


OLI / GMUG 2026 · Joint Working Group Meeting

# TreeDGS

*Aerial Gaussian Splatting for Direct DBH  
Measurement*

**Michael Wu · Co-founder & CEO, Coolant**

*Joint work with B. Shaheen, M.-H. Nguyen, B.-T. Bui, Shubham, T. Wu, M. D. Zane, M. Fairley (Coolant) & Professor J. Tompkin (Brown University)*



OUR TEAM

# Spatial AI and 3D reconstruction experts.

*A combined 36,100+ citations in spatial AI and earth observation*



**Michael Wu**  
CEO




**Dr. Bach-Thuan Bui**  
Founding Research Scientist



**Prof. Andrew Maas**  
Scientific Advisor



**Prof. James Tompkin**  
Scientific Advisor



We believe in *images*.

Our claim: aerial drone images alone can measure **every tree's diameter** in a managed pine stand.

200 ft

*flight altitude*

RGB only

*no LiDAR payload*


No calibration

*no ground tape required*

THE CORE PROBLEM

# From high up, a 20 year old stem is only ~13 pixels wide

*Thin, partly-hidden, and seen from only a few directions — this is the data any aerial model has to work with.*



# Today, every aerial workflow loses the trunk somewhere

*Whether that's low cost photogrammetry or high end aerial LiDAR*

## TERRESTRIAL LIDAR

### **Accurate trunks, but no canopy & hard to scale**

Dense trunk returns give precise DBH, but crews on the ground cannot scan stands the way a cruiser paces them.

## AERIAL LIDAR

### **Accurate canopy, but trunk returns few**

Canopy-down scanning leaves trunk points sparse and uneven exactly where breast-height is.

## SFM / MVS FROM RGB

### **Decent canopy, but no trunk information**

Classical photogrammetry is getting good at canopy and ground; but weak texture and occlusion break reconstruction on trunks at breast height.

*The open question — can we recover field-grade DBH from aerial RGB alone?*

THE INSPIRATION

# The inspiration: Tesla's *Full Self Driving*

*Image below courtesy of Tesla*

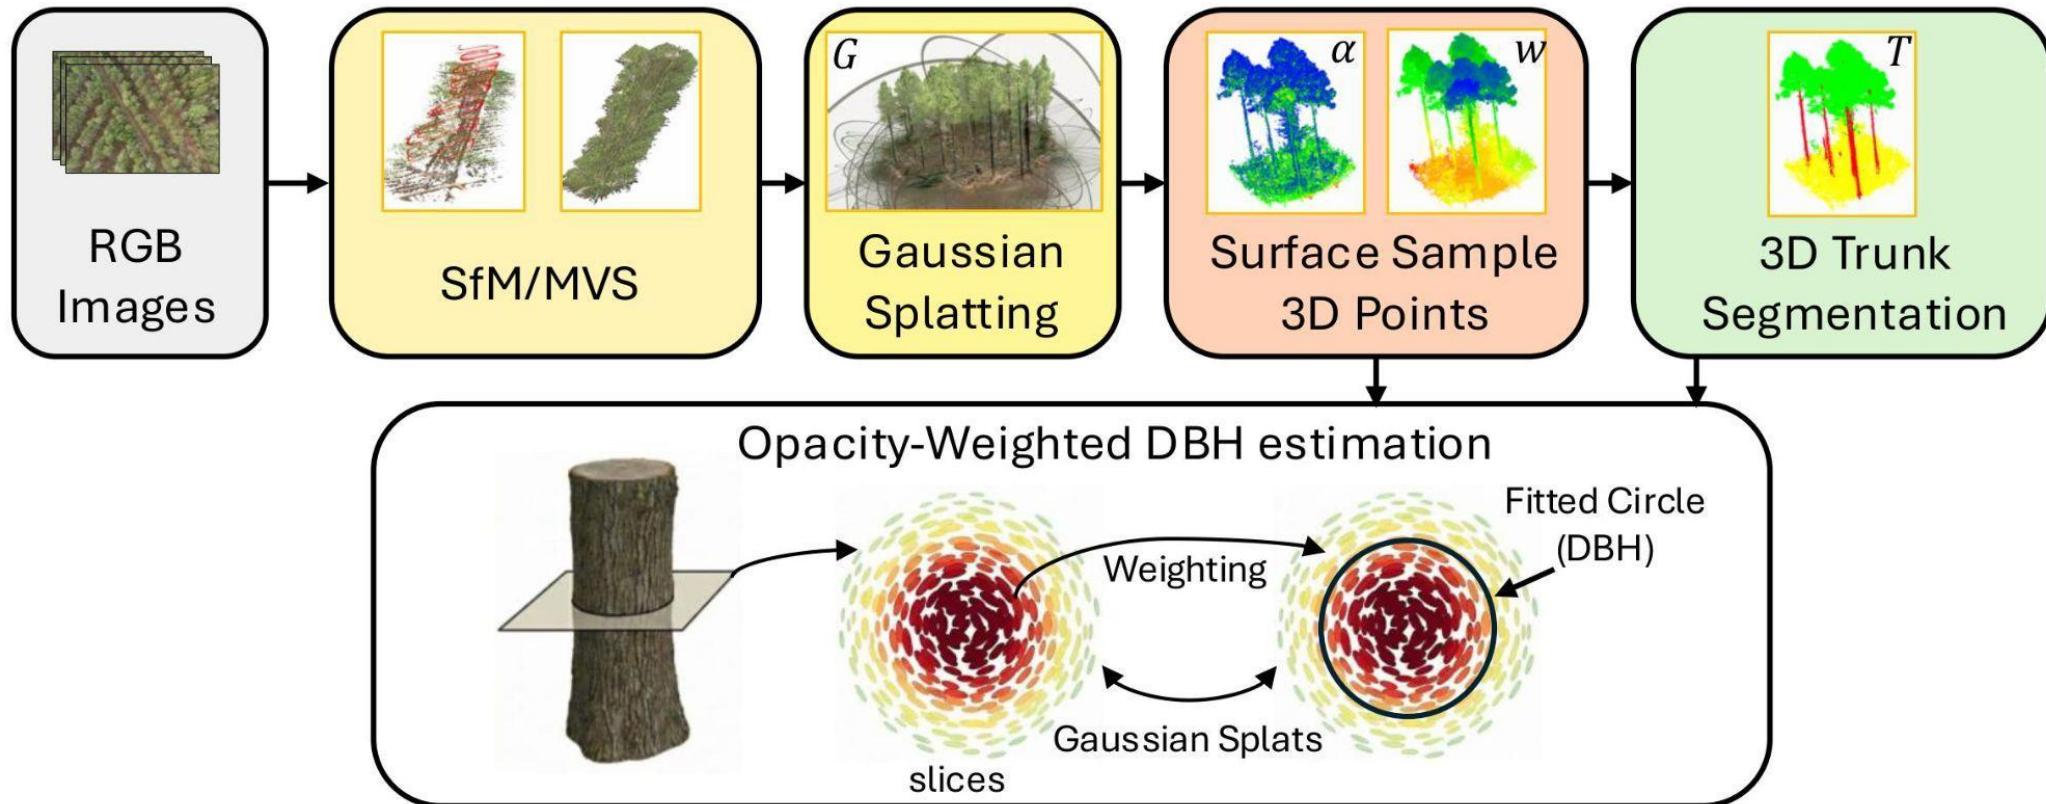






# The idea: Gaussian Splatting as a measurement device, not a renderer

*A continuous, learned 3D density field – with a confidence for every point.*



# The idea: Gaussian Splatting as a measurement device, not a renderer

*A continuous, learned 3D density field — with a confidence for every point.*

Classical 3D Gaussian Splatting learns millions of fuzzy points to render photorealistic views.

We repurpose the same representation as a **geometry and confidence source**. Each 3D Gaussian carries a learned opacity  $\alpha \in [0, 1]$  — a per-point measure of how well multiple photos agree on its presence.

Thin trunks become measurable precisely because we can ask the field for density, instead of asking an image for a pixel.

## Dense

a continuous field, not a point cloud — thin stems fill in

## Confident


opacity  $\alpha$  scores every 3D point by multi-view agreement

## RGB-only

a commodity drone is the only sensor in the loop

# Weighting samples by opacity recovers the trunk

*Points that multiple photos can't agree on get a low opacity — and drop out of the fit.*



*(a) raw sampling — sheet artifacts and floaters*




*(b) opacity-weighted sampling — clean stem geometry*


# TreeDGS fills the trunk where SfM, MVS and LiDAR drop it




(a) SfM



(b) MVS



(c) TreeDGS



(d) LiDAR

## SfM (*traditional photogrammetry*)

Sparse points — most stems missing

## MVS with Deep Learning

Canopy filled, but thin trunks still drop out

## TreeDGS

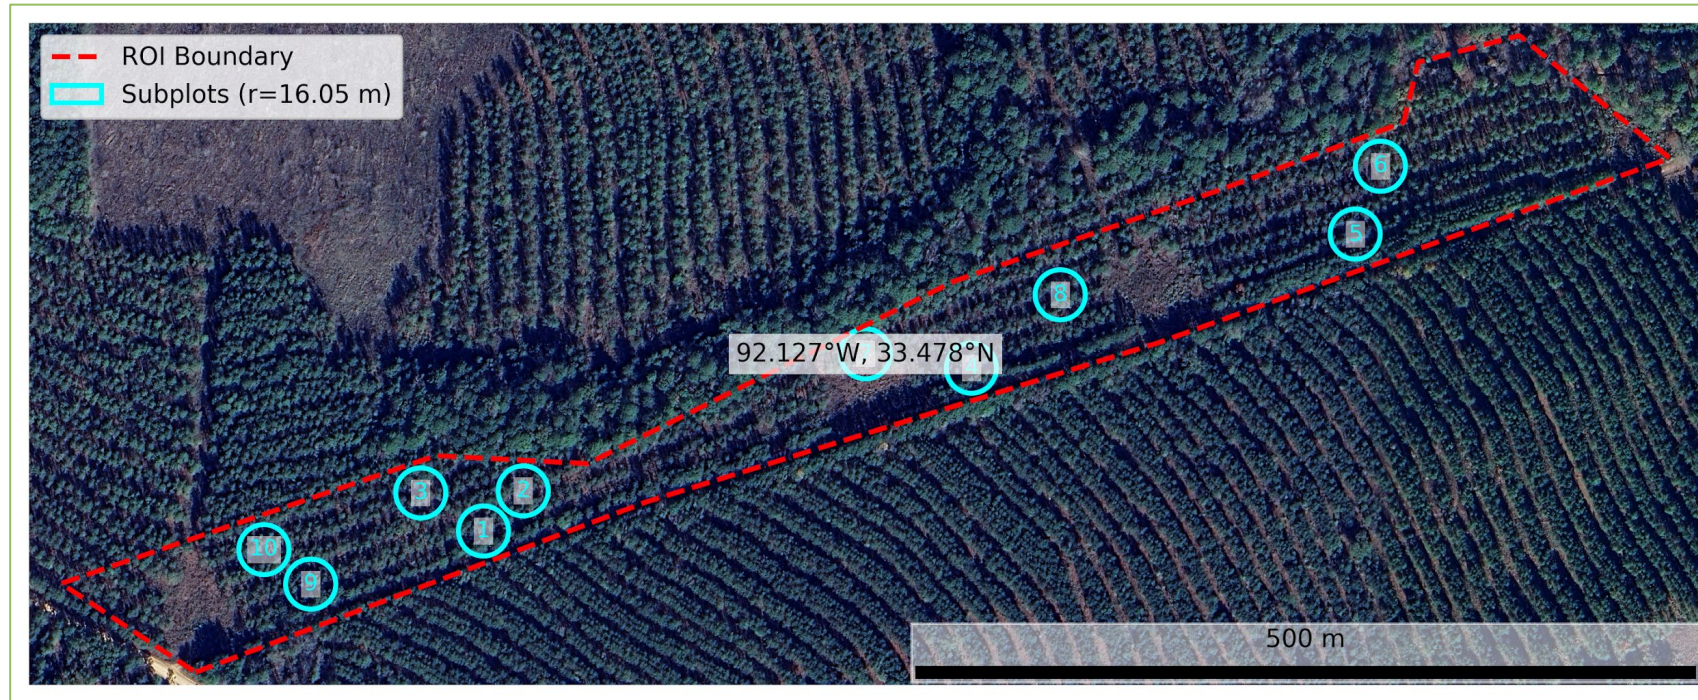
Surface-consistent, measurable trunks

## Airborne LiDAR

Good structure, great ground — but stem returns are sparse

*Same plot, four reconstruction methods — colored by height above ground.*


# Study Design: 10 x 0.2 ac plots across a 15 ac stand



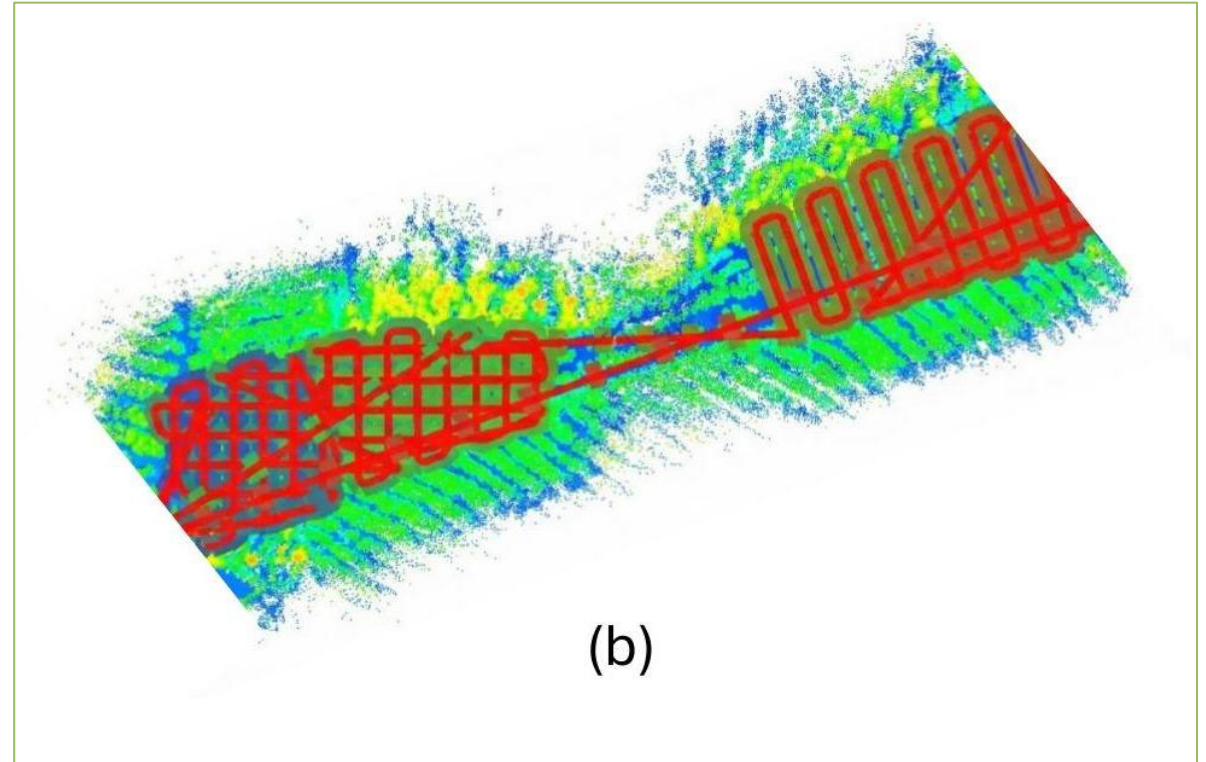
*Region of interest and 10 subplots ( $r = 16.05$  m); plot centers surveyed via RTK GNSS.*

- 210 manually paired stems
- *Pinus taeda*, mixed ages
- RGB: mixed oblique views at 70 / 90 m AGL (DJI Matrice 4E)
- LiDAR: dense lawn mower grid at ~50 m AGL (Hesai XT-32, ~1,407 pts/m<sup>2</sup>)

# Study Design: RGB vs LiDAR flight patterns



*(a) RGB: high overlap, mixed oblique views at 70 / 90 m AGL (DJI Matrice 4E).*

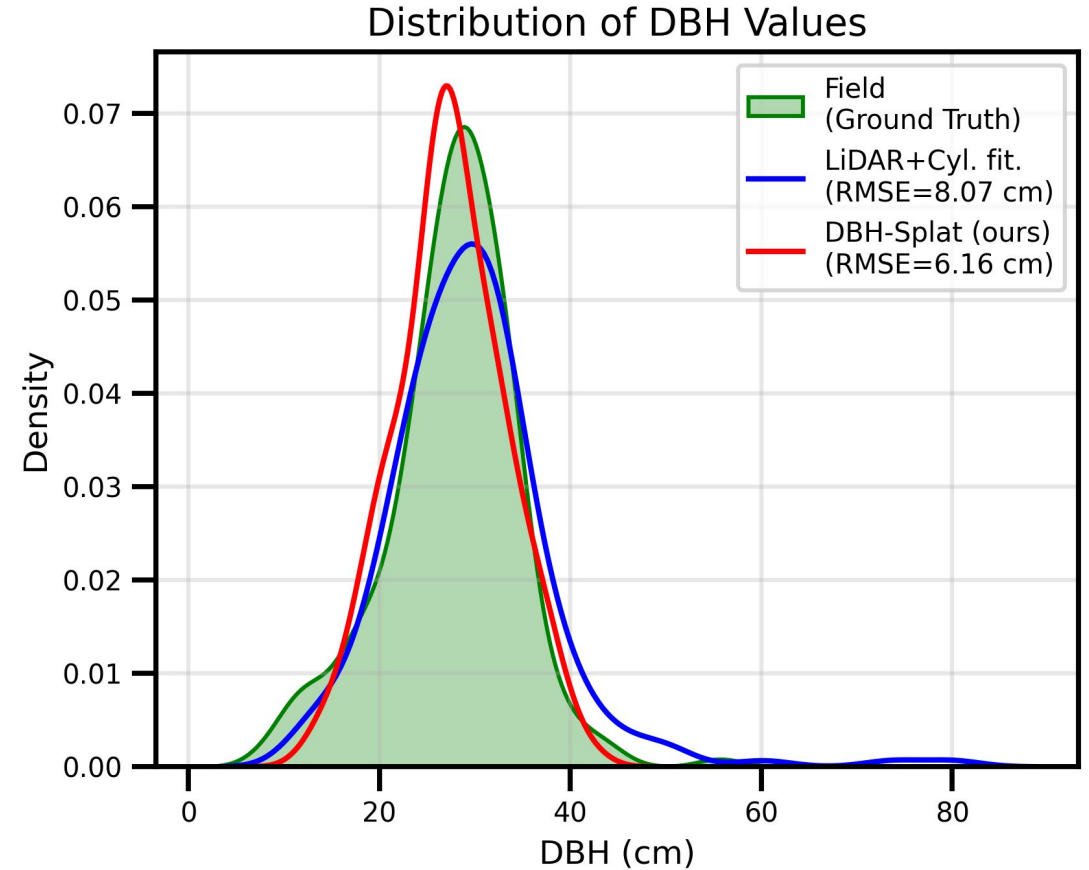


*(b) LiDAR: dense lawn-mower grid at ~50 m AGL (Hesai XT-32, ~1,407 pts/m<sup>2</sup>).*

HEADLINE RESULT

# Across 10 plots and 210 stems, TreeDGS has 39% lower error than aerial LiDAR

*Pooled RMSE on field-measured DBH, managed loblolly pine.*



*210 stems field-tape-measured at 4.5 ft · paired per-tree · 10 × 0.2-acre plots · managed loblolly pine, Arkansas*

HEADLINE RESULT

# Across 10 plots and 210 stems, TreeDGS is 39% tighter than aerial LiDAR

*Pooled RMSE on field-measured DBH, managed loblolly pine.*

TreeDGS — aerial RGB only

1.89 in

pooled RMSE · MAE 1.44 in · trunk detection success 99.7%

Aerial LiDAR + cylinder fit

3.11 in

pooled RMSE · MAE 1.98 in

*210 stems field-tape-measured at 4.5 ft · paired per-tree · 10 × 0.2-acre plots · managed loblolly pine, Arkansas*

# The error is still dropping on our internal benchmarks

*The published result is a checkpoint, not a ceiling.*

Published (this paper)

1.89 in

RMSE · 210 stems, 10 plots

Internal benchmark, today

1.0 – 1.2 in

RMSE · across expanded plots

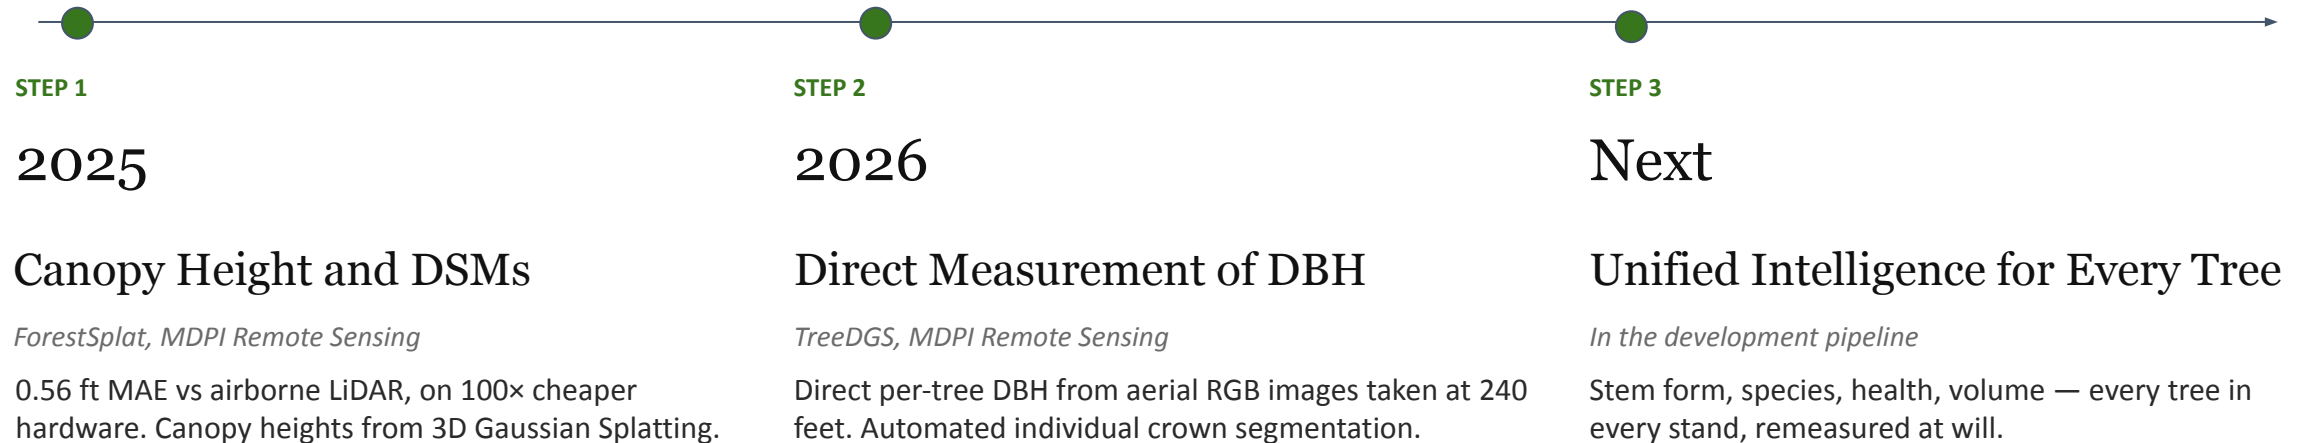
Distributional error

< 0.4 in

Wasserstein vs field distribution

*More data and smarter depth formulations tighten the error without changing the input — TreeDGS has room to run.*

# Our research journey: From surface and structure to a full spatial model



# Where we still struggle, and what's next

*Three honest limitations — each suggests its own follow-up.*

LIMITATION

## Trunk visibility

A fully-occluded stem cannot be recovered by more sampling. This method needs five to eight informative views of the breast-height band. We'd like to drop this requirement to three views.

LIMITATION

## Segmentation

The model requires perfect per tree segmentation. False-positive trunk labels from adjacent understory still bias the fit. Opacity weighting softens this, but does not remove it.

LIMITATION

## Compute

3D Gaussian Splatting optimization is heavier than a classical photogrammetry export. But the limitation is engineering now, not theory. We can process 300 ac / day currently.

*Next: jointly optimize tapered-cylinder trunk primitives inside the Gaussian field — infer DBH directly from model parameters, even under partial visibility.*

# Canopy height was step 1. DBH is step 2. The digital twin is step 3.

## Precision Silviculture

Targeted thinning, fertilization and disease response at the individual-tree level

## Awesome Growth Models

Growth models built from huge amounts of forest level datasets

## Verifiable Carbon Credits

Build trust that can be traced back layers to drone photos, orthomosaics, and 3D models

## Tree Health

See distressed stems, broken tops, and disease patterns at the individual tree level -> early detection and proactive management

## Stem-level Merchandising

Taper along the entire bole: log grades, saw-log yield and cut-to-length plans grounded in the actual stems loggers harvest


# Thank you.

*A big shout out to Jacob Strunk from Potlatch for supporting us on this research*

**Michael Wu**

Co-founder & CEO, Coolant

[michael@coolant.earth](mailto:michael@coolant.earth)



*With thanks to the Coolant team and Professor James Tompkin (Brown University). Paper & preprint available on request.*