ Within-pair spacing: 90 m
- ↔ Between beam pairs: ~3.3 km
- [--] Strong / weak beam pairs



KEY APPLICATIONS



Paired beams help estimate local surface slope.



Based on official NASA ICESat-2 mission information

Applications Across the Biosphere and Cryosphere

How GEDI and ICESat-2 enable Earth-system science

1) Forest carbon & biomass


GEDI supports aboveground biomass estimation and carbon accounting.



Measure forest height and vertical structure



Estimate biomass and carbon stocks



2) Canopy structure & habitat


Vertical structure helps quantify habitat complexity and ecosystem condition.



Resolve canopy layers and vertical profiles



Inform habitat quality and biodiversity assessments



3) Ice-sheet and glacier change


ICESat-2 tracks elevation change linked to mass balance and sea-level rise.



Measure ice-sheet and glacier elevation



Detect change over time and space



4) Sea ice, water, and coastal zones

Photon-counting lidar supports freeboard, inland-water height, and coastal studies.




Measure sea-ice freeboard and thickness



Observe lakes, rivers, and reservoirs



Quantify coastal change and inundation



Both missions provide complementary 3D information about Earth's surface and structure.



Climate



Carbon



Cryosphere



Hydrology




Ecosystems

LiDAR Imaging Problems


The resolution and sampling gap between airborne and satellite sensing

1 Airborne LiDARs (≈ 300 m)



High-resolution point clouds (X, Y, Z), <10 cm

2 Satellite LiDARs (> 400 km)



Sparse, low-photon, large-footprint sampling



From meters to
100's of km



Airborne LiDARs
(≈ 300 m)



High resolution point
clouds (X,Y,Z) (<10 cm)



Limited coverage in
space/time



Satellite LiDARs
(> 400 km)



Large footprint illumination



Global scale



Low-resolution footprints (>13 m)



Sparse measurements
(>600 m GEDI, >3.3 km ICESat-2)



Low photon count



**Drastic
improvement
needed**

The LiDAR Equation

Photon loss from orbital transmission to return detection

$$\mathcal{Y}_l \sim \underbrace{\mathcal{A}(\mathcal{X}_h; r, \mathcal{G}, \eta^2, \Psi)}_{\text{Instrument's Physics}}$$

Instrument's Physics

- r : range
- $\mathcal{G}$: geometry
- η^2 : noise / system effects
- Ψ : additional instrument physics

$$E_r(r) = E_t \eta_t \eta_r \rho T_{\text{atm}}^2(r) \frac{A_r}{\pi r^2}$$

$$N_r = N_t \eta_{\text{sys}} \rho T_{\text{atm}}^2(r) \frac{A_r}{\pi r^2}$$

System efficiency
(optics, detector,
timing, etc.)

Surface
reflectance
loss

Atmospheric
transmission
losses

Receiver
aperture
collection

Geometric
spreading /
range loss

Physical loss mechanisms



Geometric spreading
($1/R^2$)



Atmospheric
transmission T_{atm}^2



Surface reflectance ρ



Receiver aperture
 A_r



System efficiency η_{sys}
(optics, detector,
timing, etc.)



① **Transmit**
Thousands of photons

② **Surface interaction**
Hundreds reflected

③ **Detection**
Only dozens received



$$N_t \sim 10^3$$

N_t : transmitted photons per pulse

$$N_{\text{surface}} \sim 10^2$$

N_{surface} : photons reflected from the surface


$$N_r \sim 10^1$$

N_r : received return photons per pulse

Orbital LiDAR operates in a *photon-starved* regime.

Our Goal — Close the Gap Between Airborne and Satellite

From sparse satellite observations to airborne-quality 3D reconstruction



Recover dense airborne-like 3D structure from sparse satellite sampling

Concurrent Artificially-intelligent Spectrometry and Adaptive LiDAR System (CASALS)

Adaptive wavelength control and diffractive beam steering for next-generation airborne-to-orbital LiDAR



New-generation LiDAR system



Adaptive wavelength laser





Currently in airborne testing



Expected to be deployed in orbit
in the next years



Goal: dense 3D swath mapping
from space



CASALS NASA Lidar

Adaptive wavelength-scanning lidar for dense 3D swath mapping

Concurrent Artificially-intelligent Spectrometry and Adaptive Lidar System (CASALS)




Co-flown with
Headwall Hyperspec
VNIR-SWIR sensor




Airborne prototype on NASA Langley King Air

ADAPTIVE WAVELENGTH-SCANNING LIDAR



AIRBORNE PUSH-BROOM SWATH



Airborne demo:
256
parallel tracks



Platform:
NASA Langley
King Air,
8.5 km altitude



Swath width:
100 m



Footprint
diameter:
0.7 m



Spacing:
0.43 m cross-track,
0.20 m along-track



~25 photons/pulse;
cm-level precision
after averaging




Scales toward orbital swath mapping


Airborne Footprints, Sampling, and Validation

Real airborne-platform demonstration of CASALS dense swath mapping

1 Airborne demonstration over NEON validation sites




2 Footprint geometry (top-down view)



3 Airborne validation and measurement performance

(a) Example CASALS canopy profile (lidar returns)



(b) Platform and performance summary

- Platform: NASA Langley Beechcraft B200 King Air
- Cruise altitude: 8.5 km
- Air speed: 230 knots (115 m/s)
- ~25 photons per pulse detected
- 2 m pixels by averaging 4x6 groups of footprints
- Several-cm bare-ground range precision



Airborne demo scales to orbital concept: up to 1200 10 m footprints across a 7 km swath.

CASALS Airborne Campaign (Front Royal, VA, USA)


Real airborne-platform data: swath footprints, pulse ordering, and sample track validation

1 Measured airborne footprint pattern over Front Royal




CASALS measurements over map, colored points are GEDI footprints over the same scene.

2 CASALS pulse and track ordering



3 Example elevation profile from airborne returns




Computational Imaging

Satellite LiDAR Imaging Systems: Hardware Efforts

From profile sampling to dense orbital swath mapping

1 Current / emerging sampling geometries



Hardware trajectory

GEDI

Full-waveform LiDAR

25 m footprints

ICESat-2

Photon-counting ATLAS

~11–13 m footprints

CASALS


Adaptive wavelength-scanning LiDAR

10 m orbital footprints over dense grid
Dense ~6.5 km swaths from ~500 km orbit


Closing the Gap Between Airborne and Satellite?

From sparse satellite measurements to airborne-quality 3D reconstruction

(a) Satellite measurements



(b) Airborne scene / ground truth



Satellite input

Low-resolution, noisy,
compressive



Reconstruction

Bayesian inverse
problems +
generative models



Airborne target


Dense, high-resolution
3D structure

Computational Imaging

Take the low-resolution, noisy (and compressive) satellite measurements, do ‘something’, and recover a reconstruction of the high-resolution airborne scene.

What is that “Something”? — Inverse Problems

Closing the gap by combining a physics-based likelihood with a learned prior over high-resolution LiDAR structure




LiDAR Signals as HyperHeight Data Cubes

From waveforms and profiles to 3D structural representations of landscapes




HHDCs represent LiDAR data capturing spatial coordinates and height to provide detailed structural 3D information of landscapes.


- ✓ Each spatial cell stores a vertical return distribution across height.
- ✓ HHDCs preserve terrain, canopy, and vertical structure information.
- ✓ Derived products include DTM, CHM, and RH percentile metrics.



Single-Footprint Return Waveform




Profile-Derived Percentile Metrics



Satellite LiDAR HHDCs: From Waveforms to Profiles


HyperHeight Data Cubes link footprint waveforms, vertical profiles, and landscape structure

1. LANDSCAPE MEASUREMENTS



2. HYPERHEIGHT DATA CUBE (HHDC)

A **HyperHeight Data Cube (HHDC)** is a 3D representation of LiDAR data that captures spatial coordinates and height to provide detailed structural information of landscapes.



3. WAVEFORMS TO PROFILES



Each vertical bar is a 25 m diameter laser footprint.



Grayscale corresponds to the height photon count in the waveform.




Waveform distributions are composed of thousands of photons combined.




Percentiles derived from the profiles include the 2nd percentile (DTM) and 98th minus DTM (CHM), among others.

Footprints / Vertical Profiles




Waveform Profiles & Percentiles



Satellite LiDAR HHDC Formation

Let $\mathcal{X}_h \in \mathbb{R}^{c \times n \times m}$ be the high-resolution airborne tensor, and $\mathcal{Y}_l \in \mathbb{R}^{c_2 \times n_2 \times m_2}$ be the low-resolution satellite tensor.



$$\max_{\mathcal{X}_h} \log p(\mathcal{Y}_l | \mathcal{X}_h) + \log p(\mathcal{X}_h),$$

$$\mathcal{Y}_l \sim \underbrace{\mathcal{A}(\mathcal{X}_h; r, \mathcal{G}, \eta^2, \Psi)}_{\text{Instrument Physics}}$$


Forward Imaging Model

Three fundamental effects that shape satellite LiDAR measurements




When thinking about satellite measurements we need to consider **three phenomena**:

1 Gaussian Footprint Aggregation




How we lose spatial resolution

2 Photon Loss with Distance



How we lose height information

3 Noise Sources




How we degrade our signals



Forward Imaging Model -- Gaussian Footprints

Beam widening with distance in a Gaussian laser footprint



- Laser beams show a Gaussian shape.

$$I(x, z) = \frac{2P}{\pi w(z)^2} \exp\left(-2 \frac{x^2}{w(z)^2}\right),$$


$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2},$$

$$z_R = \frac{\pi w_0^2}{\lambda}$$

where z is the distance, P is the power, w_0 is the beam waist, z_R is the Rayleigh distance, and λ is the wavelength.

Forward Imaging Model — Photon Loss with Distance

Modeling how photon-starved orbital LiDAR measurements arise from altitude-dependent sampling




$$\mathbf{x}_l \sim \text{Multinomial} \left(r, \frac{\mathbf{x}_a}{\|\mathbf{x}_a\|_1} \right)$$

where r is the number of detected photons and $\frac{\mathbf{x}_a}{\|\mathbf{x}_a\|_1}$ is the true photon distribution.

Detected photons decrease with range.

$$r = \frac{E_{\text{det}} \lambda}{ch}, \quad E_{\text{det}} = \frac{E_{\text{shot}} Q A \rho \tau^2}{2\pi z^2}$$



Higher-count regions are more likely to be **sampled** at orbit.



Altitude increases photon loss and **reduces vertical detail**.



Orbital profiles can be emulated by **multinomial sampling** from airborne distributions.



Forward Imaging Model -- Noise Sources

From ideal canopy structure to photon-limited and noisy satellite observations



Poisson noise

$$r^{(i,j)} \sim \text{Poisson}\left(\frac{E_{\text{det}} \lambda}{ch}\right)$$




Gaussian noise

$$\tilde{y}_l^{(i,j)} = \mathbf{x}_l^{(i,j)} + \epsilon^{(i,j)}, \quad \epsilon^{(i,j)} \sim \mathcal{N}(0, \eta^2 \mathbf{I})$$




Noiseless canopy height profile



Ideal reflectivity / return profile




Poisson noise only



Photon-limited counting noise



Poisson + Gaussian noise



Counting noise + detector/readout noise

Forward Imaging Model – Illumination Pattern





Energy-aware sampling through programmable footprint illumination

$$\mathcal{Y}_l = \Psi \odot \tilde{\mathcal{Y}}_l$$

$\tilde{\mathcal{Y}}_l$: fully illuminated low-resolution HHDC

Ψ : illumination mask / sampling pattern

$\mathcal{Y}_l$: measured HHDC after patterned illumination



Why patterned illumination?

- Saving energy is critical for orbital systems
- HHDCs are highly compressible / sparse in suitable transforms




Compressed sensing opportunity

- Programmable Ψ enables efficient acquisition
- Recover structure from fewer illuminated footprints

Forward Imaging Model

A probabilistic chain from high-resolution airborne structure to satellite measurements

$$\mathcal{Y}_l = \underbrace{\begin{pmatrix} y_l^{(1,1)} & \cdots & y_l^{(1,m)} \\ \vdots & \ddots & \vdots \\ y_l^{(n,1)} & \cdots & y_l^{(n,m)} \end{pmatrix}}_{\in \mathbb{R}^{c \times n \times m}} \sim \mathcal{A} \left(\mathcal{X}_h; \frac{E_{\text{det}} \lambda}{ch}, \mathcal{G}, \eta^2, \Psi \right).$$



$$\mathcal{X}_a = \mathcal{X}_h * \mathcal{G}, \left\{ \begin{array}{l} r^{(i,j)} \sim \text{Poisson} \left(\frac{E_{\text{det}} \lambda}{ch} \right), \\ \mathbf{x}_l^{(i,j)} \sim \text{Multinomial} \left(r^{(i,j)}; \frac{\mathbf{x}_a^{(i,j)}}{\|\mathbf{x}_a^{(i,j)}\|_1} \right) \end{array} \right\} \Leftrightarrow \underbrace{\text{Poisson} \left(\frac{E_{\text{det}} \lambda}{ch} \frac{\mathbf{x}_a^{(i,j)}}{\|\mathbf{x}_a^{(i,j)}\|_1} \right)}_{\text{Poisson thinning}} \left| y_l^{(i,j)} \sim \mathcal{N} \left(\mathbf{x}_l^{(i,j)}, \eta^2 \mathbf{I} \right).$$




The forward imaging model defines the likelihood $p(\mathcal{Y}_l | \mathcal{X}_h)$.

Forward Imaging Model -- HHDC Image Sequence

From high-resolution HHDC to noisy, patterned satellite observations

$$\mathcal{Y}_i \sim \mathcal{A}\left(\mathcal{X}_h; \frac{E_{\text{det}} \lambda}{ch}, G, \eta^2, \Psi\right)$$

Instrument physics connects the latent high-resolution scene to the measured low-resolution tensor



High-resolution $\mathcal{X}_h$

Aggregated $\mathcal{X}_a$

Photon-sampled $\mathcal{X}_t$

Noisy $\hat{\mathcal{Y}}_i$

Patterned $\mathcal{Y}_i$

The forward model provides the likelihood

$$p(\mathcal{Y}_i | \mathcal{X}_h)$$


used to solve the inverse problem.

HHDC Diffusion Model

Forward diffusion progressively corrupts the cube

$$q(\mathbf{x}_T | \mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t | \mathbf{x}_{t-1})$$

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \mu_{t-1}, \Sigma_{t-1})$$




Train a denoising neural network p_θ that recovers a **high-resolution cube** out of pure noise.

Diffusion Model — Denoising Neural Network

A conditioned U-Net predicts the reverse diffusion step

- p_θ is implemented through a deep U-Net conditioned on the timestep and the low-resolution tensor.



Solving Inverse Problems

Posterior sampling closes the gap between sparse measurements and high-resolution structure

In order to approximate the solution to (i.e., close the gap):

$$\max_{\mathcal{X}_h} \log p(\mathcal{Y}_l | \mathcal{X}_h) + \log p(\mathcal{X}_h),$$

we sample from the posterior by solving the ODE backward in time:

$$\frac{d\mathcal{X}_h}{dt} = -t \left[\underbrace{s_\theta(\mathcal{X}_h, t)}_{\text{Prior}} + \zeta \underbrace{\nabla \log p_t(\mathcal{Y}_l | \mathcal{X}_h)}_{\text{Likelihood}} \right],$$



Prior score field
 $s_\theta(\mathcal{X}_h, t)$
from pretrained
denoising neural
network $p_\theta(\mathcal{X}_h, t)$



Likelihood /
measurement
guidance
 $\nabla \log p_t(\mathcal{Y}_l | \mathcal{X}_h)$
from forward imaging
model $p(\mathcal{Y}_l | \mathcal{X}_h)$



Posterior
sampling
guided backward
in time to draw
high-quality
samples

guided by our forward imaging model $p(\mathcal{Y}_l | \mathcal{X}_h)$ and our pretrained denoising neural network $p_\theta(\mathcal{X}_h, t)$.

Airborne LiDAR Data Resources: NASA G-LiHT and NSF NEON

Open airborne datasets for ecosystem structure, canopy height, and terrain mapping



NASA G-LiHT

Goddard's LiDAR, Hyperspectral & Thermal Imager



1 Platform

- Portable airborne imaging system
- Coincident LiDAR, hyperspectral, and thermal sensing
- Fine-scale observations (< 1 m)





2 LiDAR data available

- Classified point clouds in ASPRS LAS format
- Bare-earth elevation and canopy height models
- LiDAR metrics: height and density percentiles



3 Access / strengths

- Download via the G-LiHT Data Center Webmap
- Open NASA data policy
- Ideal for data fusion and ecosystem studies



Coincident multi-sensor data



NSF NEON AOP

National Ecological Observatory Network – Airborne Observation Platform



1 Platform

- Airborne surveys over NEON field sites during peak greenness
- Discrete and full-waveform LiDAR, imaging spectrometer, and digital camera
- Regional flight boxes with nominal 1 m LiDAR products



2 LiDAR data available

- Discrete return point clouds
- Canopy Height Model (CHM)
- Elevation products: DTM and DSM
- Derived terrain layers such as slope and aspect




3 Collection characteristics

- Minimum flight-box coverage about 100 km² per site
- Minimum 4 LiDAR shots per m² in non-overlap areas
- Up to 5 points discretized per shot



4 Access / strengths

- Open NEON Data Portal access
- Repeat observations across many U.S. field sites
- Strong integration with in-situ ecological measurements



Continental-scale ecological observatory

Where data are available



NASA G-LiHT (campaign-based)

Broad airborne LiDAR acquisitions across the U.S., Alaska, and Puerto Rico (campaign footprints vary by year) + selected campaigns in Mexico



NSF NEON AOP (site-based)

LiDAR collections at NEON field sites distributed across the U.S. and Puerto Rico



0 250 500 750 1000 km

Why these datasets matter



Canopy height



Terrain / DTM



Vegetation structure



Algorithm development & validation

Experiments – Smithsonian Environmental Research Center


- ▶ Input Footprint size: 10 m
- ▶ Input Horizontal Resolution: 6 m × 3 m
- ▶ Output Footprint size: 2 m
- ▶ Output Horizontal Resolution: 2 m × 2 m




Illumination Pattern: ≈ 20 photons per footprint at 0 dB SNR



ILLUMINATION PATTERN & EXAMPLES



Topographical & Biome Diversity Sensing



- ▶ Validated across diverse U.S. ecoregions
- ▶ Captures distinct canopy, terrain, and profile signatures
- ▶ Demonstrates robustness to regional structural variability
- ▶ Each landscape presents unique sensing nuances



Figure 1 displays a comparison of CS and Super-Resolution results across three rows (Canopy Height Model, Terrain, Height Profile) and four columns (Noisy CS Input, Denoised Super Resolved, Ground Truth, Error). The y-axis for the first two rows is Distance (m) (0 to 540), and for the third row is Height (m) (0 to 54). The x-axis for all rows is Distance (m) (0 to 270). The color bars on the right indicate Absolute Error (dB) for the first two rows and Absolute Error (m) for the third row.



Canopy
Height Model

Distance (m)




Terrain
(dB)

Distance (m)




Height Profile
(m)

Height (m)



Super Resolution Footprint and 3D Comparison



CASALS resolution




ML HR resolution

Key Findings – Comparison with Several Baselines

Method	CHM			DTM		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
3D U-Net (convolutional)	14.50	0.51	0.35	20.32	0.61	0.27
SwinIR (transformer)	16.17	0.59	0.34	20.43	0.65	0.26
PnP-ADMM+BM3D	16.71	0.40	0.29	22.38	0.60	0.31
Proposed (diffusion-guided)	18.77	0.61	0.16	25.90	0.79	0.16

Reconstruction vs Ground Truth



Super-Resolution with RGB Side Information



CASALS is multimodal by nature, combining **LiDAR** with spectral / **RGB** information.




RGB side information enables super-resolution down to **1 m** and improves structural recovery.



Performance generalizes across diverse U.S. regions including **Florida**, **SERC**, and **California**.

Florida Test Site




Quantitative Performance Across Test Sites

Metric	Florida			SERC			California		
	Guided	Side	Comb.	Guided	Side	Comb.	Guided	Side	Comb.
$PSNR_{chm}$	19.68	30	30.26	20.35	23.89	28.7	14.58	22.5	23.08
$PSNR_{dtm}$	8.7	29.58	29.83	26.18	25.26	31.57	32.32	28.77	35.31
$SSIM_{chm}$	0.32	0.94	0.94	0.53	0.79	0.89	0.41	0.89	0.9
$SSIM_{dtm}$	0.28	0.9	0.9	0.77	0.68	0.82	0.87	0.87	0.88



Side information substantially improves both **CHM** and **DTM** reconstruction, and the **combined model** yields the strongest overall performance.

CASALS Airborne Campaign



What's Next: Toward Global 3D Earth Imaging

From established super-resolution results to the next frontier: meter-resolution, global 3D sensing


Established Results




- ▶ Validated local and regional reconstructions



- ▶ Airborne-to-satellite super-resolution and HHDC modeling



- ▶ Multimodal fusion with LiDAR, RGB, and side information




Goal: globally reconstruct Earth's 3D structure at **meter-scale resolution**.

Ongoing Work

LiDAR + InSAR + Optical
meter-scale 3D Earth + confidence

- ◆ **Uncertainty-aware inversion**
predict mean 3D structure plus per-pixel confidence
- ⚡ **InSAR fusion**
phase/elevation cues sharpen terrain under canopy
- ▤ **Global robust models**
domain shift across biomes, slopes, and sensors
- 🕒 **Accelerated global inference**
scalable training and deployment at global scale

Outputs: DTM • CHM • RH-percentiles • σ / credible intervals



Sensing

photon-starved measurements, sparse coverage, multimodal acquisition



Challenges & Opportunities

Learning

generative priors, scalability, uncertainty, domain shift



Impact


carbon, biodiversity, habitat structure, climate resilience

Pure Quantile Reconstruction


- ▶ Targeted reconstruction on useful information.
- ▶ Tailored cost function.

$$\mathcal{L}_{\text{QR}}(\theta) = \sum_{\tau \in \mathcal{T}} w_{\tau} \sum_u M(u) \left[\rho_{\tau} \left(z_g(u) - \hat{z}_{g,\tau}(u) \right) + \rho_{\tau} \left(z_t(u) - \hat{z}_{t,\tau}(u) \right) \right].$$


25th Percentile



50th Percentile



98th Percentile




CHMv2: Global Canopy Height Mapping from Optical Imagery

Meta and collaborators approximate global canopy height models from RGB satellite imagery

- ▶ **Near-global, meter-resolution** canopy height mapping from high-resolution optical imagery
- ▶ **Built on DINOv3** for depth-style canopy height estimation
- ▶ **Trained against** airborne laser scanning (ALS) canopy height models
- ▶ **Validated with independent** ALS data and large-scale GEDI / ICESat-2 observations
- ▶ **Improves detail, global consistency, and reduces bias in tall forests**



CHMv2 provides a powerful 2D canopy-height approximation at planetary scale, but unlike LiDAR-based HHDC methods it does **not** recover the full vertical profile structure.



CHMv2 Workflow



1. Optical Satellite Imagery



2. DINOv3 / CHMv2 Model



3. Global CHM




Tropical Rainforest (Amazon)

RGB



CHMv2 (m)




0 60 m+

Temperate Forest (Pacific NW)

RGB



CHMv2 (m)




0 60 m+

Savanna / Mixed Landscape (Africa)

RGB



CHMv2 (m)




0 60 m+



Comparison with Different Imaging Modalities


- **Cross-modality:** LiDAR-diffusion (3D HHDC) vs. RGB-based CHMv2 (2D CHM)
- **Performance:** 2.52 m MAE & 20.5 dB PSNR vs. CHMv2's 6.69 m & 15.4 dB (2.7× MAE reduction)
- **Advantages:** Eliminates tall-canopy bias; recovers vertical profiles inaccessible to RGB




Accelerating Diffusion for Global Scale

GeM-SR: multiscale Markovian generation for fast high-fidelity reconstruction


(a) JS-RQ Tokenizer



(b) Conditioning



(c) Markov Generation



Key
Takeaways



Wavelet-domain
Decomposition

Haar DWT into $\{S_0, W_0, W_1, W_2\}$
for hierarchical tokenization.



Physics-informed
Conditioning

Inverse-Physics Stem + DWT
produce multiscale conditions
 $(c_g, c_{root}, c_0, c_1, c_2)$.



Multiscale Token
Generation

Coarse-to-fine Markov chain
generates $\hat{S}_0 \rightarrow \hat{S}_3$
across scales.




Non-Autoregressive
Markovian Refinement

Parallel token generation at each scale
with Refiner for high-fidelity SR.

GeM-SR Results: Qualitative and Quantitative Comparison

QUALITATIVE COMPARISON

-  Zoom region (1) structures
-  Zoom region (2) textures
-  4x super-resolution (LR → HR)



QUANTITATIVE COMPARISON

Bold = absolute best. Underline = best generative architectures.

Model	PSNR ↑	SSIM ↑	LPIPS ↓	FID ↓	FPS ↑	VRAM (MB) ↓	Params (M) ↓
<i>Deterministic & Quasi-Deterministic Baselines</i>							
SwinIR	25.75	0.762	0.268	101.13	732.25	884.0	23.77
Flow-Matching (4 step)	25.61	0.746	<u>0.248</u>	76.96	691.36	1086.0	<u>56.72</u>
<i>Multi-Step Generative Baselines</i>							
SR3 (50 steps)	24.88	0.730	0.224	27.47	53.99	<u>1088.9</u>	56.72
SR3 (100 steps)	24.71	0.718	0.212	14.56	26.97	<u>1088.9</u>	56.72
SR3 (200 steps)	24.60	0.709	0.205	9.68	13.49	<u>1088.9</u>	56.72
VAR-SR	24.00	0.692	0.199	17.37	197.19	2943.5	51.79
SR3 (2000-step DDPM)	24.54	0.707	0.199	7.79	1.37	1090.0	56.72
<i>GeM-SR (Remote Sensing Temperatures)</i>							
GeM-SR (ConvNeXt, Uniform)	24.95	0.713	0.207	19.44	1267.14	1730.3	27.17
GeM-SR (Transformer, Uniform)	24.92	0.716	0.204	18.67	887.47	2600.4	43.35
GeM-SR (ConvNeXt, Ascending)	24.91	0.709	0.205	16.09	1270.29	1731.0	27.17
<i>GeM-SR (Perceptual Temperatures)</i>							
GeM-SR (ConvNeXt, Uniform)	24.58	0.689	0.211	14.77	1265.29	1732.3	27.17
GeM-SR (ConvNeXt, Ascending)	24.52	0.684	0.210	12.43	1266.66	1733.8	27.17



1. Best balance of fidelity and efficiency

GeM-SR (ConvNeXt, Ascending) delivers strong PSNR/SSIM with state-of-the-art FPS and low parameters.



2. Competitive generative quality

Lower FID and LPIPS than other efficient generative baselines, matching multi-step quality with far fewer resources.



3. Improved sampling efficiency

~940× faster than SR3 (2000-step DDPM) while maintaining superior fidelity.

Computational LiDAR for Remote Sensing

Numerous opportunities for exploration

University of Delaware PhD student team



Andres Ramirez



Nestor Porras



Kyle Regan



Christian Neuman

NASA CASALS team



◆ Mark Stephan

◆ Guan Yang

◆ David Harding



With thanks to our students, collaborators, and the **CASALS** team.



*Toward the next generation
of global 3D Earth sensing.*