

# SilviBot

## Autonomous Ground-Based LiDAR Platform for Complete Forest Census *Product Introduction and Competitive Analysis*

---

Martin Hensley  
February 2026

**Abstract.** *Forest inventory in the United States relies on statistical sampling methods largely unchanged since the 1940s—covering 2–5% of standing stems and producing volume estimates with  $\pm 10$ –20% standard error. The workforce collecting this data faces structural decline: seasonal technicians earning \$19–32/hour, annual turnover exceeding 50%, and measurement quality problems that the profession has long recognized but never resolved. SilviBot is an autonomous terrestrial LiDAR platform designed to replace statistical timber cruising with complete forest census—every merchantable stem measured, georeferenced, and backed by verifiable point cloud data. This paper introduces the SilviBot product concept, describes the terrestrial laser scanning technology that enables it, and positions SilviBot within the competitive landscape of precision forestry solutions.*

## 1. The Forest Mensuration Problem

Forest mensuration—the science of measuring tree and stand attributes—is the quantitative foundation of all forest management. Every timber sale appraisal, silvicultural prescription, carbon stock estimate, and NEPA environmental analysis begins with inventory data: measurements of diameter, height, species, and condition collected in the field and extrapolated across landscapes. The quality of these measurements determines the quality of every downstream decision, from harvest scheduling to carbon accounting under international climate agreements.

The dominant methodology has remained essentially unchanged for eighty years. A crew of two to four technicians traverses sample points using angle-gauge (prism or Bitterlich) sampling, recording diameter at breast height (DBH), species, height estimates, and defect for trees that subtend the critical angle. A standard variable-radius plot cruise on a 100-acre timber sale typically samples 2–5% of standing stems, yielding volume estimates with standard errors of  $\pm 10$ –20%. The USDA Forest Service’s Forest Inventory and Analysis (FIA) program—the nation’s primary forest monitoring system—maintains approximately 355,000 permanent plots re-measured on 5–10 year cycles, investing roughly \$75 million annually in field operations.

The limitation is epistemic: sampling tells you what is *probably* there, within confidence bounds determined by sample intensity, stand variability, and crew competence. For a timber sale appraised at \$500,000, a  $\pm 15$ % confidence interval represents \$75,000 of bid risk—a direct economic cost borne by buyer or seller with no mechanism for resolution.

The workforce collecting this data is in structural decline. Seasonal field technicians at GS-5 through GS-9 pay grades (\$19–32/hour under the 2025 OPM Rest of U.S. locality pay table) work six-month temporary appointments with no health insurance, no retirement benefits, and no path to permanent employment. Annual turnover exceeds 50%. The result is a perpetual cycle of recruitment, training, and attrition that degrades data quality year over year. Industry professionals are familiar with the phenomenon of *gun-decking*—a term borrowed from naval tradition for recording data one did not actually collect—which is not fraud in any meaningful moral sense but the predictable outcome of paying experienced professionals poverty wages and providing no career ladder, no monitoring infrastructure, and no feedback loop between measurement quality and individual outcomes.

The data provenance crisis compounds the problem. Forest management decisions on federal lands carry legal weight. Timber sale Environmental Assessments under NEPA depend on inventory data to support volume estimates and silvicultural justifications—documents subject to administrative appeal and judicial review. Yet the entire evidentiary chain runs through subjective measurement methods with limited quality control. When challenged, the agency’s defense rests on procedural compliance rather than verifiable measurement accuracy. There is no point cloud to revisit. No sensor log to audit. No reproducible measurement to validate.

## 2. Terrestrial LiDAR: The Enabling Technology

Terrestrial laser scanning (TLS) has emerged over the past two decades as a transformative technology for forest measurement. A 2025 review in *Nature Communications* characterized TLS as providing a unique perspective on three-dimensional forest structure, enabling new insights into ecological processes, forest disturbances, and carbon inventories. The technology captures the 3D arrangement of plant components at millimeter-level resolution, fundamentally changing what is measurable in forest environments.

Research-grade TLS systems (RIEGL VZ-1000, FARO Focus) achieve DBH estimation accuracy with RMSE values of 1.0–2.0 cm using multi-scan approaches—substantially better than manual caliper measurements, which carry operator-dependent bias. The International TLS Benchmarking Initiative, involving over 20 national mapping agencies, confirms that multi-scan TLS delivers DTM accuracy at a mean RMSE of 12.7 cm, with tree detection rates exceeding 94% under leaf-off conditions (Liang et al., 2018).

The critical development for commercial viability is the democratization of LiDAR hardware. Apple’s integration of direct time-of-flight (dToF) LiDAR sensors into the iPad Pro and iPhone Pro lines brought 3D scanning capability to consumer price points—approximately \$1,000 for a device that doubles as a forest inventory sensor. Peer-reviewed validation has been systematic: Gollob et al. (2021) demonstrated iPad Pro tree detection rates of 90–97% for stems with DBH >10 cm. Tatsumi et al. (2023) achieved DBH RMSE of 2.27 cm and  $R^2$  of 0.963 against diameter tape reference measurements across 672 trees using ForestScanner on iPhone. A low-cost TLS prototype built from a Livox Avia sensor achieved DBH RMSE of 1.50 cm at a total hardware cost of approximately €2,050—results comparable to commercial scanners costing ten to thirty times more.

Beyond basic inventory parameters, TLS enables Quantitative Structure Modeling (QSM)—complete 3D reconstruction of individual tree architecture from point cloud data. QSM produces estimates of aboveground biomass and stem taper independent of traditional allometric equations, potentially eliminating one of the largest uncertainty sources in forest carbon accounting (Raumonen et al., 2021). This capability is not incremental. It is a different category of measurement.

## 3. SilviBot: From Sampling to Census

SilviBot is an autonomous terrestrial LiDAR platform that replaces statistical timber cruising with complete forest census. The system captures a georeferenced 3D point cloud of every merchantable stem in a stand, then processes that point cloud through an automated pipeline to produce individual-tree measurements of DBH, height, species, taper, and allometric volume—each backed by verifiable sensor data.

### 3.1 Processing Pipeline

The processing pipeline is the core intellectual property. Designed to be hardware-agnostic, it operates identically on point clouds from any terrestrial LiDAR source and executes six sequential stages:

**Ground classification and DTM generation.** Cloth Simulation Filtering (CSF) or progressive morphological filtering separates ground returns from vegetation, producing a digital terrain model that normalizes all subsequent height calculations to local ground level—correcting for slope and micro-topography.

**Individual stem detection and segmentation.** A height-normalized point cloud slice at 1.0–1.7 m above ground isolates the breast-height stem zone. DBSCAN clustering identifies individual stems with detection rates of 90–97% for stems >10 cm DBH on consumer hardware, rising to 99.4% on research-grade personal laser scanners.

**Circle fitting for DBH extraction.** Points in the 1.3–1.4 m height band are projected to 2D and fit using RANSAC or least-squares circle fitting, yielding DBH RMSE of 1.5–2.5 cm depending on sensor quality.

**Tree height estimation.** Maximum point height within each segmented crown, normalized to local DTM, provides tree height. SilviBot’s mobile continuous-scan approach captures stems from multiple angles during traversal, mitigating the canopy occlusion that limits single-scan TLS in dense stands.

**Species classification.** Bark texture features extracted from point cloud intensity and geometric properties (branching pattern, crown shape, stem taper) feed a gradient-boosted classifier trained on region-specific reference datasets. Purdue’s iForester project has demonstrated >97% segmentation accuracy using mobile Segment Anything Model (SAM).

**Allometric volume calculation.** Measured DBH, estimated height, and species assignment are applied to FIA-standard allometric equations to compute per-tree volume. For high-resolution point clouds, QSM reconstruction enables direct volumetric estimation independent of allometric equations—eliminating a major uncertainty source in conventional inventory.

### 3.2 Autonomy Stack

The mobile platform adds SLAM-based navigation for continuous forest traversal: LiDAR-inertial odometry for real-time pose estimation under GNSS-denied canopy, occupancy grid mapping with dynamic obstacle detection, adaptive scan planning that adjusts station density to stand visibility, terrain assessment for safe traversal on slopes up to 35%, and battery-aware mission planning. At its core is an **AI World Model**—a persistent spatial representation that fuses sensor data into a continuously updated digital twin of the forest, enabling autonomous path planning, real-time obstacle avoidance, and predictive decision-making as stand conditions change during traversal.

### 3.3 The Paradigm Shift

The transition from sampling to census is not an incremental accuracy improvement. It is a categorical change in the information available to forest managers. A sensor does not gun-deck. Every data point has a GPS coordinate, a timestamp, a point cloud, and a measured geometry. The chain of custody from raw scan to volume estimate is auditable. A census-derived stem map with measured DBH, estimated height, and georeferenced position for every tree enables harvest planning optimized at the individual tree level, spatially explicit growth modeling calibrated to actual stand conditions, defensible carbon accounting with measurement-level uncertainty propagation, and real-time monitoring of treatment effects through repeat scanning.

### 3.4 Development Roadmap

**Phase 1 — iPad LiDAR MVP (Year 1):** A native iOS application using ARKit for multi-station point cloud capture validates the complete processing pipeline from raw scan to inventory report. Field trials against conventional cruise data establish accuracy benchmarks. Target: >95% detection, <2.5 cm DBH RMSE, <5 minutes per plot. Total hardware cost: approximately \$1,400.

**Phase 2 — Mobile Platform (Years 2–3):** A quadruped robot (Unitree Go2 or equivalent) equipped with a solid-state LiDAR sensor (<\$5,000 total hardware) adds SLAM navigation, automated scan planning, and continuous point cloud registration during movement.

**Phase 3 — Autonomous Census (Years 3–5):** Full autonomy: deploy a unit at a stand boundary, define the survey area, retrieve a complete census inventory upon mission completion. Multi-hour missions, fleet management, and adaptive scan density. Extension to steep-slope Pacific Northwest terrain and international plantation markets.

## 4. Competitive Landscape

The global precision forestry market has been valued at approximately \$5.2–5.8 billion (2024) with projected growth to \$7.9–11.6 billion by 2029–2032 at 5.8–10.0% CAGR. North America represents the largest regional market. The forest inventory data collection segment remains fragmented among consulting firms, government agencies, and emerging technology startups—presenting a greenfield opportunity for a purpose-built census platform. SilviBot is the only solution combining ground-level census accuracy with a development roadmap to full autonomy.

Solution	Type	Cost	DBH Acc.	Automation	Census	Audit Trail
Traditional Cruise	Manual	\$5–15/ac	Variable	None	No	No
RIEGL VZ-1000	Research TLS	\$100K+	~1.0 cm	Post-proc.	Plot	Yes
GeoSLAM ZEB	Handheld PLS	€50K	~1.5 cm	Semi-auto	Plot	Yes
ForestScanner	iPad App	Free	~2.3 cm	Real-time	Manual	Partial
Arboreal Forest	iPad App	\$28/mo	~2.5 cm	Real-time	Manual	Partial
Trefos (Drone)	Airborne	\$10–50K	Canopy	Flight plan	Overstory	Yes
Treemetrics	Platform	Enterprise	Multi-src	Cloud	Remote	Partial
GreenValley Int'l	HW+SW	Premium	Integrated	Cloud	Plot	Yes
<b>SilviBot</b>	<b>Auton. TLS</b>	<b>\$5K HW</b>	<b>&lt;2.5 cm</b>	<b>Full auto</b>	<b>Yes</b>	<b>Full</b>

Table 1. Competitive comparison of forest inventory technologies.

## 4.1 Differentiation by Category

**Against traditional cruising.** SilviBot eliminates sampling error, replaces subjective human measurement with sensor-verified data, produces an auditable point cloud record, and removes workers from hazardous steep-slope terrain. At a target cost of ~\$8/acre for census-level accuracy versus \$15–25/acre for  $\pm 15\%$  statistical estimates, the value proposition is immediate for any organization purchasing forest inventory services.

**Against drone and airborne LiDAR.** Drone-based systems (Trefos, NMI ForestView) capture overstory canopy structure from above but cannot reliably measure DBH, lower stem taper, or understory conditions due to canopy occlusion. DBH is the single most important variable in timber volume estimation—it enters volume equations as a polynomial function where a 5% diameter error produces roughly 10% volume error. Ground-level measurement is not optional; it is definitional for operational forest inventory.

**Against research-grade TLS.** Systems like the RIEGL VZ-1000 (\$100K+) and FARO Focus achieve excellent accuracy but require manual setup, manual operation, and post-processing by trained specialists. They are scientific instruments, not operational inventory tools. SilviBot targets equivalent measurement quality on a \$5,000 commodity robotics platform with autonomous operation—a 20:1 cost reduction that transforms TLS from a research method into a production workflow.

**Against mobile apps (ForestScanner, Arboreal).** These iPad-based applications validate the core technical feasibility that SilviBot builds upon—ForestScanner’s peer-reviewed  $R^2$  of 0.963 for DBH demonstrates that consumer LiDAR works for forest measurement. However, these apps lack an autonomy pathway, species attribution, allometric volume calculation, and commercial platform architecture. SilviBot extends the validated measurement capability into a complete, end-to-end autonomous inventory system with full data provenance.

**Against platform integrators (Treemetrics, GreenValley).** These companies aggregate data from multiple sensor modalities—satellite, aerial, terrestrial—into analytics platforms. Their approach is valuable for landscape-scale monitoring but does not solve the ground-truth data collection problem. SilviBot is complementary: it produces the high-accuracy ground-level census data that these platforms need as calibration input but cannot generate on their own.

## 5. Comparative Advantage

SilviBot's comparative advantage derives from three interconnected elements.

**Processing pipeline IP.** Algorithms for stem detection, cylinder fitting, DBH extraction, taper modeling, species attribution, and allometric volume calculation—optimized for consumer-grade point cloud data. The pipeline is hardware-agnostic by design: it survives sensor commoditization because the software is the product, not the hardware. Any terrestrial LiDAR that produces a point cloud feeds the same pipeline.

**Forest autonomy stack.** SLAM navigation, terrain assessment, scan planning, and obstacle avoidance purpose-built for unstructured forest environments. Forest navigation requires handling downed wood, slash, dense understory, and steep slopes—conditions that no existing commercial autonomy stack addresses. The AI World Model that drives this stack—fusing real-time sensor data into a persistent spatial representation for planning and prediction—is a fundamentally different engineering problem from warehouse or sidewalk robotics.

**Calibration dataset.** Every deployment produces paired point cloud and ground-truth measurements across diverse forest types, creating a growing data asset that improves model accuracy over time and cannot be replicated without equivalent field deployment at scale.

## 6. Revenue Model and Market Opportunity

SilviBot targets four complementary revenue streams: **Software sales** of the processing pipeline as a purchased application for organizations using their own LiDAR hardware; **Hardware-as-a-Service (HaaS)** leasing of autonomous census units at \$2,000–5,000/month for agencies and large-program operators; **Census-as-a-Service (CaaS)** providing turn-key inventory at \$3–8/acre versus \$5–15/acre for traditional sampling; and **data licensing** of aggregated forest structure datasets for model calibration, carbon market applications, and remote sensing validation.

The total addressable market spans federal forest inventory programs (~\$105M annually), state agencies (~\$200M), domestic private sector operations including timber companies, TIMOs, REITs, carbon MRV, and consulting firms (~\$700M), and international markets including Canadian, European, and Australasian forestry (~\$1.2B)—totaling approximately \$2.2 billion globally. SilviBot's realistic addressable share is estimated at \$165 million in annual revenue at market maturity, with total development investment of \$1.8–3.25M through commercialization over 3–5 years.

## 7. Conclusion

The fundamental question of forest inventory has always been: *what is actually growing on this ground?* Every existing method answers with an estimate. SilviBot answers with a measurement.

The convergence of consumer-grade LiDAR, commodity quadruped robotics, SLAM navigation, and AI-driven spatial reasoning has made complete forest enumeration technically feasible and economically viable for the first time. The processing pipeline that converts raw point clouds into auditable inventory reports is the core innovation; the hardware is commodity infrastructure that will only improve and cheapen with each product generation.

SilviBot does not replace experienced foresters. It lets them do what they are actually good at—making silvicultural prescriptions, laying out sales, managing stands—using data they can trust, without destroying their bodies to collect it. The soul of forestry is not a nineteen-year-old with a semester of dendrology walking a chain on a timber sale for \$19/hour with no benefits. The soul of forestry is the person who looks at a stand and understands what happened there, what’s happening now, and what it needs next. SilviBot makes that person’s expertise more valuable, because the entire inventory pipeline depends on their ability to calibrate and validate what the sensor measures.

---

## Selected References

- Gollob, C., Ritter, T., & Nothdurft, A. (2021). Measurement of forest inventory parameters with Apple iPad Pro and integrated LiDAR technology. *Remote Sensing*, 13(16), 3129.
- Liang, X., et al. (2018). International benchmarking of terrestrial laser scanning approaches for forest inventories. *ISPRS Journal of Photogrammetry and Remote Sensing*, 144, 137–179.
- Raumonen, P., et al. (2021). Terrestrial laser scanning: a new standard of forest measuring and modelling? *Annals of Botany*, 128(6), 653–662.
- Tatsumi, S., et al. (2023). ForestScanner: A mobile application for measuring and mapping trees with LiDAR-equipped iPhone and iPad. *Methods in Ecology and Evolution*, 14(7), 1603–1614.
- Expanding forest research with terrestrial LiDAR technology. (2025). *Nature Communications*, 16, 8742.
- Purdue University HTIRC. (2024). iForester: AI-assisted smartphone app for automated tree inventory.
- USDA Forest Service. (2024). Forest Inventory and Analysis National Program.