

Building Forest Inventories with Autonomous Legged Robots — System, Lessons, and Challenges Ahead

Matías Mattamala¹, Nived Chebrolu¹, Jonas Frey^{2,3}, Leonard Freißmuth⁴, Haedam Oh¹, Benoit Casseau¹, Marco Hutter², Maurice Fallon¹

¹Oxford Robotics Institute, Department of Engineering Science, University of Oxford, Oxford, UK

²Department of Mechanical and Process Engineering, ETH Zurich, Zurich, Switzerland

³Max Planck Institute for Intelligent Systems, Tübingen, Germany

⁴Technical University of Munich, Munich, Germany

Corresponding author: Matías Mattamala (email: matias@robots.ox.ac.uk).

This work is supported in part by the EU Horizon 2020 Project 101070405 (DigiForest) and a Royal Society University Research Fellowship.

ABSTRACT Legged robots are increasingly being adopted in industries such as oil, gas, mining, nuclear, and agriculture. However, new challenges exist when moving into natural, less-structured environments, such as forestry applications. This paper presents a prototype system for autonomous, under-canopy forest inventory with legged platforms. Motivated by the robustness and mobility of modern legged robots, we introduce a system architecture which enabled a quadruped platform to autonomously navigate and map forest plots. Our solution involves a complete navigation stack for state estimation, mission planning, and tree detection and trait estimation. We report the performance of the system from trials executed over one and a half years in forests in three European countries. Our results with the ANYmal robot demonstrate that we can survey plots up to 1 ha plot under 30 min, while also identifying trees with typical Diameter at Breast Height (DBH) accuracy of 2 cm. The findings of this project are presented into five lessons and challenges. The lessons discuss the achievements of the trials, the robustness achieved by the locomotion and state estimation systems, and the benefits of testing in varied environments. The challenges present the key issues we faced including the limitations of our sensing system, the need for more integrated perception and control algorithms, and the applicability of these solutions for other real-world tasks. By sharing these lessons and challenges, we offer insight and new directions for future research on legged robots, navigation systems, and applications in natural environments. Additional videos can be found in <https://dynamic.robots.ox.ac.uk/projects/legged-robots>

INDEX TERMS Autonomous Robots, Environmental Monitoring, Forestry, Legged Robots, Simultaneous Localization and Mapping (SLAM)

I. INTRODUCTION

SYSTEMATICALLY mapping forests is an important task in modern forestry [1] and environmental sciences [2]. These maps can be used to understand how a forest contributes to global biomass estimates, the impact of natural and human-induced hazards, and to enable sustainable management within the forestry industry. Obtaining these maps is a laborious process. While remote sensing techniques such as uncrewed platforms or satellites [3] have scaled the extent of

this task [4], ground measurements are still required under the canopy, to determine fine-level tree traits or calibrate forest models from aerial data [5].

Acquiring ground-level measurements is a laborious process. Traditional techniques used in forestry and environmental sciences rely on tape measures, calipers, and hand-held laser rangefinders [6]. Nowadays, technologies such as Terrestrial Laser Scanning (TLS) [7] provide high-accuracy, dense reconstructions of forest plots, further enabling the



FIGURE 1. We deployed our autonomous legged robot system in three different countries—Finland, Switzerland, and the United Kingdom—in five campaigns between 2023 and 2024. By testing in different types of forests—mixed, deciduous, and coniferous—across different seasons, we gained valuable insights into the current capabilities of legged platforms for forestry applications.

study of biomass estimates and tree structure. However, collecting TLS data still requires important human efforts in the field, therefore it is constrained to forest plots pre-defined by foresters [5].

Robotic solutions have the potential to scale up the mapping process even further. Core technologies such as Simultaneous Localization And Mapping (SLAM) have enabled online, large-scale 3D mapping. Handheld SLAM devices have been used to produce ground-level 3D reconstructions of forest plots [8], [9], which in forestry is commonly known as Mobile Laser Scanning (MLS). Further, small-scale drones [10], [11] and wheeled platforms [12] have also been used for under-canopy forest inventory missions.

Legged robots can potentially provide significant advantages to aerial and wheeled platforms in terms of mobility, endurance, and lower soil impact compared to tracked platforms [13]. Their advanced mobility allow them to execute tasks in environments designed for humans, such as industrial sites [14], but also outdoors. Recent applications in crop fields and vineyards [15], as well as open-field scientific sampling applications [16], have demonstrated their potential.

In this work, we study the application of legged platforms for forestry applications, particularly for *forest inventory*. From a forestry perspective, the ability to autonomously survey forest plots from the ground level can provide data

that is not captured from Aerial Laser Scanning (ALS) or satellite images. Automating this process with autonomous robots has the potential to facilitate and scale up ground-level surveying, which is otherwise conducted with TLS and MLS. From a robotics perspective, this requires an integrated autonomy stack, in alignment with the high-level surveying objectives. The system should be able to operate in a highly unstructured, natural environment with minimal intervention by humans.

Our main contribution is an integrated autonomy system for online forest inventory with legged robots and an evaluation of its performance in diverse forest environments. Our solution builds upon a rugged legged platform—the ANYbotics ANYmal, an integrated multi-sensor unit, and state-of-the-art components from locomotion systems to navigation and forest inventory algorithms. We present the development and improvement of the autonomy system, through five field campaigns executed between May 2023 and July 2024 in three European countries—Finland, Switzerland, and the UK—in the context of the DigiForest EU Horizon project¹, shown in Fig. 7. On each campaign we deployed a legged platform autonomously in the field, mapping areas up to 1 ha.

A publicly available report [17] presented a preliminary report on this system. In this work we extend that analysis with additional results from two new campaigns conducted in

¹<https://www.digiforest.eu>

Wytham Woods, UK, as well as Stein am Rhein, Switzerland. We report on the the autonomy performance, mapping, and tree inventory capabilities observed in the field. Lastly, we discuss the findings and lessons learned throughout the campaigns.

II. RELATED WORK

We first review state-of-the-art robotics systems for forestry applications and environmental sciences. Then we focus on the uses of legged platforms for field applications in industrial sites, environmental monitoring, and forestry.

A. Robotics for Forestry and Environmental Applications

Robotics has the potential to upscale environmental measurement through technologies and platforms that can automate intensive or time-consuming manual labor.

Mapping technologies such as **SLAM** can accelerate ground-based forest measurements by integrating LiDAR scans from a mobile sensing payload. The produced maps are post-processed to extract the desired attributes, such as the tree positions and **DBH** [8], [9], [18]. The incremental nature of **SLAM** and the data acquisition process has motivated the concurrent construction of maps and inventories. Chen et al. [10] and Proudman et al. [19], demonstrated real-time tree detection and **DBH** estimation using small-scale quadcopters and handheld devices, respectively. In this work, we explore this idea by developing an autonomy system that builds upon similar online forest inventory approaches [20].

On the other hand, robotic platforms can support surveying tasks by freeing humans from collecting the data themselves [21]. In the context of aerial surveying and **ALS**, autonomous fixed-wing aircraft have been used to obtain Light Detection and Ranging (LiDAR) point clouds [22]. Other platforms such as small-scale helicopters [23] and more recently hexacopter platforms [24] have been used for over-canopy data collection. Since these platforms fly over the trees, they do not have to deal with challenges such as collision-free navigation. Instead, the questions concern how to survey a defined area, maximizing coverage given energy [25] and time constraints [26].

Using robotic platforms for under-canopy data collection presents more challenges, as it requires robots to operate in cluttered, unstructured spaces. For aerial platforms, it is not possible to use fixed-wing aircraft, hence smaller quadcopters or hexacopters are preferred. Teleoperated small-scale drones have been used for more than ten years for forest surveying [27], [28]. More recently, some applications of autonomous drones have been developed for forest inventory [10], [11] and DNA sampling [29]. Nevertheless, their small form-factor introduces additional constraints to their capabilities, since the payloads and battery size must be reduced, consequently limiting their autonomy, the complexity of the algorithms, and the extent of the missions they can accomplish [30].

An interesting alternative to aerial platforms is tethered systems, such as SlothBot [31]. This has been designed for long-term monitoring of forests with minimal energy requirements. Other approaches combined aerial platforms with a tether to navigate through forest canopies [32], overcoming some of the navigation challenges that a purely aerial platform would face. Nonetheless, many challenges still exist for the implementation and adoption of such systems in monitoring applications.

Ground-based platforms are a competitive approach to providing under-canopy measurements while carrying heavier sensing payloads, enabling longer surveys for forestry or monitoring applications. Wheeled and tracked robot platforms have been used for forest data collection for more than 15 years [8], [33]. These approaches mainly rely on LiDAR **SLAM** solutions to map the forest [12]. Larger-scale wheeled platforms have been demonstrated in the forestry domain to support harvesting operations and achieve autonomous tree cutting [34]–[36]. However, ground platforms have larger footprint and mass, and by using rugged wheels and tracks they can damage the soil via trampling effects [37], [38].

In this work, we propose to use lighter, legged robotic platforms for forestry survey tasks—up to 50 kg. The focus on smaller platforms is motivated by emerging global forestry policies that advocate for minimal impact, sustainable approaches, such as continuous cover forestry [39].

B. Field Applications of Legged Robots

Since our focus is on legged platforms, we discuss their applications in different fields. From an industrial perspective, legged platforms are being used for inspection and monitoring potentially dangerous or remote sites. For example, they have been deployed in offshore oil and gas facilities [14], [40], [41], where their advanced mobility capabilities allow them to navigate complex, human-made environments without requiring the personnel on-site. Nuclear facilities represent another important area where these robots are being deployed [42].

Exploring underground environments is a key application for mining and disaster response. The recent DARPA SubT Challenge [43], held from 2017 to 2021, demonstrated the advantages of legged platforms in overcoming most of the challenges of these unstructured environments. In fact, the majority of the teams participating in the finals used legged platforms as part of their solutions [44], [45], including the winning team [46].

In open-field applications, legged platforms have been used for agricultural monitoring [15], [47], litter collection in beaches [48], as well as load-carrying tasks [49]. A significant achievement was the autonomous deployment of the LS3 robot over hundreds of kilometers in forests and other natural environments [50].

For scientific applications, legged platforms have been deployed for soil sampling [51] and measurement of aeolian processes [16]. The K-Rock amphibious robot [52] was used




FIGURE 2. Our vision for autonomous forest inventory with legged robots. Our system autonomously drives a legged platform in an unknown forest plot, while recording environmental data used for tree segmentation and trait estimation.

for wildlife filming of crocodiles in Africa; the ANYmal robot has also been used to study social interactions with monkeys [53]. Potential applications for space exploration have also been studied in planetary analog settings [54].

Forestry applications of legged platforms have been discussed since the eighties [13]. Their potential advantages in terms of heavier payload carrying and lower soil impact have been the main advantages for their use. Despite this potential, the field remains under-explored, with few real-world demonstrations. The HEAP hybrid wheeled-legged platform, adapted for tree harvesting, is one of the few examples [35]. In this work, we provide an experimental demonstration of the capabilities of smaller size legged platforms for mapping and inventory tasks in forests.

III. SYSTEM

A. Overview

Our goal is to develop a system for autonomous forest inventory using a 50 kg legged platform. Our vision for this system is presented in Fig. 2: A human operator defines an unexplored patch of the forest—a *plot*—they are interested in surveying, which is the input for the mission planning system. Based on this input, the mission planner proposes a preliminary plan for the robot to navigate the environment, given some specifications, e.g., maximizing coverage and consistency (via loop closures).

The robot executes the mission plan using a reactive local planning approach to negotiate obstacles and other unknowns of the environment. While the local planner ensures safe navigation, it also communicates any discrepancies between the planned route and actual surroundings back to the mission planner, enabling dynamic global re-planning of the survey mission as needed. The main data collected by the legged robot consists of 3D scans, which are continuously integrated by the SLAM system running onboard. These scans are used for online forest inventory to determine attributes such as the

DBH of the surrounding forests, and reported in real-time to the operator, providing live feedback on the mission.

The mission concludes with the robot returning to the starting position. At this point, the complete dataset can then be retrieved, analyzed, and refined in post-processing to generate the forest inventory. The inventory is the main input for the Decision Support System (DSS) used by foresters to manage a woodland [55].

Fig. 3 shows the system architecture of our solution, involving state estimation, legged autonomy, and forest inventory. The specific tasks and modules are explained in the following sections.

B. Robotic Platform

In this work, we used the ANYbotics ANYmal platform, specifically *ANYmal C* (released 2020) and *ANYmal D* (released 2022), which are shown in Fig. 4 along with their main sensors.

For ANYmal C, we used the configuration used by Team CERBERUS in the DARPA SubT Challenge [46]. The default sensors (Epson M-G365 Inertial Measurement Unit (IMU), Velodyne VLP-16 LiDAR, and 4 Intel Realsense cameras) were preserved and an additional Alphasense multi-camera unit was added (7 cameras). Networking was limited to the default WiFi communication antenna. Computing resources on the ANYmal C included two CPUs—the Locomotion PC (LPC) and Navigation PC (NPC)—along with a Jetson Xavier GPU. This configuration was used during the 2023 campaigns.

For ANYmal D, we used the default sensors available on the robot, along with the LPC and NPC. No GPU is available in this configuration. We integrated a multi-sensor payload, *Frontier*, which consisted of an Intel NUC CPU as well as a customized Alphasense multi-camera unit (3 cameras) and a Hesai QT64 Wide Field-of-View (FoV) LiDAR). Additionally, we introduced a Rajant Breadcrumb module which improved networking in the forest via a mesh network. This configuration has been used for all the campaigns starting in 2024.

Lastly, an additional Operator PC (OPC) laptop was used to run the graphical interface and to command the robot, via WiFi or the Breadcrumb mesh network.

C. State Estimation

Our state estimation system has evolved significantly since the first campaign in May 2023. Here, we present our current system used in deployments from late 2023 onward, while the former approach is briefly discussed in the corresponding campaign in Sec. IV.

Our current solution involves four modules: (1) LiDAR-Inertial Odometry, (2) LiDAR SLAM, (3) dense mapping, and (4) local terrain mapping, which provides representations for forest inventory as well as navigation. Modules (1)–(3) were executed on the NPC for ANYmal C, or on the Frontier

unit for the ANYmal D setup. Module (4) was executed on the Jetson board, or on the [NPC](#).

1) LiDAR-Inertial Odometry

For odometry, we used a LiDAR-inertial factor graph-based odometry solution, VILENS [56], which provides high-frequency pose and velocity estimates (~ 200 Hz) in a fixed inertial frame. The system performs sensor fusion of relative pose displacements from LiDAR using the Iterative Closest Point (ICP) algorithm [57], aided by preintegrated IMU measurements [58].

We did not use proprioceptive sensing from joint encoders in our estimator, as we observed that the LiDAR-inertial approach was effective in forest environments. The tree distribution and the ground points provided sufficient constraints to localize the robot reliably for navigation purposes.

2) LiDAR SLAM

We used a factor graph-based LiDAR SLAM system [59] to obtain a consistent estimate of the robot pose in a global fixed frame. Our SLAM system incrementally builds a pose graph, where each node stores a LiDAR scan and 6 DoF pose. The edge factors in the pose graph are formed from incremental odometry estimates from the LiDAR-Inertial system, spaced approximately 1 m apart.

To mitigate drift as the robot navigates, we incorporated loop closure factors. For this, we detect loop closures by checking the most recent scan against those linked to pose graph nodes within a radius of 10–15 m from the robot. The candidate loop closures are geometrically verified using ICP, and the successful candidates are added as additional factors in the pose graph.

The pose graph is optimized online in an incremental fashion, using the Incremental Smoothing and Mapping (iSAM2) algorithm [60]. Since the LiDAR scans are rigidly attached to the poses, the 3D map is automatically corrected when the poses are optimized. However, we observed that while this representation was sufficient for localization purposes, it was not dense enough for detecting trees. We then introduced an additional dense mapping module to provide more complete representations without impacting the performance of the lightweight pose graph [SLAM](#).

3) Dense Mapping

In order to exploit the continuous data stream from the LiDAR, we implemented a local sub-mapping module. This module accumulates all the incoming LiDAR scans within a fixed-distance window, generating dense scans in the sensor frame, which we refer to as *data payloads*. These data payloads are gravity-aligned and published at a lower frequency (~ 0.1 Hz), accumulating clouds for about every 20 m traveled. The data payloads are the main input for carrying out online tree traits analysis (see Sec. B).

4) Local Terrain Mapping

In addition to the large-scale mapping modules, we also perform local terrain mapping and use its output as the main representation for local planning. In our deployments with ANYmal C, we used a GPU-based implementation running on the onboard Jetson computer [61]. With ANYmal D we used the default terrain mapping available in the manufacturer's software stack, which runs on the [NPC](#). In both cases, we used a 2.5D grid-based representation with a resolution of 4 cm.




FIGURE 3. System overview. The Autonomy system executes the mission from a survey reference provided by a human operator. State estimation provides a consistent scene representation and dense clouds used for estimating tree traits. Forest inventory segments the point cloud data and estimates tree traits from point cloud data. The main output of the system is a forest inventory database with the main attributes of the surveyed forest.




FIGURE 4. Different configurations of the legged platform used across the campaigns. *Left:* ANYmal C configuration, using the same sensor payload of the DARPA SubT Challenge. *Right:* ANYmal D configuration, using the *Frontier* multi-sensor unit.

D. Autonomy System

Our autonomy system is hierarchical, with multiple levels, which span from the user interface for the human operator (top level) to the locomotion algorithms (lower level).

1) Survey Interface

The survey interface operates as the highest level of planning within our system, allowing the human operator to define the robot deployment. We implemented it as a graphical user interface (GUI) on RViz using interactive markers, which runs on the *OPC*. Given the robot's initial pose, the GUI allows the operator to define a survey area and initialize a *survey plan*, given by a boustrophedon decomposition, commonly known as the lawn mower pattern [62]. A similar strategy has been established for *TLS* data collection protocols [5].

The survey plan is constructed as a sequence of equally-spaced 6 DoF waypoints. The pattern ensures multiple loop closures by enforcing a minimum distance constraint between neighboring path segments, which leads to a more consistent map. This plan does not encode the traversability of the plot or the reachability of the waypoints, which is handled by the lower levels of the autonomy system.

The survey interface also provides mechanisms to interrupt the mission at any time, which triggers a safe stopping behavior via the lower levels. It also enables the operator to modify and resume the plan from any specified goal.

2) Mission Planner

The mission planner manages the execution of the survey plan proposed by the operator. It was executed on the *NPC* for ANYmal C, and on the *Frontier* for ANYmal D.

The main task of the mission planner is scheduling the waypoints to be followed by the local planner, and executing any high-level action from the operator, such as safety stops. Additionally, it processes the progress signals from the local planner to determine if a waypoint is unreachable.

In these cases, it triggers a re-planning strategy proposing alternative waypoints.

3) Local Planning

The local planner generates velocity commands for the low-level locomotion controller to reach a specified waypoint. To achieve this, it relies on the local scene representation built by the local terrain mapping module.

First, a traversability estimation system processes the local terrain map and extracts geometric features to score the terrain's navigability. The traversability score s_{trav} is specifically tailored to enable navigation in forest environments with branches and twigs, although visual traversability approaches could also be used (see Challenge 2 in Sec. VII).

This traversability information is used by a reactive local planner [63], which combines different vector fields for navigation. We extended the approach to compute a Geodesic Distance Field (GDF) on a continuous cost-to-go $c_{\text{to-go}}$ obtained by the heuristic:

$$c_{\text{to-go}} = w_{\text{trav}} (1.0 - s_{\text{trav}}) + w_{\text{unkn}} s_{\text{unkn}}, \quad (1)$$

where s_{unkn} is a fixed score assigned to empty (unknown) cells in the terrain map, and w_{trav} , w_{unkn} are user-defined weights balancing the influence of known and unknown cells.

The local planner continuously monitors progress toward each waypoint to detect potential unfeasible situations, for example if the proposed waypoint is in the middle of bushes. If a waypoint is determined unreachable, the local planner reports this situation to the mission planner, triggering a re-planning of the next waypoint.

Both the local planner and the traversability analysis modules were always executed on the *NPC* of all the versions of the ANYmal we used.

4) Locomotion Controller

The local planner outputs a 3 DoF velocity command—linear (x, y), angular (yaw)—, which is executed by a low-level




FIGURE 5. Survey Interface: GUI for the operator, implemented on RViz using interactive markers. The operator defines the area to be surveyed, while the survey plan is automatically proposed by the Mission Planner. The grey point cloud displayed corresponds to the *data payload*.

Reinforcement Learning (RL)-based locomotion controller. In our experiments, we used both the perceptive controller presented by Miki et al. [64], as well as the blind, learning-based controller from ANYbotics. We used them with no additional modifications to the architecture or training environment to operate in the forest. In both cases, these ran on the [LPC](#).

E. Online Forest Inventory

Our system integrates a state estimation-driven forest inventory pipeline, presented in Freißmuth et al. [20]. The approach achieves near real-time tree segmentation and reconstruction during the mission execution. The system was only tested on the ANYmal D campaigns, where it was executed as part of the Frontier software stack.

The main steps are outlined below; for additional technical details, please refer to the original publication [20].

1) Terrain Modeling

The input for the online forest inventory step consists of data payloads generated at 0.1 Hz (see Sec. C). For each payload, we extract a terrain model using the *cloth simulation filtering* method by Zhang et al. [65]. Please note that this terrain model is different to the local terrain mapping system used for navigation. Here, the terrain model extracted from the data payload enables us to segment the points belonging to the terrain from the potential tree points.

2) Tree Segmentation

To segment trees, we use a method inspired by Cabo et al. [66]. We first normalize the payload cloud z-values using

the terrain model, and then filter the data to retain only points within a slice where tree foliage is not expected. We then use a Voronoi-based clustering to obtain tree cluster candidates. We refine the points associated to each tree by fitting a cylinder model and retaining the points within a distance from the cylinder's surface. Finally, the filtered tree points per tree candidates are de-normalized by height to obtain tree candidate clouds in the original reference frame.

3) Spatio-temporal Aggregation

Our system is designed to incrementally build a forest inventory, effectively improving the tree estimates and terrain model as the robot autonomously collects more data. We exploit the pose graph representation of the LiDAR-SLAM system, by attaching local terrain models and tree instance clouds to it, and updating them with new data after adding new poses or loop closure updates. Specifically, the global terrain model is updated by locally averaging the points around the latest terrain map using a weighted average. The tree instance clouds are managed accordingly to decide if new instances correspond to an existing tree (due to proximity) or a novel instance. We additionally check coverage conditions to ensure that high points are captured, and even coverage from different viewpoints.

4) Tree Reconstruction and Trait Estimation

Lastly, we estimate the traits required by the forest inventory. To *reconstruct* each tree, circles are fitted along the stem to form a representation of the trunk as a series of oblique cone frustums. From these reconstructions, we recover geometric traits such as the [DBH](#) and tree height. The traits are

Campaign Location	Mission	Date	Area [m^2]	Robot Setup		Sensing	State estimation	Autonomy	Tree Analysis
				Platform					
Evo, Finland	Evo-01	2023-05-03	40×25	ANYmal C		Velodyne VLP-16, Epson IMU	CompSLAM odometry	Perceptive locomotion	None
	Evo-02	2023-05-03	40×25						
	Evo-03	2023-05-04	40×35						
	Evo-04	2023-05-04	35×35						
	Evo-05	2023-05-05	70×25						
Wytham Woods, UK	WyO-01	2023-10-06	20×20	ANYmal C		Velodyne VLP-16, Epson IMU	VILENS odometry , VILENS-SLAM	Perceptive locomotion	None
Forest of Dean, UK	Dea-01	2024-02-19	125×30	ANYmal D		Frontier (Hesai QT-64)	VILENS odometry, VILENS-SLAM	Perceptive locomotion	Integrated
Wytham Woods, UK	WyJ-01	2024-06-25	12×12	ANYmal D		Frontier (Hesai QT-64), Rajant mesh network	VILENS odometry, VILENS-SLAM	Blind locomotion	Integrated
	WyJ-02	2024-06-25	20×15						
	WyJ-03	2024-06-25	20×15						
	WyJ-04	2024-06-25	22×20						
	WyJ-05	2024-06-25	22×15						
Stein am Rhein, Switzerland	SaR-01	2024-07-08	26×22	ANYmal D		Frontier (Hesai QT-64), Rajant mesh network	VILENS odometry, VILENS-SLAM	Perceptive locomotion	Integrated
	SaR-02	2024-07-09	16×20						
	SaR-03	2024-07-09	16×20						
	SaR-04	2024-07-09	16×20						

* This mission was manually interrupted.

TABLE 1. Summary of the campaigns and missions executed between May 2023 and July 2024. We specify the locations, date, and area of the survey defined by the operator. We also report changes in the robot setup (platform, sensing, and software stack). We have **highlighted** each component that changed with respect to the previous campaign. Refer to Sec. IV for further details.

	A	B	C	D	E
1	ID	pos_x [m]	pos_y [m]	DBH [cm]	Reco Height [m]
2	0	6.841	-31.948	28.2	11.61
3	2	8.808	-29.324	87.2	9.10
4	6	25.947	-13.013	34.3	9.63
5	11	33.801	-6.123	43.6	3.17
6	12	33.169	-7.483	57.6	5.42
7	13	30.505	-9.006	66.8	7.36
8	14	19.038	-17.501	54.1	5.25
9	18	12.649	-27.656	95.9	7.36
10	19	16.858	-24.609	48.6	9.25
11	22	20.617	-15.207	67.4	5.95
12	26	12.758	-27.199	72.1	6.01
13	27	16.546	-25.241	85.0	9.87
14	28	20.688	-14.91	79.4	11.15

FIGURE 6. Output forest inventory in spreadsheet format. This example was produced using data from the Forest of Dean campaign. It reports the detected tree positions in a fixed frame, **DBH**, and height of the detected trees.

visually reported to the operator and also exported as a *forest inventory*—a spreadsheet representing the forest as a list of tree locations and associated tree traits (see Fig. 6).

IV. FIELD CAMPAIGNS

A. Campaign description

Between 2023 and 2024, we executed five field campaigns in different forests in the UK, Finland, and Switzerland. The campaigns consisted of 16 individual missions, which are summarized in Tab. 1.

Five missions (Evo-01 to Evo-05) were executed in conifer forests in Evo, Finland in May 2023. Other two were ran in the UK, in mixed (WyO-01, at Wytham Woods, in October 2023) and oak forests (Dea-01, at the Forest of Dean, in February 2024). Three runs during the summer season were executed in a different area of Wytham Woods (WyJ-01 to WyJ-05, June 2024), as well as in a mixed forest in Stein am Rhein, Switzerland (July 2024).

The 2023 campaigns developed the autonomy system and in 2024 we focused on the integration of the online forest inventory system. Fig. 7 illustrates some examples of the robot deployed in the field.

B. Deployment methodology

All the missions followed the same procedure. A *robot operator* was in charge of defining the survey mission and supervising the overall execution. A *safety operator*, carrying a remote controller to overrule the autonomy system, followed the robot at a safe distance during the mission to intervene and avoid undesired situations (e.g., collisions and *cul-de-sacs*). The interventions were the main signal to evaluate the autonomy system.

C. Campaign 1: Evo, May 2023 (Evo)

Our first campaign, in Finland, aimed to evaluate the proposed autonomy system in real forest environments. The ANYmal C robot was equipped with the Velodyne VLP-16



FIGURE 7. Examples of the autonomous deployments across the campaigns, illustrating the differences in the forest and seasonal conditions. (a) Coniferous forest in Evo, Finland. (b) Mixed forest in Wytham Woods, UK. (c) Oak forest in the Forest of Dean, UK. (d) Mixed forest in Wytham Woods. (e) Mixed forest in Stein am Rhein, Switzerland.

(30° FoV) which is not well suited for forest inventory. We used CompSLAM [67] as the main odometry system. CompSLAM does not perform online pose graph optimization and loop-closure detection to ensure consistency.

Most of the areas in the forest were flat plots with no major undergrowth as the campaign was carried out in late autumn (Fig. 7a). This simplified the locomotion and navigation challenges as trees and bushes were evident. Missions EVO-01 and EVO-02 were executed in the same plots, obtaining consistent results. Mission EVO-03 was executed in a nearby boggy area. The robot became trapped in a bog until the mission was eventually terminated (Fig. 13). Missions EVO-04 and EVO-05 were executed in different, larger areas of the forest. The missions lasted between 10–20 min, with the robot walking up to 600 m on the longest missions.

D. Campaign 2: Wytham Woods, October 2023 (WyO)

The second campaign was executed in a mixed forest plot in Wytham Woods, Oxford, UK. The goal was to improve the state estimation solution problems reported in the Evo campaign. We transitioned to the current state estimation solution of using VILENS for odometry, and VILENS-SLAM for LiDAR SLAM to ensure consistency. The testing area was a sloped plot with little undergrowth and a few loose branches, shown in Fig. 7b.

E. Campaign 3: Forest of Dean, February 2024 (Dea)

The third campaign, executed in an oak plantation in the Forest of Dean in the UK, tested the integration of the forest inventory system and the hardware transition to ANYmal D. We also improved the sensor payload by introducing the Frontier multi-sensor unit (Sec. B). This switched from the 30° vertical FoV VLP-16 to the Hesai QT64 LiDAR with a 104° FoV, enabling better coverage of the tree canopy.

This campaign was also executed during winter. Minimal undercanopy and the leaf-fall conditions simplified navigation and reconstruction (Fig. 7c). We chose a flat testing area with few boggy patches and sharp inclines. In this campaign we executed our largest mission, Dea-01, a 125 m × 30 m survey area, which corresponded to a 0.96 ha plot.

F. Campaign 4: Wytham Woods, June 2024 (WyJ)

Our fourth campaign was carried out during the summer in Wytham Woods. We did not test in the same area as WyO-01 because significant undergrowth (about 1 m high) and dense bushes impeded access. We instead tested in another area, close to an access road, where the undergrowth was about 50 cm tall, see Fig. 7d.

For the deployments, we also introduced the Rajant Bread-Crumb mesh network. In previous deployments we found the internal WiFi antenna was not sufficient to communicate with

Mission Name	Autonomy		Interv. [#]	MDBI [m]	MTBI [s]	Forest Inventory	
	Time [s]	Dist. traveled [m]				Area covered [ha]	Trees [#]
Evo-01	575.6	270.3	2	84.8	176.2	0.33	✗
Evo-02	432.0	233.6	0	233.6	432.4	0.32	✗
Evo-03 [†]	816.8	301.1	7	30.5	72.2	0.37	✗
Evo-04	988.4	336.6	7	39.4	114.2	0.37	✗
Evo-05	1275.5	609.7	10	51.1	99.6	0.64	✗
WyO-01	436.9	215.0	2	69.9	136.4	0.29	✗
Dea-01	1283.5	665.4	8	66.0	127.8	0.96	97
WyJ-01	239.5	82.4	4	15.8	45.2	0.16	28
WyJ-02	447.7	163.3	4	28.8	77.0	0.20	46
WyJ-03	418.1	139.4	3	33.9	101.6	0.20	41
WyJ-04	556.6	179.4	5	27.3	80.8	0.23	46
WyJ-05	576.1	196.9	7	22.9	66.2	0.22	52
SaR-01	828.7	286.5	7	30.3	88.8	0.26	66
SaR-02	458.3	149.1	0	149.1	458.3	0.20	43
SaR-03	461.5	152.1	0	152.1	461.5	0.20	44
SaR-04	432.4	140.2	1	78.6	215.3	0.20	36

[†] This mission was manually interrupted.

✗ Forest inventory system was not available.

TABLE 2. Summary of the results for all campaigns. We report the autonomy metrics, and the performance of the forest inventory system.

the robot and to allow the robot operator to monitor mission execution. With a small network of BreadCrumb units (generally four), we could improve robot-to-base connectivity.

In this campaign, we deployed the robot 5 times (WyJ-01–WyJ-05), aiming to map the same environment every time. For this campaign we changed the locomotion controller to the blind stock controller from ANYbotics, which performed better in the dense undergrowth present in this area.

G. Campaign 5: Stein am Rhein, July 2024 (SaR)

Our last campaign was executed in Switzerland in the second week of July. The *Oberwald* forest is located beside the northern border with Germany. Much of the forest was untraversable to the robot due to dense foliage—being difficult even for humans to traverse. We selected a plot on a mild slope that contained less undergrowth.

The robot was deployed four times (SaR-01–SaR-05) in the chosen forest plot, where it covered only a small region of the area (approximately 0.2 ha). The objective of these deployments was (1) to evaluate the consistency of the tree detection system, and (2) to demonstrate the system as part of the wider DigiForest research project.

V. RESULTS

We report the main results obtained with our system for three different aspects: the autonomy system, the forest inventory performance, and potential future applications. The following sections discuss them in detail, presenting the corresponding evaluation metrics, data, and findings of each case. Tab. 2 provides a summary of the main outcomes.

A. Analysis of the Autonomy System

Our first analysis focuses on the level of autonomy achieved by our system. To measure this, we used metrics from the autonomous driving literature [68]. Particularly, we chose the Mean Distance Between Interventions (MDBI) and Mean Time Between Interventions (MTBI), which provide a mechanism to assess the time and distance the robot operated autonomously when compared to the total mission duration and distance traveled.

To compute these metrics, we analyzed the logs of each mission to determine the total time taken on the deployments, as well as the duration of each command given by the safety operator. In principle, each action should be considered an intervention event. However, due to the sparsity of the commands and the variability of the different safety operators to react during the mission, we decided to aggregate all the safety commands within 10 s time windows to define single *intervention* events. This is shown in Fig. 8.

We matched the intervention signals to the geometrical paths recorded on each mission, which we will denote $\mathcal{M}$ in this section. This allowed us to obtain N time and distance segments in which the robot operated autonomously on each mission, denoted by d_{auto} and t_{auto} . Then, we computed the **MDBI** and **MTBI** of each mission as follows:

$$\text{MDBI}_{\mathcal{M}} = \frac{1}{N} \sum_{i \in \mathcal{M}} d_{\text{auto}}^i, \quad (2)$$

$$\text{MTBI}_{\mathcal{M}} = \frac{1}{N} \sum_{i \in \mathcal{M}} t_{\text{auto}}^i. \quad (3)$$

Tab. 2 reports the number of interventions as well as the **MDBI** and **MTBI** for all the campaigns. While these metrics




FIGURE 8. Summary of the autonomy results obtained across the five campaigns. We illustrate the periods where the robots operated fully autonomously (■), and where it was controlled by the safety operator (■). The aggregated safety actions are defined as *intervention* events (■).

have been used to describe the performance of autonomous systems, we observed that the intervention events were usually rare, and short in time. To obtain a better understanding of the interventions events, we additionally aggregated all the intervention events by time and distance to build histograms of their distributions. This is shown in Fig. 9.

In general, we observed that when all the missions are considered, the interventions do not follow a normal distribution but Poisson’s in both time and distance. Most of the interventions were effectively short: the robot was manually driven for less than 20 s, or corrected for less than 5 m. A few sparse events in Evo (Evo-03) and the Forest of Dean (Dea-01) required the safety operator to take control for longer periods, specially in the larger scale missions.

We further revised the Time Between Interventions (TBI) and Distance Between Interventions (DBI) from a distribution point of view, which is shown in Fig. 10. We removed all the missions where no interventions were reported.

We noticed that the distributions are not Gaussian either—this suggests that using the MTBI and MDBI might not be an ideal metric in these applications. In fact, they also loosely follow a Poisson distribution, reflecting our field observation that the safety operator tended to intervene frequently for short periods of time, only in some parts of the missions. The safety operator’s actions mainly aimed to “help the robot” when the local planner was not able to find a path, or when




FIGURE 9. Histogram of the duration (top) and distance traveled (bottom) during the interventions. Most of the interventions reported were short—20 s on average.

the mission planner did not produce significant progress by the proposed re-planning strategy.

As a general conclusion of this analysis, we note that since we tested in different conditions with an evolving solution, it is difficult to claim specific achievements in terms of the autonomy of the system. The state estimation stack changed across campaigns, and the differences in the scenery across seasons also introduced new challenges that were not considered in previous missions. Nevertheless, the system achieved similar performance across campaigns *on average*. Our solution enabled the robot to navigate autonomously above 80 % of the distance traveled, or 90 % of the mission time. The main limitations of our system are discussed in detail in Sec. VI and Sec. VII.

B. Analysis of Forest Inventory System

Our second analysis focuses on the performance of the forest inventory system. We report results for the Forest of Dean, Wytham Woods, and Stein am Rhein campaigns.

Tab. 2 summarizes the general metrics of the missions. To estimate the scanning coverage, we considered an *effective range* of 15 m, which was the maximum LiDAR range to obtain sufficiently dense point clouds for the forest inventory system. Using this value, the estimates of the area covered across the different missions span from 0.16 ha to 0.96 ha. The typical scanning speed was 1.8 ha/h, with the robot walking at an average speed of 0.6 m/s across the different campaigns.




FIGURE 10. Distribution of time (top) and distance (bottom) between interventions for all missions.

In missions with repeated operation in the same plot, such as Wytham Woods (June) and Stein am Rhein, we detected a consistent number of trees, as reported in Tab. 2. While we did not have ground truth TLS measurements for many of the test sites, in our previous work [20] we determined that the forest inventory can achieve average DBH accuracy of 2 cm. This was additionally confirmed in experiments in Stein am Rhein by comparing our estimates against manual measurements with tree calipers. An extract of the forest inventory obtained in mission Dea-01 of the Forest of Dean campaign is shown in Fig. 6.

Fig. 11 shows visualizations of the forest inventory system when deployed in the Forest of Dean, Wytham Woods (2024), as well as Stein am Rhein. The detected trees are visualized by colored point clouds. The estimated terrain model used to segment the ground is also visualized. We reported differences in the tree species and vegetation level due to the different seasons. In the Forest of Dean, the trees look well-defined, as they were all oak trees without leaves. In Stein am Rhein, it was a mixed forest with a majority of large European beeches, which have a well-defined stem that is easily detected by our system. In contrast, for Wytham Woods tree segmentation was more difficult, mainly due to the mixed species and multi-stem trees that broke some of the assumptions of our tree analysis system.

Our results generally demonstrate the feasibility of achieving online forest inventory with the autonomous legged platform. This is subject to the autonomy performance of the platform, as discussed in Sec. A, as well as the sensing

specifications of the MLS unit being carried. These aspects are further discussed in Sec. VI.

C. Relocalization in Prior Maps

In our last experiment we report preliminary experiments regarding relocalization in forests with the legged platform for future continuous monitoring. The initial proof of concept of this idea was executed during the Forest of Dean campaign. We collected an initial map using a Frontier unit mounted on a backpack, which was mounted around 2.5 m above the ground level. A human carried the sensor within the plot to build an initial map, covering an area above 1 ha.

We additionally recorded two sequences with the ANYmal D platform teleoperated. We recorded the same signals as in the other missions reported in the paper, but since the robot is 60 cm height, it had a lower viewpoint with respect to the prior map.

In post-processing, we evaluated a LiDAR-based localization system, which we presented in prior work [69]. Fig. 12 shows an example of the relocalization sequences. The general overview illustrates the scale of the forest plot. The blue path indicates the prior trajectory recorded with the MLS system at 2.5 m height approximately. The yellow path shows the path followed in one of the teleoperated runs, as estimated by our proposed localization system. The differences in the paths' height are further shown in close-ups (a) and (b).

VI. LESSONS LEARNED

In this section, we reflect on the lessons we learned during the five campaigns, highlighting positive observations during the deployments.

A. Lesson 1 – Hardware: Some Platforms are Mature for Natural Deployments

In our experiments we focus on the ANYmal platform for two main reasons. First, its ROS-based development ecosystem allowed us to exploit the current developments in locomotion control, which make falls a rare event during deployment, as well as a straightforward integration with external sensing. Second, the ANYmal is a rugged, IP67 certified platform, which gave us confidence deploying it in natural environments without risking damaging the robot. We tested it in environments with branches, twigs, grass, mud, and undergrowth, without hardware problems.

What we learned: The ANYmal-legged robot platform is mature enough to be reliably deployed outdoors in natural environments. Other commercially available platforms, such as the Unitree B2 or Boston Dynamics Spot, offer comparable capabilities in terms of ruggedness, locomotion capability, and extendability.




FIGURE 11. Illustrative example of the different outputs of our forest inventory pipeline in different autonomous deployments. (a) Forest of Dean (Dea-01), (b) Wytham Woods (Wyt-02), (c) Stein am Rhein (SaR-03). We can observe the differences in the tree segmentations due to the diverse species present in the forests considered.

B. Lesson 2 – State estimation: LiDAR-inertial State Estimation Provided Satisfactory Performance in Forests

Also motivated by the research outcomes of the DARPA SubT challenge, we started with a state estimator that fused multiple odometry sources—legged-inertial, visual-inertial, LiDAR, LiDAR-inertial. The DARPA SubT test environments challenged the different modalities in different ways: darkness affected visual solutions, while degenerate structures such as tunnels caused LiDAR-based solutions to fail. Sensor fusion approaches such as CompSLAM [67] or SuperOdometry [70] dealt with this problem by running different odometry estimators in parallel and choosing the most reliable depending on the environmental conditions.

For the forestry application we focused on, where the environment is geometrically rich, a simpler solution based on LiDAR-inertial odometry was sufficient. We decided not to use leg odometry nor visual-odometry in our state

estimation pipeline, in order to avoid additional challenges introduced by the terrain rigidity or the variable lighting conditions under the canopy.

What we learned: LiDAR was sufficient to achieve accurate and reliable state estimate in forest environments. When combined with a simple pose graph-based SLAM approach, it produced a globally consistent estimate for the autonomy system.

C. Lesson 3 – Autonomy: A Modular Formulation Simplified Development and Integration

The approach we adopted for the autonomy system was composed of different levels—from locomotion up to operator-guided mission planning. The modules simplified the development and allowed us to reuse components, such as the locomotion and local planning algorithms.




FIGURE 12. Relocalization example in the Forest of Dean. The blue path ■ represents the prior trajectory recorded with a **MLS** system carried on a backpack at 2.5 m height. The yellow path ■ is executed by the robot in a subsequent localization run. (a) and (b) are close-ups to two segments of the sequence, illustrating the differences between the reference and executed path.

The mission planning strategy based on a reference lawn mower pattern simplified the design of the high-level modules. Then, most of the efforts were delegated to the lower-levels responsible for detecting traversable areas and reacting accordingly. Furthermore, the forest inventory system did not require an active data collection strategy but instead adapted as the SLAM pose graph improved.

What we learned: The modular autonomy system was simple to develop and practical for debugging, making it a good starting point for our deployments. The modularity also revealed that we need further improvements in certain modules (e.g., locomotion, traversability), as well as a better feedback mechanism between them to improve the autonomy performance.

D. Lesson 4 – Forestry Use-case: Competitive Inventory Performance can be Achieved from Cheaper Sensors

We were interested in testing how robotics could contribute to fundamental challenges in forestry, such as forest inventory, and potentially offer a path to scale it up. This required a re-think about how such inventories are obtained, and if it was possible to achieve comparable results with the sensors that are commonly used in robotics.

After the Evo campaign, we evaluated the performance of the onboard LiDAR sensors on the ANYmal robot for performing **DBH** estimation. However, the narrow field of view of the sensor produced sub-standard results. This motivated us to find alternatives to increase the portion of the tree the robot could observe. In fact, using a wide-angle LiDAR, produced overall better results with our online forest

inventory pipeline [20]. The same LiDAR was additionally used for the state estimation pipeline, making it an integrated solution.

What we learned: Our preliminary results demonstrate that we can obtain good **DBH** estimates from cheaper sensors, when compared to **TLS**. By collecting the data online and providing direct feedback to the forestry users, we can potentially improve the data acquisition process during the mission.

E. Lesson 5 – Assessment: Testing in Different Environments Provides a Holistic View of the System

We recognize that it is challenging to fairly compare the conditions and results of the campaigns—different countries, seasons, forests, and plots within them, as well as the inherent variability expected when deploying complex robotic systems in autonomous trials. However, they allowed us to obtain a broad overview of the performance of legged systems across different natural domains. By changing the environment for each trial, we could identify new challenges for the different components of the system.

We observed that testing our system was generally easier in winter, where deciduous trees were leafless and undergrowth had died off. This was an ideal starting point for our deployments, because the challenges focused on the system and the localization capabilities. As we moved towards spring and summer with a more mature system, we encountered challenges to our locomotion system from the dense vegetation.

What we learned: Continuous deployment is essential. Starting with simpler environments enabled us to develop an initial working system, which was iterated over as we increased the difficulty of the testing conditions.

VII. CHALLENGES AHEAD

The lessons presented our main outcomes and findings, highlighting the main positive aspects of our campaigns. Next we discuss the challenges and avenues for future work.

A. Challenge 1 – Hardware: The Current Sensing Modalities might be Insufficient for Resilient Autonomy

In our campaigns we mainly relied on LiDAR and inertial sensing for odometry and mapping. However, we did not fully exploit modalities such as vision, haptics or touch. Visual sensing can provide further cues and insights about the environment in situations that cannot be identified using LiDAR alone. Particularly, during summer, the high grass hides objects such as bushes and fallen trees, which could be more easily detected using vision. However, we believe that High Dynamic Range (HDR) imaging is required to equalize images in challenging under-canopy lighting conditions.

Locomotion and navigation are aspects that can improve by additional touch sensing. While the current RL-based controllers provide mechanisms to sense external forces, we did not exploit this information. We reported two cases where such signals could be of benefit, shown in Fig. 13: the robot getting stuck in a bog-hole and also getting entangled with undergrowth. Being able to detect differences in the stiffness of the terrain [71], or new sensors and design of legs or feet [72] could provide new avenues to face these challenges.

B. Challenge 2 – State estimation: Multi-modal Representations are Required

In Lesson 2 we did not report significant challenges for odometry estimation and SLAM. However the interventions reported suggest there are limitations in the representations used in our system for navigation. A straightforward next step for our system is to introduce vision as part of the local map representations, in order to introduce semantic information as discussed in Challenge 1. While we did not integrate our prior work on vision-based traversability estimation [73] to keep the complexity of the system bounded, it is a logical next step for future work. This aligns with current efforts in off-road navigation [74]–[76]. However, we also envision that proprioceptive information must be incorporated, as explored in recent works [77].

Multi-modal representations exploiting vision and LiDAR can also provide additional benefits for downstream tasks such as tree analysis and trait estimation. This is further discussed in Challenge 4.

(a) Evo: Robot stuck in a bog



(b) Stein am Rhein: Robot tripped in undergrowth



FIGURE 13. Challenges Ahead: Some examples of failure cases we observed in Evo and Stein am Rhein. (a) The robot was trapped in bog, where it sank to the body-level and was unable to continue its autonomous mission (Evo-03). (b) The robot fell after tripping due to dense undergrowth. This occurred while being manually operated in an area where we were unable to deploy the robot autonomously.

C. Challenge 3 – Autonomy: Closing the Gap Between the Components can Improve Performance

In Lesson 3, we noted that the hierarchical architecture presented several advantages from an integration and development perspective. However, we also observed that the modularization potentially limited the performance of the autonomy system.

We believe that a redesign of the autonomy system is required to better integrate the different levels of information encoded at each module. For instance, neither the local nor mission planner were driven by the capabilities of the locomotion system. We believe that to achieve further improvements, the local plan should be conditioned on the locomotion capabilities and proprioceptive sensing, as explored in recent work [78].

Further, we could exploit more information from the task execution. This would be useful not in exploration settings but when revisiting an existing prior map. In those cases

our mission planner could be populated with details about ground-level structure, semantics, and traversability. This would open new questions such as long term planning and operation in the site of interest while considering coverage, energy usage, and time constraints.

D. Challenge 4 – Forestry Use-case: Estimating More Trait Parameters is a Next Step

The forest inventory system we introduced in Sec. III, mainly relied on LiDAR data. Therefore, we focused on geometric traits, such as DBH and tree height, which could be estimated from LiDAR scans aggregated in time and space. However, real world applications in forestry or environmental science require different advanced sensors and algorithms.

Species, health, bends, and defects, among others, are fundamental traits for foresters to assess each tree, which cannot be estimated from purely geometric information. As discussed in Challenges 1 and 2, additional sensing modalities must be integrated on the robot to make progress in this direction. These modalities can be loosely coupled (e.g., vision can be used to detect trees and appearance traits, and later associated to geometric traits such as DBH) [79]–[81], or tightly integrated into multi-modal representations that enable its estimation in a holistic manner.

Furthermore, the trees we considered were mostly single-stem trees, which can easily break some of the assumptions taken for its detection, as qualitatively shown in the WyJ experiments (Fig. 11). Mixed forests with deciduous trees, in the summer season, are challenging enough due to dense vegetation. Rainforests are even more difficult cases, with denser structures and challenging lighting conditions. These environments present many opportunities for developing more advanced algorithms for online trait estimation during data collection.

E. Challenge 5 – Assessment: Systematic Evaluation is Required to Measure Progress

We finally observe that further studies are needed to assess performance and improve the proposed solution in a systematic manner. We argue that while our current study sampled various relevant environments and failure cases, we need to reassess the deployment methodology. The work of Bajracharya et al. [82] is a good example to follow.

The four previous challenges presented some directions to improve in. This requires the definition of a few, easily accessible testing environments for continuous deployment, which need to be coupled with a well-defined set of metrics. For instance, the challenges we reported for the autonomy system can be improved by designing simpler experiments in more systematic settings, such as *navigating from A to B* in the same forest plot across different seasons.

Lastly, we argued that small legged platforms can present advantages to other wheeled robots used for forest inventory tasks. While we demonstrated initial steps using legged robots for this purpose, their advantages should be proven

experimentally, quantitatively comparing aspects such as soil damage or the time required to survey the same area. This requires further work in collaboration with foresters and environmental scientists.

VIII. CONCLUSION

In this work, we presented the development and field deployments of an autonomous forest inventory system using a quadruped robot. Through five campaigns in forests in Finland, the UK, and Switzerland across different seasons, we assessed the feasibility of using legged robot for this task.

The system we proposed achieved initial steps toward this goal, autonomously building forest inventories built autonomously up to 1 ha in under 30 min. We analyzed the performance of the system in terms of the autonomy achieved, as well as its potential for forestry and environmental monitoring tasks.

The main outcomes of our deployments are captured in five lessons and challenges. The lessons highlight the maturity of our legged platform, state estimation and forest inventory systems, as well as the advantages of the system modularity and continuous deployment.

The challenges, on the other hand, identify the main limitations and issues we observed with our current system, namely: multi-modal sensing and state estimation, navigation systems to robustly operate in natural environments, and the need for systematic evaluation. We hope this work can foster future research into mobile robot forest inventory.

IX. ACKNOWLEDGMENTS

We acknowledge PreFor Oy for organizing the Evo campaign, Forest Research UK for arranging the Forest of Dean campaign, and WSL for their preparations for the Stein am Rhein campaign. We thank Daniel Marques for his support with the Forest of Dean and Wytham Woods experiments; Wayne Tubby, Matthew Graham, and Tobit Flatscher for hardware integration support; Ethan Tao, Uljad Berdica, and Frank Fu for field footage; Christina Kassab for feedback on the manuscript.

REFERENCES

- [1] C. L. Brack, "Forest Measurement and Modeling - Measuring trees, stands and forests for effective forest management," Australian National University, Tech. Rep., 2001. [Online]. Available: <http://fennerschool-associated.anu.edu.au/mensuration/index.htm>
- [2] K. Calders, J. Adams, J. Armston, H. Bartholomeus, S. Bauwens, L. P. Bentley, J. Chave, F. M. Danson, M. Demol, M. Disney, R. Gaulton, S. M. Krishna Moorthy, S. R. Levick, N. Saarinen, C. Schaaf, A. Stovall, L. Terryn, P. Wilkes, and H. Verbeeck, "Terrestrial laser scanning in forest ecology: Expanding the horizon," *Remote Sensing of Environment*, vol. 251, p. 112102, 2020.
- [3] R. J. Hall, "The Roles of Aerial Photographs in Forestry Remote Sensing Image Analysis," in *Remote Sensing of Forest Environments: Concepts and Case Studies*, M. A. Wulder and S. E. Franklin, Eds. Boston, MA: Springer US, 2003, pp. 47–75.
- [4] X. Liang, V. Kankare, J. Hyypä, Y. Wang, A. Kukko, H. Haggren, X. Yu, H. Kaartinen, A. Jaakkola, F. Guan, M. Holopainen, and M. Vastaranta, "Terrestrial laser scanning in forest inventories," *ISPRS J. Photogramm. Remote Sens.*, vol. 115, pp. 63–77, 2016.

- [5] L. Duncanson, J. Armston, M. Disney, V. Avitabile, N. Barbier, K. Calders, S. Carter, J. Chave, M. Herold, N. Macbean, *et al.*, "Global aboveground biomass product validation best practices protocol. version 1.0," *Best practice protocol for satellite derived land product validation. Land Product Validation Subgroup (Working Group on Calibration and Validation, Committee on Earth Observation Satellites)*, 2020.
- [6] G. J. Hamilton, *Forest mensuration handbook*. Forestry Commission, 1988.
- [7] R. O. Dubayah and J. B. Drake, "Lidar Remote Sensing for Forestry," *Journal of Forestry*, vol. 98, no. 6, pp. 44–46, June 2000.
- [8] M. Miettinen, M. Ohman, A. Visala, and P. Forsman, "Simultaneous localization and mapping for forest harvesters," in *IEEE Int. Conf. Robot. Autom.*, 2007, pp. 517–522.
- [9] J.-F. Tremblay, M. Béland, R. Gagnon, F. Pomerleau, and P. Giguère, "Automatic three-dimensional mapping for tree diameter measurements in inventory operations," *J. Field Robot.*, vol. 37, no. 8, pp. 1328–1346, 2020.
- [10] S. W. Chen, G. V. Nardari, E. S. Lee, C. Qu, X. Liu, R. A. F. Romero, and V. Kumar, "Sloam: Semantic lidar odometry and mapping for forest inventory," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 612–619, 2020.
- [11] X. Liu, G. V. Nardari, F. C. Ojeda, Y. Tao, A. Zhou, T. Donnelly, C. Qu, S. W. Chen, R. A. F. Romero, C. J. Taylor, and V. Kumar, "Large-scale autonomous flight with real-time semantic slam under dense forest canopy," *IEEE Robot. Autom. Lett.*, vol. 7, no. 2, pp. 5512–5519, 2022.
- [12] M. Pierzchała, P. Giguère, and R. Astrup, "Mapping forests using an unmanned ground vehicle with 3D LiDAR and graph-SLAM," *Comput. Electron. Agr.*, vol. 145, pp. 217–225, Feb. 2018.
- [13] D. J. Todd, *Applications of walking machines*. Boston, MA: Springer US, 1985, pp. 169–177.
- [14] C. D. Bellicoso, M. Bjelonic, L. Wellhausen, K. Holtmann, F. Günther, M. Tranzatto, P. Fankhauser, and M. Hutter, "Advances in real-world applications for legged robots," *J. Field Robot.*, vol. 35, no. 8, pp. 1311–1326, 2018.
- [15] C. Quail, E. Emonot-de Carolis, and F. Auat Cheein, "Legged robots in the agricultural context: Analysing their traverse capabilities and performance," in *Annual Conference of the IEEE Industrial Electronics Society*, 2023, pp. 01–07.
- [16] F. Qian, D. Jerolmack, N. Lancaster, G. Nikolich, P. Reverdy, S. Roberts, T. Shipley, R. S. Van Pelt, T. M. Zobeck, and D. E. Koditschek, "Ground robotic measurement of aeolian processes," *Aeolian Research*, vol. 27, pp. 1–11, 2017.
- [17] M. Mattamala, N. Chebrolu, B. Casseau, L. Freißmuth, J. Frey, T. Tuna, M. Hutter, and M. Fallon, "Autonomous forest inventory with legged robots: System design and field deployment," in *ICRA Workshop on Field Robotics*, 2024.
- [18] J. Tang, Y. Chen, A. Kukko, H. Kaartinen, A. Jaakkola, E. Khoramshahi, T. Hakala, J. Hyypä, M. Holopainen, and H. Hyypä, "SLAM-Aided Stem Mapping for Forest Inventory with Small-Footprint Mobile LiDAR," *Forests*, vol. 6, no. 12, pp. 4588–4606, Dec. 2015.
- [19] A. Proudman, M. Ramezani, S. T. Digumarti, N. Chebrolu, and M. Fallon, "Towards real-time forest inventory using handheld lidar," *Robotics and Autonomous Systems*, vol. 157, p. 104240, 2022.
- [20] L. Freißmuth, M. Mattamala, N. Chebrolu, S. Schaefer, S. Leutenegger, and M. Fallon, "Online Tree Reconstruction and Forest Inventory on a Mobile Robotic System," in *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2024.
- [21] K. Anderson and K. J. Gaston, "Lightweight unmanned aerial vehicles will revolutionize spatial ecology," *Frontiers in Ecology and the Environment*, vol. 11, no. 3, pp. 138–146, 2013.
- [22] L. P. Koh and S. A. Wich, "Dawn of Drone Ecology: Low-Cost Autonomous Aerial Vehicles for Conservation," *Tropical Conservation Science*, vol. 5, no. 2, pp. 121–132, June 2012.
- [23] J. R. Kellner, J. Armston, M. Birrer, K. C. Cushman, L. Duncanson, C. Eck, C. Fallegger, B. Imbach, K. Král, M. Krůček, J. Trochta, T. Vrška, and C. Zraggen, "New Opportunities for Forest Remote Sensing Through Ultra-High-Density Drone Lidar," *Surveys in Geophysics*, vol. 40, no. 4, pp. 959–977, July 2019.
- [24] D. R. A. Almeida, A. M. Almeida Zambrano, E. N. Broadbent, A. L. Wendt, P. Foster, B. E. Wilkinson, C. Salk, D. D. A. Papa, S. C. Stark, R. Valbuena, E. B. Gorgens, C. A. Silva, P. H. S. Brancalion, M. Fagan, P. Meli, and R. Chazdon, "Detecting successional changes in tropical forest structure using GatorEye drone-borne lidar," *Biotropica*, vol. 52, no. 6, pp. 1155–1167, Nov. 2020.
- [25] K. Shah, G. Ballard, A. Schmidt, and M. Schwager, "Multidrone aerial surveys of penguin colonies in antarctica," *Sci. Robot.*, vol. 5, no. 47, p. eabc3000, 2020.
- [26] D. C. Schedl, I. Kurmi, and O. Bimber, "An autonomous drone for search and rescue in forests using airborne optical sectioning," *Sci. Robot.*, vol. 6, no. 55, p. eabg1188, 2021.
- [27] Y. Lin, J. Hyypä, and A. Jaakkola, "Mini-UAV-Borne LIDAR for Fine-Scale Mapping," *IEEE Geosci. Remote Sens. Lett.*, vol. 8, no. 3, pp. 426–430, May 2011.
- [28] E. Hyypä, J. Hyypä, T. Hakala, A. Kukko, M. A. Wulder, J. C. White, J. Pyörälä, X. Yu, Y. Wang, J.-P. Virtanen, O. Pohjavirta, X. Liang, M. Holopainen, and H. Kaartinen, "Under-canopy UAV laser scanning for accurate forest field measurements," *ISPRS J. Photogramm. Remote Sens.*, vol. 164, pp. 41–60, June 2020.
- [29] E. Aucone, S. Kirchgeorg, A. Valentini, L. Pellissier, K. Deiner, and S. Mintchev, "Drone-assisted collection of environmental DNA from tree branches for biodiversity monitoring," *Sci. Robot.*, vol. 8, no. 74, p. eadd5762, 2023.
- [30] X. Zhou, X. Wen, Z. Wang, Y. Gao, H. Li, Q. Wang, T. Yang, H. Lu, Y. Cao, C. Xu, and F. Gao, "Swarm of micro flying robots in the wild," *Sci. Robot.*, vol. 7, no. 66, p. eabm5954, 2022.
- [31] G. Notomista, Y. Emam, and M. Egerstedt, "The slothbot: A novel design for a wire-traversing robot," *IEEE Robot. Autom. Lett.*, vol. 4, no. 2, pp. 1993–1998, 2019.
- [32] S. Kirchgeorg, E. Aucone, F. Wenk, and S. Mintchev, "Design, modeling, and control of avocado: A multimodal aerial-tethered robot for tree canopy exploration," *IEEE Trans. Robot.*, vol. 40, pp. 592–605, 2024.
- [33] A. Tominaga, H. Eiji, and A. Mowshowitz, "Development of Navigation System in Field Robot for Forest Management," in *Int. Conf. Soft Comput. Intell. Syst. and Int. Sym. Adv. Intell. Syst.*, Toyama, Japan, Dec. 2018, pp. 1142–1147.
- [34] J. Rossmann, M. Schluse, C. Schlette, A. Buecken, P. Krahwinkel, and M. Emde, "Realization of a highly accurate mobile robot system for multi purpose precision forestry applications," in *2009 International Conference on Advanced Robotics*, June 2009, pp. 1–6.
- [35] E. Jelavic, D. Jud, P. Egli, and M. Hutter, "Robotic precision harvesting: Mapping, localization, planning and control for a legged tree harvester," *Field Robotics*, 2021.
- [36] E. Jelavic, T. Kapgen, S. Kerscher, D. Jud, and M. Hutter, "Harveri : A Small (Semi-)Autonomous Precision Tree Harvester," in *Innovation in Forestry Robotics: Research and Industry Adoption (ICRA 2022)*, May 2022.
- [37] S. Fountas, T. A. Gemtos, and S. Blackmore, *Robotics and Sustainability in Soil Engineering*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 69–80.
- [38] A. Calleja-Huerta, M. Lamandé, O. Green, and L. Munkholm, "Vertical and horizontal stresses from a lightweight autonomous field robot during repeated wheeling," *Soil and Tillage Research*, vol. 233, p. 105790, Sept. 2023.
- [39] European Commission, "Commission communication: New EU forest strategy for 2030," 2021.
- [40] M. Ramezani, M. Brandao, B. Casseau, I. Havoutis, and M. Fallon, "Legged robots for autonomous inspection and monitoring of offshore assets," in *Offshore Technology Conference*. OTC, 2020.
- [41] G. Xin, C. Tiseo, W. Wolfslag, J. Smith, O. Cebe, Z. Li, S. Vijayakumar, and M. N. Mistry, "Variable autonomy of whole-body control for inspection and intervention in industrial environments using legged robots," in *IEEE Intl. Conf. Autom. Sci. Eng.*, 2020, pp. 1415–1420.
- [42] M. Staniaszek, T. Flatscher, J. Rowell, H. Niu, W. Liu, Y. You, M. Gadd, M. Mattamala, D. D. Martini, L. Pitt, A. Schutz, R. Skilton, M. Fallon, and N. Hawes, "Autoinspect: Towards long-term autonomous industrial inspection," *IEEE Trans. Fiel. Robot.*, 2024, submitted.
- [43] T. H. Chung, V. Orekhov, and A. Maio, "Into the Robotic Depths: Analysis and Insights from the DARPA Subterranean Challenge," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 6, no. 1, 2023.
- [44] N. Hudson, F. Talbot, M. Cox, J. Williams, T. Hines, A. Pitt, B. Wood, D. Frousheger, K. L. Surdo, T. Molnar, R. Steindl, M. Wildie, I. Sa, N. Kottege, K. Stepanas, E. Hernandez, G. Catt, W. Docherty,

- B. Tidd, B. Tam, S. Murrell, M. Bessell, L. Hanson, L. Tychsen-Smith, H. Suzuki, L. Overs, F. Kendoul, G. Wagner, D. Palmer, P. Milani, M. O'Brien, S. Jiang, S. Chen, and R. Arkin, "Heterogeneous Ground and Air Platforms, Homogeneous Sensing: Team CSIRO Data61's Approach to the DARPA Subterranean Challenge," *Field Robotics*, vol. 2, no. 1, pp. 595–636, 2022.
- [45] A. Agha, K. Otsu, B. Morrell, D. D. Fan, R. Thakker, A. Santamaria-Navarro, S. Kim, A. Bouman, X. Lei, J. A. Edlund, M. F. Ginting, K. Ebadi, M. Anderson, T. Pailevanian, E. Terry, M. T. Wolf, A. Tagliabue, T. S. Vaquero, M. Palieri, S. Tepsuporn, Y. Chang, A. Kalantari, F. Chavez, B. T. Lopez, N. Funabiki, G. Miles, T. Touma, A. Buscicchio, J. Tordesillas, N. Alatur, J. Nash, W. Walsh, S. Jung, H. Lee, C. Kanellakis, J. Mayo, S. Harper, M. Kaufmann, A. Dixit, G. Correa, C. Lee, J. Gao, G. Merewether, J. Maldonado-Contreras, G. Salhotra, M. S. da Silva, B. Ramtoulou, Y. Kubo, S. A. Fakoorian, A. Hatteland, T. Kim, T. Bartlett, A. Stephens, L. Kim, C. Bergh, E. Heiden, T. Lew, A. Cauligi, T. Heywood, A. Kramer, H. A. Leopold, H. C. Choi, S. Daftry, O. Toupet, I. Wee, A. Thakur, M. Feras, G. Beltrame, G. Nikolakopoulos, D. H. Shim, L. Carlone, and J. Burdick, "Nebula: Quest for robotic autonomy in challenging environments; TEAM costar at the DARPA subterranean challenge," *CoRR*, vol. abs/2103.11470, 2021.
- [46] M. Tranzatto, T. Miki, M. Dharmadhikari, L. Bernreiter, M. Kulkarni, F. Mascarich, O. Andersson, S. Khattak, M. Hutter, R. Siegwart, and K. Alexis, "CERBERUS in the DARPA subterranean challenge," *Sci. Robot.*, vol. 7, no. 66, 2022.
- [47] L. Milburn, J. Gamba, M. Fernandes, and C. Semini, "Computer-vision based real time waypoint generation for autonomous vineyard navigation with quadruped robots," in *IEEE Intl. Conf. Auton. Robot Syst. Compet.*, 2023, pp. 239–244.
- [48] L. Amatucci, G. Turrissi, A. Bratta, V. Barasuol, and C. Semini, "VERO: A vacuum-cleaner-equipped quadruped robot for efficient litter removal," *J. Field Robot.*, vol. 41, no. 6, pp. 1829–1842, 2024.
- [49] D. Wooden, M. Malchano, K. Blankespoor, A. Howard, A. A. Rizzi, and M. H. Raibert, "Autonomous navigation for bigdog," in *IEEE Int. Conf. Robot. Autom.*, 2010, pp. 4736–4741.
- [50] D. M. Bradley, J. K. Chang, D. Silver, M. Powers, H. Herman, P. Rander, and A. Stentz, "Scene understanding for a high-mobility walking robot," in *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2015, pp. 1144–1151.
- [51] C. G. Wilson, F. Qian, D. J. Jerolmack, S. Roberts, J. Ham, D. Koditschek, and T. F. Shipley, "Spatially and temporally distributed data foraging decisions in disciplinary field science," *Cognitive Research: Principles and Implications*, vol. 6, no. 1, p. 29, 2021.
- [52] K. Melo, T. Horvat, and A. J. Ijspeert, "Animal robots in the african wilderness: Lessons learned and outlook for field robotics," *Sci. Robot.*, vol. 8, no. 85, p. eadd8662, 2023.
- [53] C. Canteloup, J. Lee, S. Zimmermann, M. Alvino, M. Montenegro, M. Hutter, and E. van de Waal, "When monkeys meet an anymal robot in the wild," *bioRxiv*, 2024.
- [54] P. Arm, G. Waibel, J. Preisig, T. Tuna, R. Zhou, V. Bickel, G. Ligeza, T. Miki, F. Kehl, H. Kolvenbach, and M. Hutter, "Scientific exploration of challenging planetary analog environments with a team of legged robots," *Sci. Robot.*, vol. 8, no. 80, p. eade9548, 2023.
- [55] M. Segura, D. Ray, and C. Maroto, "Decision support systems for forest management: A comparative analysis and assessment," *Comput. Electron. Agr.*, vol. 101, pp. 55–67, Feb. 2014.
- [56] D. Wisth, M. Camurri, and M. Fallon, "Vilens: Visual, inertial, lidar, and leg odometry for all-terrain legged robots," *IEEE Trans. Robot.*, vol. 39, no. 1, pp. 309–326, 2023.
- [57] P. J. Besl and N. D. McKay, "A method for registration of 3-d shapes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 2, pp. 239–256, 1992.
- [58] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza, "On-manifold preintegration for real-time visual-inertial odometry," *IEEE Trans. Robot.*, vol. 33, no. 1, pp. 1–21, 2017.
- [59] M. Ramezani, G. Tinchev, E. Iuganov, and M. Fallon, "Online LiDAR-SLAM for Legged Robots with Robust Registration and Deep-Learned Loop Closure," in *IEEE Int. Conf. Robot. Autom.*, 2020, pp. 4158–4164.
- [60] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. J. Leonard, and F. Dellaert, "iSAM2: Incremental smoothing and mapping using the Bayes tree," *Int. J. Robot. Res.*, vol. 31, no. 2, pp. 216–235, 2012.
- [61] T. Miki, L. Wellhausen, R. Grandia, F. Jenelten, T. Homberger, and M. Hutter, "Elevation mapping for locomotion and navigation using gpu," in *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2022, pp. 2273–2280.
- [62] H. Choset, "Coverage of known spaces: The boustrophedon cellular decomposition," *Auton. Robots*, vol. 9, no. 3, pp. 247–253, 2000.
- [63] M. Mattamala, N. Chebrolu, and M. F. Fallon, "An efficient locally reactive controller for safe navigation in visual teach and repeat missions," *IEEE Robot. Autom. Lett.*, vol. 7, no. 2, pp. 2353–2360, 2022.
- [64] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Sci. Robot.*, vol. 7, no. 62, 2022.
- [65] W. Zhang, J. Qi, P. Wan, H. Wang, D. Xie, X. Wang, and G. Yan, "An easy-to-use airborne lidar data filtering method based on cloth simulation," *Remote. Sens.*, vol. 8, no. 6, p. 501, 2016.
- [66] C. Cabo, C. Ordóñez, C. A. López-Sánchez, and J. Armesto, "Automatic dendrometry: Tree detection, tree height and diameter estimation using terrestrial laser scanning," *Int. J. Appl. Earth Obs. Geoinformation*, vol. 69, pp. 164–174, 2018.
- [67] S. Khattak, H. Nguyen, F. Mascarich, T. Dang, and K. Alexis, "Complementary multi-modal sensor fusion for resilient robot pose estimation in subterranean environments," in *Int. Conf. on Unmanned Aircraft Systems (ICUAS)*, 2020, pp. 1024–1029.
- [68] D. Paz, P. Lai, N. Chan, Y. Jiang, and H. I. Christensen, "Autonomous vehicle benchmarking using unbiased metrics," in *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2020, pp. 6223–6228.
- [69] H. Oh, N. Chebrolu, M. Mattamala, L. Freißmuth, and M. Fallon, "Evaluation and deployment of lidar-based place recognition in dense forests," in *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, vol. abs/2403.14326, Oct. 2024.
- [70] S. Zhao, H. Zhang, P. Wang, L. Nogueira, and S. Scherer, "Super odometry: Imu-centric lidar-visual-inertial estimator for challenging environments," in *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2021, pp. 8729–8736.
- [71] G. B. Margolis, X. Fu, Y. Ji, and P. Agrawal, "Learning to see physical properties with active sensing motor policies," in *Conf. on Robot Learning*, vol. 229, 2023, pp. 2537–2548.
- [72] S. Godon, A. Ristolainen, and M. Kruusmaa, "Robotic feet modeled after ungulates improve locomotion on soft wet grounds," *Bioinspiration Biomimetics*, vol. 19, no. 6, p. 066009, nov 2024.
- [73] J. Frey, M. Mattamala, N. Chebrolu, C. Cadena, M. Fallon, and M. Hutter, "Fast Traversability Estimation for Wild Visual Navigation," in *Robot.: Sci. Syst.*, 2023.
- [74] X. Meng, N. Hatch, A. Lambert, A. Li, N. Wagener, M. Schmittle, J. Lee, W. Yuan, Z. Q. Chen, S. Deng, G. Okopal, D. Fox, B. Boots, and A. Shaban, "Terrainet: Visual modeling of complex terrain for high-speed, off-road navigation," in *Robot.: Sci. Syst.*, 2023.
- [75] J. Frey, S. Khattak, M. Patel, D. Atha, J. Nubert, C. Padgett, M. Hutter, and P. Spieler, "Roadrunner - learning traversability estimation for autonomous off-road driving," *IEEE Trans. Fiel. Robot.*, vol. abs/2402.19341, 2024.
- [76] X. Cai, S. Ancha, L. Sharma, P. R. Osteen, B. Bucher, S. Phillips, J. Wang, M. Everett, N. Roy, and J. P. How, "EVORA: deep evidential traversability learning for risk-aware off-road autonomy," *IEEE Trans. Robot.*, vol. 40, pp. 3756–3777, 2024.
- [77] J. Chen, J. Frey, R. Zhou, T. Miki, G. Martius, and M. Hutter, "Identifying terrain physical parameters from vision - towards physical-parameter-aware locomotion and navigation," *IEEE Robot. Autom. Lett.*, vol. 9, no. 11, pp. 9279–9286, 2024.
- [78] C. Zhang, J. Jin, J. Frey, N. Rudin, M. Mattamala, C. Cadena, and M. Hutter, "Resilient legged local navigation: Learning to traverse with compromised perception end-to-end," in *IEEE Int. Conf. Robot. Autom.*, 2024.
- [79] M. Carpentier, P. Giguère, and J. Gaudreault, "Tree species identification from bark images using convolutional neural networks," in *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2018, pp. 1075–1081.
- [80] H. Woo, I. Kim, and B. Choi, "Computer vision techniques in forest inventory assessment: Improving accuracy of tree diameter measurement using smartphone camera and photogrammetry," *Sensors & Materials*, vol. 33, 2021.
- [81] V. Grondin, J.-M. Fortin, F. Pomerleau, and P. Giguère, "Tree detection and diameter estimation based on deep learning," *Forestry: An International Journal of Forest Research*, vol. 96, no. 2, pp. 264–276, Oct. 2022.

- [82] M. Bajracharya, J. Borders, R. Cheng, D. M. Helmick, L. Kaul, D. Kruse, J. Leichty, J. Ma, C. Matl, F. Michel, C. Papazov, J. Petersen, K. Shankar, and M. Tjersland, "Demonstrating mobile manipulation in the wild: A metrics-driven approach," in *Robot.: Sci. Syst.*, 2023.



Matias Mattamala (Member, IEEE) is a Postdoctoral Researcher in the Dynamic Robot Systems Group at the University of Oxford. He received his M.Sc. in Electrical Engineering from the Universidad de Chile in 2018, and his Ph.D. in vision-based legged robot navigation from the University of Oxford in 2023. His research interests are in the foundations and systems for robot autonomy—representations, perception, action, and learning—with applications to field robotics.



Nived Chebrolu (Member, IEEE) is a Postdoctoral Research Associate at the Oxford Robotics Institute, University of Oxford, UK. He received his M.Sc. in Robotics from Ecole Centrale de Nantes (ECN), France, and the University of Genoa, Italy in 2015, and his Ph.D. from the University of Bonn in 2021, where he developed registration techniques for agricultural robotic applications. His research interests are in developing robust localization and mapping techniques for field robotics applications.



Jonas Frey (Student Member, IEEE) is a Ph.D. student in the Robotic Systems Lab at ETH Zurich. He received his M.Sc. in Robotics, Systems & Control in 2021 from ETH Zurich. He is also affiliated with the Max Planck Institute through the MPI ETH Center for Learning Systems. His research interests lie in the field of perception, navigation and locomotion, and how it can be used for the deployment of mobile robotic systems.



Leonard Freißmuth (Student Member, IEEE) is a M.Sc. student at the Technical University of Munich, Germany. He was a visiting researcher at the University of Oxford in 2023-2024, working on the DigiForest EU project.



Haedam Oh (Student Member, IEEE) is a Ph.D. student in the Dynamic Robot Systems Group at the University of Oxford. He received his MEng from the University of Oxford in 2024. His research interests include localization and mapping with LiDAR and vision.



Benoit Casseau (Member, IEEE) is a Robotics Software Engineer at the Oxford Robotics Institute, Oxford, UK. He received his M.Sc. in computer science and applied mathematics in 2013 from École Nationale Supérieure d'Informatique et de Mathématiques Appliquées, Grenoble, France.



Marco Hutter (Member, IEEE) is an Associate Professor for Robotic Systems and Director of the Center for Robotics at ETH Zurich. He received his M.Sc. and Ph.D. from ETH Zurich in 2009 and 2013 in the field of design, actuation, and control of legged robots. His research interests are in the development of novel machines and actuation concepts together with the underlying control, planning, and machine learning algorithms for locomotion and manipulation. Marco is the recipient of an ERC Starting Grant, PI of the NCCRs robotics, automation, and digital fabrication, winner of the DARPA SubT Challenge, and a co-founder of several ETH Startups such as ANYbotics and Gravis Robotics. He is also the Director of the Boston Dynamics AI Institute Zurich office.



Maurice Fallon (Senior Member, IEEE) received the B.Eng. degree in electrical engineering from University College Dublin, Dublin, Ireland, in 2004 and the Ph.D. degree in acoustic source tracking from the University of Cambridge, Cambridge, U.K., in 2008. From 2008 to 2012, he was a Postdoc and a Research Scientist with MIT Marine Robotics Group working on SLAM. Later, he was the Perception Lead of MIT's team in the DARPA Robotics Challenge. Since 2017, he has been a Royal Society University Research Fellow and an Associate Professor with the University of Oxford, Oxford, U.K. He leads the Dynamic Robot Systems Group, Oxford Robotics Institute. His research interests include probabilistic methods for localization, mapping, multisensor fusion, and robot navigation.