



Digital Analytics and Robotics for Sustainable Forestry

CL4-2021-DIGITAL-EMERGING-01

Grant agreement no: 101070405

DELIVERABLE 6.3

Benchmarking different survey platforms

Due date: month 36 (September 2025)

Deliverable type: R

Lead beneficiary: PreFor

Dissemination Level: PUBLIC

Main author: Markus Karppinen, Nived Chebrolu, Maurice Fallon, Martin Jacquet

Contents

1	Introduction	3
2	Platforms	3
2.1	Above-canopy aerial platforms	3
2.2	Below-canopy aerial platforms	4
2.2.1	RMF drone	4
2.2.2	Vision-only drone	4
2.2.3	Leica drone	5
2.3	Terrestrial platforms	5
2.3.1	Frontier Backpack	6
2.3.2	Anymal Quadruped	9
3	Conclusions	13

1 Introduction

During the DigiForest project, several robotic survey platforms were developed to suit forestry applications. These included aerial systems operating both above the canopy (PreFor) and below the canopy (developed by TUM and NTNU). In addition, terrestrial survey platforms were tested, including the AnyMAL robot and the Frontier mobile mapping system developed by the University of Oxford.

All platforms were evaluated during field integration trials conducted in Evo, Finland, in May 2023 and July 2025, as well as in Switzerland in March 2023 and July 2024. In this deliverable performance of each of the systems are described by speed, coverage, accuracy and specific focus on repeatability of the autonomous missions of the robotic based survey platforms.

2 Platforms

2.1 Above-canopy aerial platforms

Above-canopy data collection was carried out using the *PreFor* aerial platform to enable large-scale forest mapping. The data was provided in point cloud format to support the development of methods for merging aerial and terrestrial datasets, and to enhance understanding of forest dimensional structure for robotic navigation. The aerial data was also used to assist in the localization of the *SAHA* machine during autonomous navigation testing.

Data collection was conducted using PreFor’s in-house developed scanning systems. Early in the project (2023 campaigns), a Velodyne HDL-32E LiDAR sensor was used, while later campaigns employed a Hesai XT32-M2X. Data acquisition was performed using unmanned aerial vehicles (UAVs), with the LiDAR payloads mounted directly onto the aircraft. The choice of UAV significantly affected the efficiency and productivity of data collection. However, productivity, measured in scanned hectares per hour, does not differ significantly between engine types, as gasoline-powered UAVs require more complex take-off procedures.

Three UAV platforms were used during the campaigns:

- **Avartek Boxer Hybrid** – powered by a gasoline engine,
- **Aurelia X6 MAX** – electric-powered,
- **DJI Matrice 600** – electric-powered.

The performance comparison of these platforms is presented in Table 1.

Table 1: Summary of the used platforms for above canopy aerial data collection.

Platform	Cost (e)	Flight time (min)	Productivity (ha/hour)
Avartek Boxer Hybrid	60 000	120	25
Aurelia X6 MAX	12 000	40	20
DJI Matrice 600	6 000	40	20

The main outcome of the above-canopy platform was the generation of point cloud data covering large forest areas. In total, 120 hectares of forest were scanned in Finland and 80 hectares in Switzerland. During the project, a fully automatic methodology

for merging aerial and terrestrial datasets was developed and integrated into PreFor's cloud-based software platform. The point cloud data produced by the Frontier mobile mapping device is directly compatible with this system. Accuracy of the above-canopy platform in producing forest inventory is described in deliverable *D5.1: Report on forest inventory accuracy*.

2.2 Below-canopy aerial platforms

Several platform were exploited by the various partners for below-canopy exploration, with varying sensing capabilities and cost. This section presents each platform and its specific capabilities.

The expected outcome of below-canopy surveying missions is LiDAR data for mapping and tree-level segmentation, inventory and inspection.

2.2.1 RMF drone

The Resilient Micro-Flyer [**rmf**] is a class of aerial robots developed at NTNU. Because of their small scale and collision resilience, they are well suited for under-canopy exploration and surveying.

The main RMF system is a drone with size $38 \times 38 \times 24$ cm and weights 1.45 kg.

The platform flight time for one battery is around 8 minutes, which allows to collect data over a surface of around 0.25 ha. A bigger version of the same system has also been developed in the scope of the Digiforest project. This larger ($52 \times 52 \times 30$ cm) and heavier (2.6 kg) system accommodates a flight time up to around 12min, allowing to cover a larger surface per battery time. Both version of the system are pictures in Fig. 1

It embarks a Ouster OS1 LiDAR for mapping and data acquisition, as well as a front facing RGB (or RGBD for the larger system) camera for inspection.



Figure 1: Both RMF-class drones used in the project. The smaller (left) can be ferried in the marsupial system, while the larger (right) allows a longer flight time and large coverage per battery.

2.2.2 BLK2FLY

The BLK2FLY drone is a commercial product developed by Hexagon to obtain high-quality reconstructions of indoor and outdoor environments. The drone has a size of $60 \times 53 \times 19$ cm and weights 2.6 kg including the battery. It includes a BLK LiDAR that can generate up to 420000 points per second while its measurements range from 0.5 to

25m. As for the visual perception, it includes 5 cameras to cover different directions of the drone. The estimated flight-time is of 10 minutes.



Figure 2: The BLK2FLY drone, equipped with BLK LiDAR and visual cameras, is an ideal embodiment for forest mapping thanks to its ability to map in all kind of terrains.

2.2.3 Vision-only drone

The vision-only drone is an effort from TU Munich to develop an MAV that can perform under-canopy exploration and surveying by solely relying on passive-visual sensing, effectively reducing the cost of the deployed robot since it does not require an additional LiDAR. The drone is based on commercially available components: HolybroS500 V2 ARF-Kit airframe, an mRo PixRacer R15 flight controller, an Intel RealSense D455 camera and an NVIDIA Jetson Orin NX 16GB computer for on-board processing. The total weight is around 1.2kg and the flight time is around 15-20 minutes. The coverage of this MAV is smaller compared to that of the RMF as it does not have a LiDAR and relies only on visual sensing, reducing the reliability of its depth estimates. An illustration of the MAV can be found in Fig. 2.



Figure 3: Vision-based MAV used in the project. This system does not use a LiDAR, providing a cheaper option for forest surveying as a complement to other platforms.

2.3 Terrestrial platforms

In this section, we summarize the performance of the two main platforms for terrestrial mapping: The Frontier payload mapping device, which has been developed by Oxford, deployed on a backpack and also used integrated on the Anybotics Anymal D quadruped robot. Copies of the Frontier device were made for and used by University of Bonn, PreFor and the Free University of Bozen-Bolzano during the project.

While we recognize that there is a significant reliability and utility difference between a person carried Frontier device and the Anymal robot, the focus afforded by the DigiForest project allowed us to demonstrate that mapping technology developed for the Frontier (and easily carried by a human) could also be immediately deployed on the Anymal with little to no modification. We clearly state that while the Anymal provides promise and potential as an autonomous terrestrial forest inventory solution there remain major challenges in the face of deploying this solution commercially - in particular its limited ability to traverse complex forest terrain and its high cost.

Nonetheless, during the project both platforms were well tested and a complete, multi-session mapping pipeline was built to which either platform could feed data. The pipeline can also co-register that data into a large, combined and georeferenced point cloud map. This point cloud is then suitable for individual tree segmentation and for detecting DBH, height and species as described in other deliverables during the project.

In the following section we will discuss recent research to improve the place recognition capability of the system — specifically using the forest inventory. Then in Section 2.3.2 we will discuss the overall performance capability of the autonomous forest inventory system developed for Anymal.

Table 2: Comparison between a human carried backpack with a Frontier payload and the Anymal Quadruped.

Platform	Cost (€)	Operation time (min)	Productivity (ha/hour)
Backpack w/ Frontier	20 000	60	3
Anybotic Anymal D	150 000	60	1.8

2.3.1 Frontier Backpack

The frontier backpack system was developed with the goal of being a rapid multi-sensor data capturing setup for developing and testing perception capabilities. We developed and tested multi-session SLAM map merging capabilities (described in deliverable D4.2) using the longitudinal datasets captured using the backpack setup. We have recorded longitudinal datasets in Stein-am-Rhein over three periods during the project, March 2023, October 2023 and July 2024 respectively. These datasets have been shared publicly [[digiforests2025icra](#)]. We have also made longitudinal recordings at Evo, Finland over two periods, May 2023 and June 2025.

Analysis of the real-time forest inventory

The Frontier device computes a forest inventory in real-time using the RealTime-Trees pipeline [[Freissmuth2024](#)]. This pipeline segments individual trees from the LiDAR point cloud, and fits stacks of cone frustums to each tree stem. The output of the system is a list of tree positions, diameters at breast height (DBH), and heights.

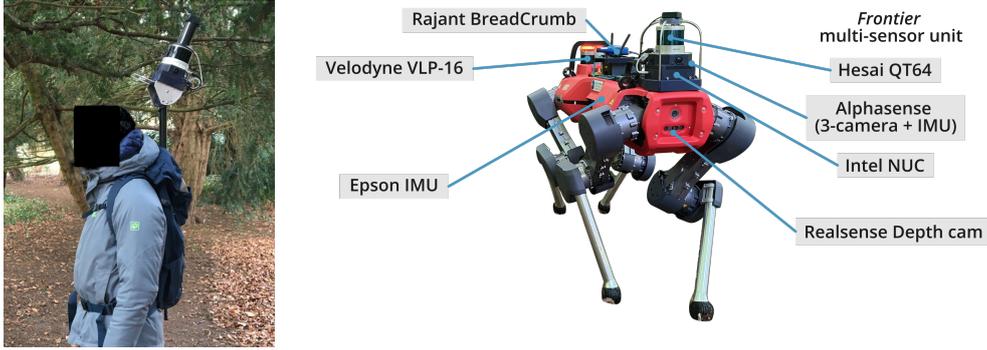


Figure 4: Left: Oxford's Frontier mapping device backpack carried and Right: Deployed on the ANYmal robot with other sensors also mentioned.

	Plot	Conifer	Mixed	Deciduous	All
	Detection Recall \uparrow	98.3%	98.6%	97.4%	98.1%
DBH	RMSE [cm] \downarrow	1.18	2.22	2.38	1.93
	Bias [cm] \downarrow	0.02	0.34	2.05	0.65
	Std [cm] \downarrow	1.17	2.18	1.04	1.81
RMSE Stem	Diameter [cm] \downarrow	2.91	3.14	3.12	3.02
	Center [cm] \downarrow	5.78	9.56	14.88	8.05
Mean Height	Ours [m] \uparrow	8.36	6.30	3.30	6.12
	TLS [m] \uparrow	17.16	15.22	5.94	10.22

Table 3: Evaluation of our pipeline on three different plots. We report the RMSE of DBH estimates (relative to manual measurements) and the RMSE of the stem diameter and curvature measured along the entire stem (relative to a TLS based model). Additionally, we measured the mean height of the reconstructed stems and the detection recall of the clustering algorithm.

The details of the pipeline are described in *D5.1: Report on forest inventory accuracy*.

We have evaluated this pipeline using backpack data acquired at Stein-am-Rhein (Switzerland), Evo (Finland) and Forest of Dean (UK). The accuracy of the inventory was evaluated using ground truth data obtained from manual measurements and terrestrial laser scanning (TLS). The results are summarized in Tab. 3, where we note that the RMSE of DBH estimates is below 2 cm while being able to automatically detect up to 98% of the trees in the plot.

As we expect to have real-time performance, we have profiled the computation time of the whole pipeline. Fig. 4 presents the average runtime of the components of our pipeline for our different datasets. With an overall mean runtime of 9.69 s and a standard deviation of 4.41 s, the algorithm is able to run approximately twice as fast as the capture frequency.

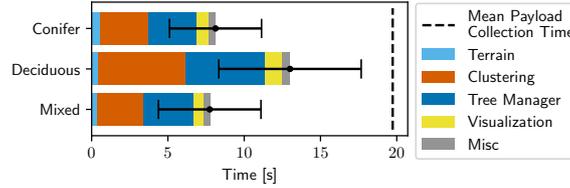


Figure 5: Runtime of individual components of our pipeline given by mean and standard deviation. With a time budget of 19.8s per payload, our approach runs almost twice as fast as required.

We also report system memory usage (Tab. 4) which, with an average of 5.69 G, is well within the capabilities of our system. Perhaps the biggest advantage of our approach is in storage space. The average storage requirements of the final output data are in the range of 160 M for 0.7 hectare including the point clouds of every tree, and only 6 M when storing just the reconstruction results. This compares to 13.5 G for the raw data and hundreds of gigabytes for a traditional TLS system, which makes our system three orders of magnitude more efficient in terms of storage. This detail is critical as data acquisition is arduous and done in remote locations. Furthermore, compared to the intrinsic value of the raw measurements, the cost of storage is immense.

		Conifer	Deciduous	Mixed
Memory		5.18 G	6.75 G	5.15 G
Storage	w/ clouds	158.0 M	187.7 M	143.1 M
	w/o clouds	1.5 M	9.2 M	8.0 M

Table 4: Memory (during runtime) and Storage requirements (afterwards) for final output data (with and without storing the point clouds for every tree).

Analysis of place recognition using forest inventory

In our previous work [oh2024evaluation], we evaluated the performance of several state-of-the-art place recognition methods in forest environments. These place recognition methods typically need to store the entire point cloud map or a large number of descriptors for each map segment, which can be inefficient in terms of storage and computation. To address this, we developed a new method called Tree-Loc [mjung2025treeloc], that builds localization method only using the forest inventory, which is a compact representation of the environment.

We evaluated Tree-Loc on on datasets captured at *Evo*, *Stein am Rhein*, and another publicly available dataset *Wild-Places*[knights2023wild]. We reported Recall@1 (R@1), maximum F1 score, and AUC under a 10m threshold for place recognition in Fig. 5. For localization, we used Recall@50cm (R@50), defined as pose estimation with translation error (TE) ≤ 0.5 m and rotation error (RE) $\leq 5^\circ$, along with the success rate (SR), the ratio of true positives meeting the R@50 criteria.

A key contribution of this work is its storage efficiency. we compared storage requirements across three sequences collected at the same site as *Evo*: two missions in 2023 (*Mission01* and *Mission02*) using Hesai XT32, and one in 2025 (*Mission03*) using Hesai QT64. As shown in Tab. 5, the global tree database required only tens of kilobytes while encoding rich spatial information for each tree, including the stem

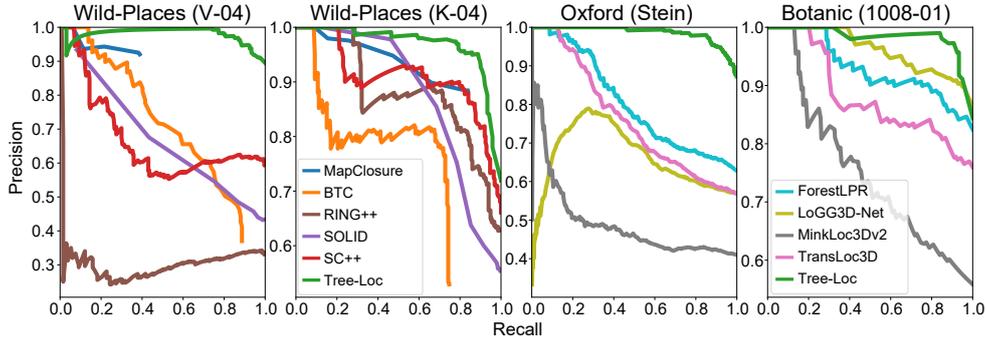


Figure 6: Precision–Recall curves comparing with algorithmic methods (left) and learning-based methods (right). Tree-Loc achieves a higher AUC and maintains high precision at high recall, thereby avoiding the sharp drop in precision observed in baselines. Table 5: Comparison of Storage Requirements for Localization

Mission	# Scenes	Raw PCD	BTC Desc.	SC++ Desc.	Tree-Loc
Mission01	452	1.7 GB	478.1 MB	3.6 MB	69.9 KB
Mission02	788	3.0 GB	857.6 MB	6.5 MB	136.7 KB
Mission03	222	0.2 GB	241.2 MB	1.8 MB	61.0 KB

axis, center, DBH, base height, and observation count. This is faster than other SOTA approaches that slowed by processing heavy point clouds. A key advantage of the global tree database is that its size scales with the number of unique trees rather than trajectory length, making it more suitable for long-term mapping. Moreover, its tree-level representation supports incremental updates: new missions can be aligned and integrated into the existing map.

The results of our analysis suggest strong performance compared to other state-of-the-art (SOTA) methods, with accuracy comparable to learning-based methods. It enables fast localization within 50 ms, and supports multi-session map maintenance with a compact representation, 3-4 orders of magnitude smaller than methods using dense point clouds or descriptor maps.

2.3.2 Anymal Quadruped

In the following section we focus on the field performance of the forest inventory system developed for the Anymal quadruped. Specifically the system incorporates the single session SLAM system (presented in earlier deliverables specifically *D4.2: Report describing map server — for robot and base station mapping*) plus a free space exploration planner which allowed the robot to work its way around obstacles (including trees) when carrying out a human-directed inspection pattern (previously reported in *D3.3: Per-platform Traversability-Aware Navigation in Forests*). Meanwhile running onboard the robot during these missions was the real-time tree segmentation, fitting and collation system which we call RealTimeTrees. This module produces a database of the individual tree positions, diameters, heights and more recently species with accuracy competitive with TLS based pipelines in real-time while operating on the robot or backpack. The RealTimeTrees pipeline was presented in *D5.1: Report on forest inventory accuracy*.

As such, the performance of the localization, mapping and inventory subsystems is the same when running on the robot as when running in the backpack (as described



Figure 7: During Digiforest, we deployed the ANYmal autonomous legged robot system in Finland, Switzerland, and the United Kingdom. We gained valuable insights into the capabilities of legged platforms for forestry inventory in different forest types.

in the previous section) with the key different being the mobility of the robot and its ability to access different physical locations. Thus in the following we zero in on the robots mobility and the frequency with which we need to intervene to reset its autonomy system using normal autonomy metrics.

The following presents the benchmarking section of our complete summary paper on this system which was recently published:

- *Building Forest Inventories with Autonomous Legged Robots - System, Lessons, and Challenges Ahead*, Matias Mattamala, Nived Chebrolu, Jonas Frey, Leonard Freißmuth, Haedam Oh, Benoit Casseau, Marco Hutter, Maurice Fallon. IEEE Transactions on Field Robotics, June 2025.

We report the main results obtained with the system for three different aspects: the autonomy system, the forest inventory performance, and potential future applications. We present the evaluation metrics, data, and findings of each case. Table 6 provides a summary of the main outcomes.

Analysis of the Anymal 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 [Paz2020]. Particularly, we chose the MDBI and 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. 7 (left).

We matched the intervention signals to the paths recorded on each mission, which

Mission Name	Autonomy	Dist. traveled [m]	In-terv. [#]	MDBI [m]	MTBI [s]	Forest Inventory Area	
	Time [s]					covered [ha]	Trees [#]
Evo-01	575.6	270.3	2	84.8	176.2	0.33	X
Evo-02	432.0	233.6	0	233.6	432.4	0.31	X
Evo-03 [†]	816.8	301.1	7	30.5	72.2	0.36	X
Evo-04	988.4	336.6	7	39.4	114.2	0.36	X
Evo-05	1275.5	609.7	10	51.1	99.6	0.58	X
WyO-01	436.9	215.0	2	69.9	136.4	0.29	X
Dea-01	1283.5	665.4	8	66.0	127.8	0.93	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	0.22	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.

^x Forest inventory system was not available.

Table 6: Summary of the results for all campaigns. We report the autonomy metrics, and the performance of the forest inventory system.

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:

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

$$\text{MTBI}_{\mathcal{M}} = \frac{1}{N} \sum_{i \in \mathcal{M}} t_{\text{auto}}^i. \quad (2)$$

Tab. 6 reports the number of interventions as well as the MDBI and MTBI for all the campaigns. While these metrics 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. 8a.

In general, we observed that when all missions are considered, the interventions follow a Poisson distribution in both time and distance, concentrated around short interventions. The plots indicate that most of the interventions required to manually drive the robot for less than 20 s, or 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, especially in the larger scale missions.

We further revised the TBI and DBI from a distribution point of view, which is shown in Fig. 8b. We removed all the missions where no interventions were reported.

We note that the TBI and DBI are also Poisson-distributed, 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 the mission planner did not produce significant progress by the proposed re-planning strategy. This figure suggests that having a full picture of the distribution provides

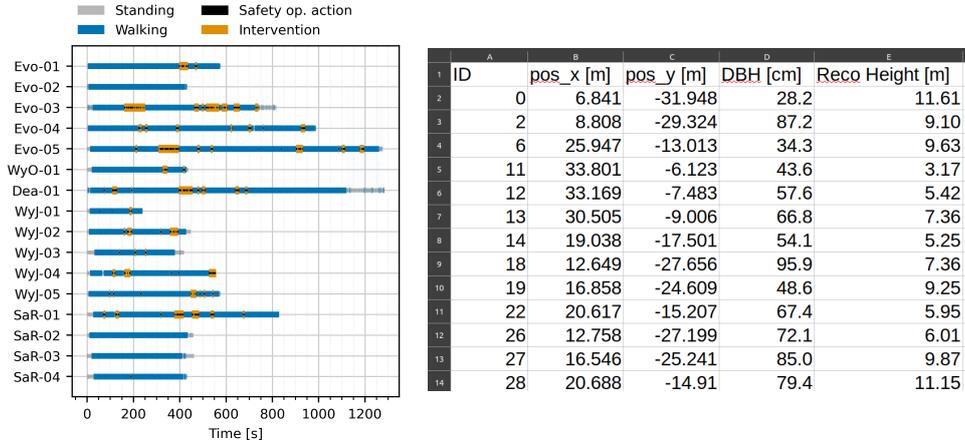
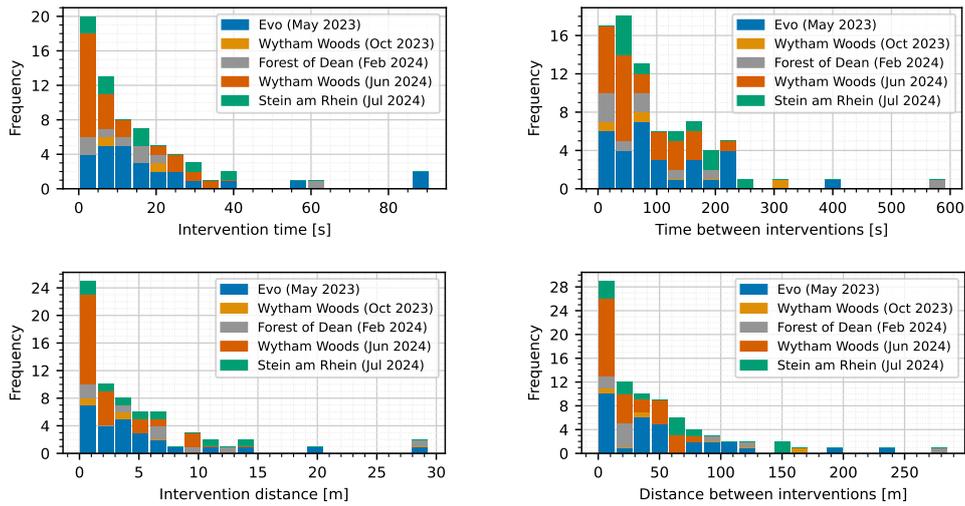


Figure 8: Left: Summary of the autonomy results obtained across the five campaigns. Periods of full autonomy (■), controlled by the safety operator (■). Safety *intervention* events (■). Right: Forest inventory in spreadsheet format produced from Forest of Dean campaign showing tree positions, DBH, and height of the detected trees.



(a) Histogram of duration (top) and distance traveled (bottom) **during** interventions. Most interventions were short—<20 s. (b) Distribution of time (top) and distance (bottom) **between** interventions for all missions.

Figure 9: Statistical analysis of autonomy and interventions.

additional insights, complementary to the summaries provided by the MTBI and MDBI.

As a general conclusion of this analysis, we acknowledge 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 limitations of our system are discussed further in Sec. 3.

Analysis of the Anymal 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. 6 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 0.93 ha. The typical scanning speed was 1.8 ha h^{-1} , with the robot walking at an average speed of 0.6 m s^{-1} across the different campaigns.

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. 6. While we did not have ground truth TLS measurements for many of the test sites, in our related prior work [Freissmuth2024] we studied the accuracy of our online inventory system. Therefore, we expect that the forest inventory can achieve an average accuracy of DBH 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. 7 (right).

Fig. 9 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.

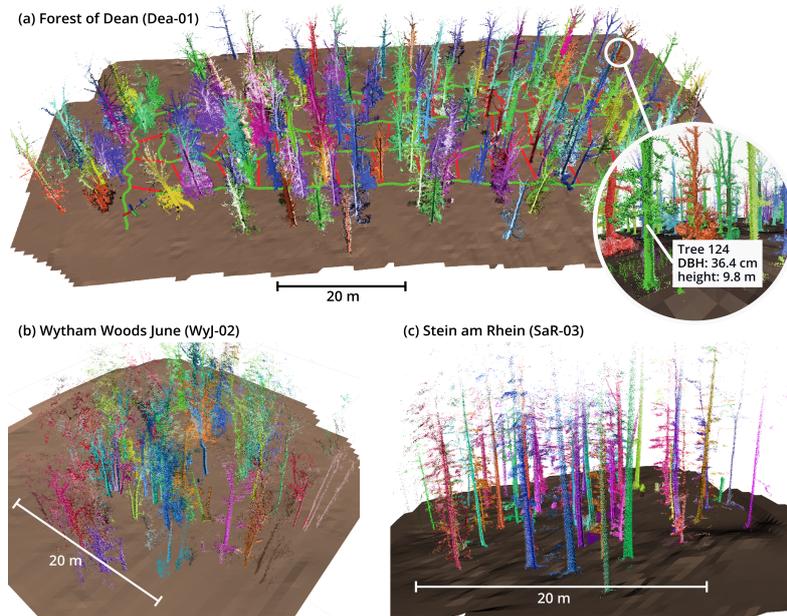


Figure 10: Illustrative example of the different outputs of our forest inventory pipeline in different autonomous deployments. (a) Forest of Dean (Dea-01), (b) Wytham Woods (WyJ-02), (c) Stein am Rhein (SaR-03). We can observe the differences in the tree segmentations due to the diverse species present in the forests considered.

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 earlier, as well as the sensing specifications of the MLS unit being carried. These aspects are further discussed in Sec. 3.

3 Conclusions

The DigiForest project has demonstrated that a wide range of robotic survey platforms, above-canopy aerial systems, below-canopy aerial drones, and terrestrial devices—can each contribute unique and complementary strengths to sustainable forestry applications. Above-canopy UAVs proved highly effective for large-scale forest mapping, providing consistent point cloud data that can be merged with terrestrial scans to create comprehensive forest inventories. Below-canopy platforms, such as the RMF drones, showed promise for detailed tree-level surveying in cluttered environments, albeit with limited endurance. Terrestrial systems, particularly the Frontier backpack and the Anymal quadruped, enabled high-accuracy, real-time forest inventories with strong potential for longitudinal monitoring and repeatability.

Benchmarking across diverse field campaigns in Finland, Switzerland, and the UK highlighted both the progress made and the remaining challenges. While autonomous navigation and tree inventory pipelines reached performance levels competitive with traditional TLS-based methods, issues such as battery endurance, terrain traversability, and operational cost continue to limit scalability and commercial deployment. Importantly, the integration of aerial and terrestrial data into unified point cloud maps marks a significant step forward, enhancing the accuracy, efficiency, and reusability of forest survey data.

Overall, the results confirm the feasibility of robotic-based survey platforms for forestry while also identifying clear research priorities for the future. These include extending platform autonomy, improving adaptability to complex terrains, and refining data fusion techniques. Addressing these challenges will be crucial for enabling robust, cost-effective, and sustainable forest monitoring solutions at scale.