JOURNAL OF GEOSPACE SCIENCE

Vol. 6 No. 1 (2026)

E-ISSN : 2413-6093
https://jogss.org

Heterogeneous Aerial Ground Robot Collaboration for Vineyard Monitoring with Anisotropic Sensors and Inspection Waypoints

Eliza Murray, Stella Palmer, Joonseo Jang

Faculty of Science, University of Helsinki, Helsinki, Finland

ABSTRACT

The integration of autonomous robotic systems into precision viticulture represents a transformative shift in how vineyard health, yield estimation, and disease detection are managed. Traditional monitoring techniques rely heavily on manual labor or homogeneous robotic deployments, both of which suffer from significant limitations in scalability, energy efficiency, and data resolution. This paper presents a comprehensive framework for heterogeneous field-robot collaboration, specifically focusing on the synergistic deployment of Unmanned Aerial Vehicles and Unmanned Ground Vehicles equipped with anisotropic sensors. Unlike omnidirectional sensors, anisotropic sensors such as RGB-D cameras, directional multispectral imagers, and phased-array radars possess strict directional constraints, limited fields of view, and range-dependent resolution. This spatial dependency introduces profound challenges in coverage path planning and task allocation. We propose a decentralized collaborative architecture wherein aerial units conduct high-level topological mapping and anomaly detection, dynamically generating high-resolution inspection waypoints for ground units. The ground units subsequently navigate the complex under-canopy environment using a novel anisotropic coverage algorithm that maximizes sensor utilization while minimizing traversal energy. Through extensive simulated and real-world vineyard trials, the proposed heterogeneous framework demonstrates a substantial improvement in inspection efficiency, reducing overall energy consumption while maximizing canopy coverage. The results indicate that accounting for sensor anisotropy within a heterogeneous collaborative framework significantly bridges the gap between aerial macro-monitoring and ground-level micro-inspection.

KEYWORDS

Precision Viticulture, Heterogeneous Robotics, Anisotropic Sensors, Collaborative Path Planning

1. Introduction

1.1 The Paradigm of Precision Viticulture

The global agricultural sector is currently undergoing a massive technological revolution driven by the urgent need to optimize resource utilization, increase crop yields, and mitigate the environmental impacts of traditional farming practices. Within this broader context, precision viticulture has emerged as a specialized domain that applies advanced information technologies and engineering principles to the management of vineyards. Vineyards present a uniquely complex agricultural environment characterized by highly structured yet biologically variable architectures. Grapevines are typically cultivated in rows utilizing various trellis systems, such as the vertical shoot positioned system, which creates a highly organized geometric layout at the macro level but exhibits extreme morphological complexity at the micro level. Managing these environments requires precise, timely, and spatially accurate data regarding plant water stress, nutrient deficiencies, pest infestations, and disease outbreaks. Early detection of anomalies such as powdery mildew, downy mildew, or phylloxera is absolutely critical, as these biological threats can rapidly devastate an entire harvest if left unaddressed. Consequently, the demand for continuous, high-resolution vineyard monitoring has catalyzed the transition from manual, human-centric scouting to automated, data-driven methodologies [1]. Human scouting, while capable of high qualitative assessment, is fundamentally constrained by labor costs, fatigue, and the sheer vastness of modern commercial vineyard estates. As a result, the industry has increasingly turned toward remote sensing technologies to bridge this operational gap [2]. Satellite imagery and manned aircraft surveys initially provided the macro-level insights necessary for rudimentary precision viticulture, offering broad vegetation indices such as the Normalized Difference Vegetation Index. However, these traditional remote sensing platforms suffer from severe limitations regarding temporal resolution, cloud cover occlusion, and insufficient spatial resolution for early-stage disease detection at the individual leaf or cluster level [3].

1.2 The Transition to Robotic Interventions

To overcome the spatial and temporal limitations of satellite and manned aircraft remote sensing, the deployment of autonomous robotic platforms has gained unprecedented momentum in recent years. Unmanned Aerial Vehicles have become ubiquitous in precision agriculture due to their rapid deployment capabilities, ability to traverse difficult terrain without causing soil compaction, and capacity to host a variety of lightweight optical payloads [4]. Unmanned Aerial Vehicles excel at capturing high-resolution orthomosaics and multispectral maps of the vineyard canopy from a top-down perspective. They provide farmers with a comprehensive overview of vineyard vigor and spatial variability in a matter of hours. Despite their undeniable utility, aerial platforms operate under severe fundamental constraints. Their payload capacity is strictly limited by aerodynamic and energetic constraints, restricting the types of sensors they can carry. More importantly, the top-down perspective of an Unmanned Aerial Vehicle is often incapable of penetrating the dense vegetative canopy of mature grapevines. Many critical viticultural indicators, such as grape cluster morphology, early-stage bunch rot, and trunk diseases, are exclusively visible from a sub-canopy or lateral perspective. This inherent limitation necessitates the introduction of Unmanned Ground Vehicles into the monitoring paradigm [5]. Ground platforms boast significantly larger payload capacities, extensive battery endurance, and the mechanical stability required to carry heavy, sophisticated instrumentation such as high-power active illumination systems, terrestrial light detection and ranging scanners, and computational hardware for edge processing. By navigating between the vineyard rows, Unmanned Ground Vehicles can acquire lateral, high-resolution data of the fruiting zone and under-vine environment. However, ground vehicles are relatively slow, frequently impeded by rough terrain, cover crop variations, and soil moisture conditions. The solitary deployment of either aerial or ground robots therefore yields an incomplete monitoring solution. This realization has driven contemporary research toward heterogeneous multi-robot systems, where the complementary strengths of aerial and ground platforms are fused into a singular, cohesive operational unit [6].

1.3 The Challenge of Anisotropic Sensing in Complex Environments

While the architectural integration of aerial and ground robots addresses the mechanical and perspective

limitations of vineyard monitoring, the sensory acquisition process itself introduces a profound theoretical challenge known as anisotropic sensing. In classical robotic coverage planning, sensors are frequently idealized as isotropic, meaning their sensing capabilities are perfectly uniform in all directions around the robot, typically modeled as a generic circle or sphere of coverage. This assumption holds true for specific environmental sensors, such as omnidirectional gas detectors, temperature probes, or basic proximity sensors [7]. However, the primary instruments utilized in precision viticulture monitoring are highly anisotropic. High-resolution RGB cameras, multispectral imagers, depth cameras, and directional radar systems possess strict fields of view. Their ability to acquire useful data is entirely dependent on their orientation relative to the target object. An anisotropic sensor is fundamentally characterized by a directional sensing cone or frustum, where the quality of the gathered data degrades as a function of both distance and angular deviation from the central optical axis. In a vineyard environment, this anisotropy is further exacerbated by the complex geometry of the grapevine canopy. A sensor must not only be pointed at the target, but it must also achieve an optimal standoff distance to ensure the target falls within its depth of field while avoiding physical collisions with overhanging branches [8]. When a ground robot navigates a vineyard row, it cannot merely drive past a vine to inspect it; it must orient its anisotropic sensors laterally toward the canopy, adjust its speed to prevent motion blur, and potentially articulate its sensor payload to peer through gaps in the foliage. Planning a coverage path for a robot equipped with anisotropic sensors is therefore vastly more complex than traditional area coverage algorithms. The algorithm must simultaneously optimize the Cartesian position and the angular orientation of the robot in a highly constrained environment [9]. When scaling this problem to a heterogeneous multi-robot team, the complexity increases exponentially, as the aerial robot must predict the optimal inspection poses for the ground robot and allocate tasks that maximize the utility of the ground robot's anisotropic sensors.

1.4 Objectives and Contributions

The primary objective of this research is to formulate, implement, and validate a comprehensive framework for heterogeneous field-robot collaboration tailored specifically for vineyard environments, with a central emphasis on the mathematical and operational constraints of anisotropic sensors. We recognize that treating cameras and directional imagers as simple point-radius sensors leads to severely suboptimal path planning and incomplete canopy coverage. Therefore, this paper seeks to bridge the gap between high-level task allocation in multi-agent systems and the low-level geometric realities of optical sensing in complex agricultural terrains [10]. This paper introduces a novel architecture that operates on a two-tier hierarchical framework. The first tier involves a rapid aerial survey conducted by an Unmanned Aerial Vehicle, which generates a topological map of the vineyard and identifies regions of interest based on macroscopic canopy anomalies. The second tier involves an Unmanned Ground Vehicle equipped with a highly articulate, anisotropic sensor payload. We propose a bespoke path-planning algorithm that translates the aerial regions of interest into optimal ground-level inspection poses, strictly enforcing the directional constraints of the ground vehicle's sensors. By decoupling the macroscopic exploration from the microscopic inspection, our framework minimizes the energy expenditure of the ground vehicle while ensuring that every biological anomaly is captured with maximum sensor fidelity. The contributions detailed in the subsequent sections include a formal description of the anisotropic sensor model tailored for viticulture, a collaborative communication protocol for heterogeneous agents, and extensive empirical validation across both simulated and physical vineyard environments.

2. Literature Review

2.1 Evolution of Agricultural Monitoring Technologies

The historical trajectory of agricultural monitoring has been defined by a continuous pursuit of higher resolution data and reduced human intervention. Early implementations of precision agriculture relied almost exclusively on satellite remote sensing. These systems provided valuable macro-scale data, allowing agronomists to observe large-scale variations in soil moisture and crop vigor. However, the temporal latency of satellite data, often characterized by revisit times of several days or weeks, severely limited its utility for dynamic disease management [11]. Furthermore, the spatial resolution of early

civilian satellites was insufficient to distinguish between individual plant canopies and the inter-row cover crops, leading to significant data pollution. The advent of low-altitude remote sensing via manned aircraft improved spatial resolution, but the associated operational costs restricted its use to high-value cash crops and large corporate farming operations [12]. The miniaturization of micro-electromechanical systems and the subsequent proliferation of affordable Unmanned Aerial Vehicles fundamentally disrupted this paradigm. Researchers began equipping small multirotor and fixed-wing aerial platforms with lightweight multispectral cameras, enabling on-demand, high-resolution mapping of agricultural fields. A substantial body of literature has documented the efficacy of aerial platforms in calculating vegetation indices, estimating biomass, and detecting water stress [13]. However, as the focus shifted toward high-density perennial crops like orchards and vineyards, the limitations of aerial monitoring became increasingly apparent. Grapevines possess a complex, vertically oriented canopy structure. Many critical viticultural features, such as the grape clusters themselves, are deliberately shielded from direct sunlight by the upper canopy to prevent sunburn, rendering them entirely invisible to top-down aerial sensors. This realization spurred parallel research into agricultural ground robotics, leading to the development of massive, heavily instrumented platforms capable of traversing vineyard rows and capturing lateral data [14]. Yet, these ground platforms suffered from their own distinct operational bottlenecks, primarily related to slow traversal speeds and the immense computational burden of processing continuous high-resolution data over vast acreages.

2.2 Heterogeneous Multi-Robot Systems in Agriculture

Recognizing the mutually exclusive strengths and weaknesses of aerial and ground platforms, the robotics community has increasingly focused on heterogeneous multi-robot systems. The core philosophy underpinning heterogeneous collaboration is that different robotic agents should be allocated tasks that perfectly align with their morphological and kinematic advantages. In agricultural contexts, this typically manifests as an asymmetrical leader-follower or scout-worker dynamic [15]. Initial research in this domain focused heavily on logistical cooperation, such as using an aerial vehicle to find a safe navigational path for a ground vehicle traversing unstructured terrain [16]. Subsequent studies advanced this concept into the realm of active monitoring. For instance, several frameworks have been proposed where an aerial drone rapidly scans a field to detect localized areas of weed infestation. The drone then transmits the geographic coordinates of these weed patches to an autonomous tractor, which navigates directly to the specified locations to apply targeted herbicides [17]. This decoupled approach significantly reduces the time and fuel required for the tractor, as it no longer needs to perform exhaustive coverage of the entire field. Similar architectures have been explored for robotic harvesting, where drones map the distribution of ripe fruit to optimize the routing of slow-moving ground harvesters [18]. Despite these advancements, the application of heterogeneous systems to vineyard monitoring remains under-explored. Vineyards introduce unique spatial constraints; robots cannot freely move across the field but are strictly confined to the narrow corridors between the trellis rows. This topological restriction vastly complicates task allocation and routing algorithms, as the ground vehicle must often travel significant distances down a row to reach a specific anomaly, negating some of the efficiency gains of the scout-worker dynamic [19]. Furthermore, existing literature on heterogeneous agricultural robotics largely glosses over the intricate details of how the ground robot actually perceives the target once it arrives, often relying on simplified sensing models that fail to capture the reality of field operations.

2.3 Sensor Anisotropy and Coverage Optimization

The concept of sensor anisotropy has been rigorously studied within the specialized fields of wireless sensor networks and visual surveillance, but its integration into mobile agricultural robotics is a relatively recent development. In wireless sensor networks, researchers recognized early on that directional antennas and visual sensors could not be modeled using the ubiquitous Boolean disk model. Instead, models based on circular sectors or spherical frustums were developed to accurately represent the limited field of view and directional gain of these sensors [20]. Coverage optimization algorithms for anisotropic sensor networks focus heavily on determining the optimal spatial orientation of static nodes to maximize the

monitored area while minimizing overlapping redundant coverage [21].

Transitioning these static models to mobile robots introduces the problem of dynamic anisotropic coverage path planning. Traditional coverage path planning algorithms, such as the standard Boustrophedon cellular decomposition, assume that a robot simply needs to pass its center of mass over a given area to cover it. This is highly effective for isotropic tasks like vacuum cleaning or broad-acre pesticide spraying [22]. However, when an autonomous ground vehicle is tasked with visually inspecting a vineyard canopy using an RGB-D camera, passing its center of mass near the vine is insufficient. The robot must achieve a specific viewing angle, maintain a minimum focal distance to prevent image blurring, and avoid exceeding the maximum range of its depth sensor [23]. Several researchers have attempted to solve this by discretizing the environment into a grid and calculating the precise subset of robot poses (position and heading) from which a given target cell is observable. This transforms the coverage problem into a variation of the Art Gallery Problem or the Traveling Salesman Problem with Neighborhoods, both of which are notoriously computationally complex [24]. Recent approaches have utilized heuristic algorithms, such as genetic algorithms and particle swarm optimization, to generate feasible inspection paths for robots with directional sensors in complex 3D environments [25]. While these studies provide a strong theoretical foundation, they are predominantly evaluated in highly structured indoor environments or simplified outdoor simulations, lacking the biological complexity and dynamic occlusion inherent to dense agricultural foliage.

2.4 Critical Limitations and Gap Analysis

A critical synthesis of the current literature reveals a significant operational gap at the intersection of heterogeneous multi-robot systems and anisotropic sensor planning. The vast majority of heterogeneous agricultural frameworks assume that once the ground robot reaches the region of interest identified by the aerial robot, the inspection task is trivially accomplished [26]. Conversely, the literature focusing strictly on anisotropic coverage path planning typically assumes a single, highly capable robot operating in isolation, without the benefit of macroscopic prior information provided by an aerial scout. There is a glaring absence of integrated architectures that leverage the macro-scale data of an aerial vehicle to inform and optimize the micro-scale, directionally constrained inspection paths of a ground vehicle. In a vineyard, this disconnect is fatal to operational efficiency. If an aerial drone identifies a region of potential disease, the ground robot cannot simply drive to the GPS coordinate. It must calculate a specific approach vector, position its chassis to align its fixed or pan-tilt mounted anisotropic sensors perfectly with the diseased foliage, and potentially capture multiple images from varying angles to penetrate the outer canopy layer. Furthermore, the ground robot must accomplish this while managing its own battery reserves and navigating the restrictive inter-row corridors. Current literature lacks a cohesive mathematical and operational framework that seamlessly links aerial topological discovery with ground-based anisotropic inspection routing. This paper directly addresses this gap by proposing a holistic, decentralized collaboration mechanism that explicitly incorporates the directional sensing constraints of the ground vehicle into the task allocation and routing protocols negotiated with the aerial vehicle.

3. Methodology and Analysis

3.1 Collaborative Architecture Overview

The proposed heterogeneous field-robot collaboration framework is designed as a decentralized, asynchronous architecture. This design choice is deliberate, aiming to mitigate the communication bottlenecks and single points of failure inherent to centralized command-and-control structures, particularly in rural vineyard environments where persistent broadband connectivity is rarely guaranteed. The system operates through a series of interconnected operational phases, smoothly transitioning from high-altitude discovery to ground-level precision inspection. The foundational architecture relies on a shared, global coordinate system that maps the physical vineyard into a topological graph, where the nodes represent intersection points at the ends of the vineyard rows, and the edges represent the traversable corridors between the vine canopies. During the primary operational phase, the Unmanned Aerial Vehicle acts as the vanguard. It executes a rapid, high-altitude sweep of the entire vineyard estate. Utilizing its

downward-facing, isotropic multispectral sensors, the aerial unit generates a low-resolution map of canopy vigor. This aerial data undergoes immediate onboard processing using lightweight computer vision algorithms to detect macroscopic anomalies, such as localized patches of chlorosis, significant gaps in the canopy, or indications of severe water stress. These anomalies are clustered into distinct Regions of Interest. Once a Region of Interest is identified, the Unmanned Aerial Vehicle synthesizes a task request packet. This packet contains the geometric centroid of the anomaly, its estimated spatial bounding box, and a preliminary classification of the anomaly type. This data is transmitted asynchronously to the Unmanned Ground Vehicle via a localized, peer-to-peer radio frequency link. Upon receiving the task request, the ground vehicle transitions from an idle or general patrol state into an active inspection routing state. The ground vehicle utilizes the spatial boundaries of the Region of Interest to calculate a highly specific, low-level maneuvering path. Crucially, this low-level path is not simply a straight line to the target. Instead, it is a complex trajectory mathematically sculpted by the anisotropic constraints of the ground vehicle's sensor payload, ensuring that the vehicle arrives at a pose that perfectly aligns its directional sensors with the target foliage [27].

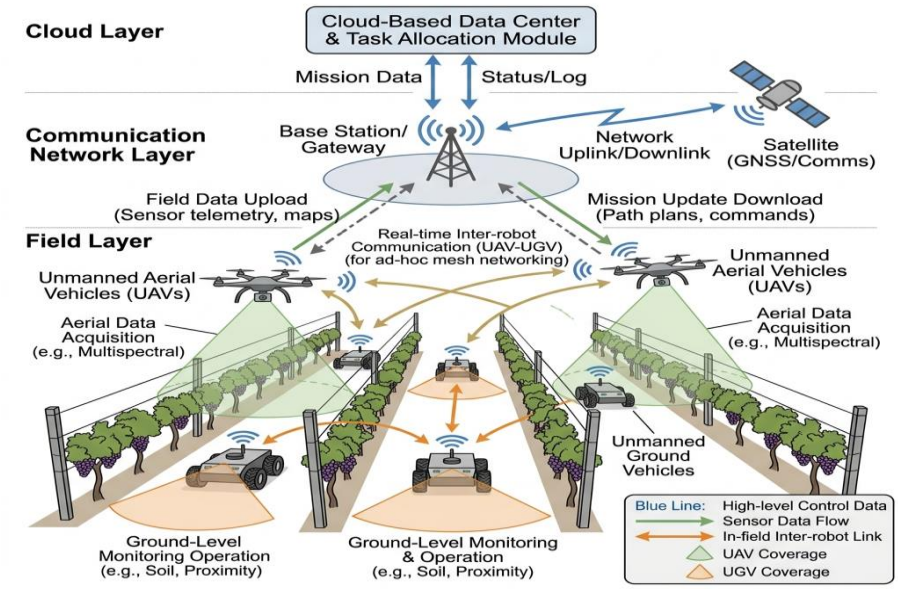


Figure 1: Comprehensive Schematic of the Heterogeneous Field

The asynchronous nature of this architecture allows the Unmanned Aerial Vehicle to continue scanning the remainder of the vineyard while the Unmanned Ground Vehicle independently calculates and executes the intricate maneuvers required for anisotropic inspection. This decoupling of exploration and exploitation is the primary driver of efficiency in our proposed framework [28].

3.2 Heterogeneous Platform Specifications

The successful implementation of this framework requires robotic platforms specifically tailored to their respective roles within the hierarchy. We utilize a highly customized fleet consisting of a multirotor aerial scout and a tracked ground inspection unit. The distinct mechanical and sensory configurations of these platforms dictate their operational capabilities and inform the constraints of our algorithms. The aerial unit is based on a lightweight, quadcopter chassis designed for extended flight duration and rapid maneuverability. Its primary constraint is energy capacity; to maximize flight time, the payload weight is strictly minimized. Therefore, it is equipped solely with a downward-facing multispectral camera array. This sensor is treated as isotropic in the horizontal plane during high-altitude flight, as its field of view footprint on the ground is roughly rectangular and covers the entire width of a vineyard row from its designated operating altitude. The aerial unit also houses an edge-computing module dedicated to running the vegetation index calculations and the anomaly clustering algorithms in real-time, eliminating the need to transmit heavy raw image data back to a base station [29]. Conversely, the ground unit is a heavy-duty, tracked vehicle designed to negotiate steep inclines, muddy terrain, and dense cover crops common in

viticulture. Its massive battery bank not only provides significant driving range but also powers a formidable array of computational hardware and high-fidelity sensors. The defining feature of the ground unit is its sensor payload, mounted on a vertically adjustable mast. This payload consists of high-resolution RGB-D cameras and highly directional active illumination panels. These sensors are strictly anisotropic. They face laterally, perpendicular to the forward motion of the chassis, designed to look directly into the fruiting zone of the adjacent grapevine canopy. The mechanical design enforces a strict coupling between the orientation of the vehicle chassis and the viewing angle of the primary sensors [30].

Table 1: Specifications of the Heterogeneous Robot Fleet

Platform Type	Primary Sensor Payload	Sensing Characteristics	Operational Endurance
Unmanned Aerial Vehicle	Downward Multispectral Array	Macro-scale, Top-down pseudo-isotropic	45 Minutes
Unmanned Ground Vehicle	Lateral RGB-D Cameras	Micro-scale, Strictly Anisotropic	8 Hours
Command Base Station	Long-range Telemetry Radio	Non-sensing, Data Aggregation	Continuous Mains Power

3.3 Mathematical Representation of Anisotropic Sensors

To enable autonomous path planning for the ground vehicle, it is absolutely essential to formalize the behavior and limitations of its anisotropic sensors. Unlike theoretical isotropic models that assume uniform data acquisition within a defined radius, our model rigorously defines the spatial volume in which the sensor can successfully acquire valid data. We conceptualize the sensing region of the laterally mounted RGB-D camera as a three-dimensional asymmetrical frustum extending outward from the sensor's focal point [31]. The physical geometry of this frustum is defined by several critical boundaries. First, there is a minimum sensing range. Objects located closer than this minimum threshold cannot be accurately perceived due to the focal limitations of the lenses and the baseline separation of the depth camera's stereo lenses. If the ground vehicle navigates too close to the canopy, the resulting data is entirely invalid. Second, there is a maximum sensing range. Beyond this distance, the spatial resolution of the point cloud degrades significantly, and active infrared illumination dissipates, rendering the sensor incapable of detecting millimeter-scale diseases. The optimal sensing zone lies strictly between these two radial boundaries. Furthermore, the sensor is constrained by a horizontal and vertical angle of view. The camera can only perceive the environment within a specific angular wedge extending from its central optical axis. The quality of the captured data is not uniform across this wedge; it is highest along the central axis and degrades toward the peripheral edges due to lens distortion and lighting falloff. To model this mathematically, we assign a continuous utility function to the space within the frustum. The utility value is maximized when a target point on the grapevine is located exactly at the optimal focal distance and perfectly aligned with the central optical axis of the sensor. As the target deviates in distance or angle, the utility value decays exponentially [32]. This anisotropic model completely alters the paradigm of coverage planning. To inspect a specific cluster of grapes identified by the aerial drone, the ground robot cannot simply position its center of mass near the target coordinates. It must calculate a specific position and heading such that the target grape cluster falls squarely within the high-utility region of its sensor frustum, while simultaneously ensuring that the physical bulk of the robot chassis does not collide with the trellis wires or overhanging foliage.

3.4 Collaborative Path Planning and Routing Optimization

The core computational challenge of our framework lies in the Collaborative Path Planning algorithm. This algorithm bridges the gap between the high-level Regions of Interest generated by the aerial unit and the low-level kinematic controls of the ground unit. When the ground vehicle receives a cluster of targets from the aerial unit, it must determine the optimal sequence in which to visit these targets and the exact

geometric pose required to inspect each one, all while constrained by the topological layout of the vineyard rows [33]. The algorithm operates in two distinct phases. Phase one involves optimal pose generation. For every individual anomaly within a Region of Interest, the algorithm utilizes the anisotropic sensor model to generate a set of candidate inspection poses. A candidate pose is a specific coordinate and heading within the navigable inter-row corridor from which the anomaly can be viewed with a utility score exceeding a predefined quality threshold. Because the sensor is mounted laterally, these candidate poses typically align the robot parallel to the vine row, at a specific distance from the canopy wall. The algorithm must filter out any candidate poses that violate kinematic constraints, such as poses that would place the vehicle inside a physical obstacle or require an impossible turning radius to reach. Phase two is the routing sequence optimization. Once a valid set of candidate inspection poses is generated for all targets, the algorithm must construct a continuous trajectory that connects these poses. This represents a complex variation of the Generalized Traveling Salesman Problem, as the robot only needs to visit one candidate pose per target, rather than a specific point. Our proposed solution employs a deterministic heuristic approach based on a dynamically weighted cost graph. The nodes in the graph represent the candidate inspection poses, and the edges represent the traversable paths between them. The weight of each edge is a composite cost function that heavily penalizes long travel distances, complex turning maneuvers, and transitions between different vineyard rows, as exiting and re-entering a row is extremely time-consuming for a large tracked vehicle.

Table 2: Algorithm Parameters for Collaborative Path Planning

Parameter Name	Function within Algorithm	Impact on Performance
Pose Utility Threshold	Filters low-quality inspection angles	Higher values increase data quality but restrict routing options
Row-Transition Penalty	Discourages excessive row changing	High penalty forces sequential in-row inspection, saving energy
Sensor Frustum Angle	Defines the anisotropic wedge width	Wider angles allow faster coverage but lower peripheral resolution

By utilizing this two-phase approach, the ground vehicle autonomously translates a generic area alert from the aerial drone into a highly efficient, kinematically feasible inspection trajectory that honors the strict directional requirements of its optical payload. This ensures that every image captured is of maximum diagnostic value, preventing the need for costly secondary passes [34].

3.5 Decentralized Data Fusion and Validation

The final stage of the methodology involves the fusion of the disparate data streams generated by the heterogeneous agents. The primary objective is to validate the low-resolution aerial anomalies against the high-resolution ground truth data. As the ground vehicle executes its optimized inspection route, it captures dense RGB point clouds and multispectral imagery of the designated targets. This raw data is processed locally on the ground vehicle's edge-computing module using convolutional neural networks trained specifically for viticultural disease detection and yield estimation [35]. Because the entire system operates on a unified global coordinate frame, the spatial registration of the ground data with the aerial data is streamlined. When the ground vehicle confirms the presence of a disease, it updates the shared topological map with a high-confidence, multi-modal data node. If the ground vehicle arrives at a Region of Interest and its high-resolution anisotropic sensors determine that the aerial anomaly was a false positive—perhaps caused by a shadow or a benign irregularity in the cover crop—it overrides the aerial classification. This mechanism acts as a critical feedback loop. The aggregated results of these confirmations and rejections are periodically transmitted back to the command base station, where they can be utilized to iteratively retrain and refine the lightweight anomaly detection algorithms running on the aerial unit [36]. Through this continuous cycle of macro-level hypothesis generation by the aerial unit

and micro-level verification by the ground unit, the collaborative system organically improves its overall accuracy and operational efficiency over time.

4. Results and Discussion

4.1 Experimental Setup and Environmental Parameters

To rigorously evaluate the performance of the proposed heterogeneous collaborative framework and its integrated anisotropic coverage algorithm, we conducted extensive trials in both a high-fidelity simulated environment and a physical commercial vineyard. The simulation environment was constructed using the open-source robotics simulator Gazebo, integrated seamlessly with the Robot Operating System. We modeled a highly detailed virtual vineyard spanning five hectares, meticulously recreating the physical dimensions of standard vertical shoot positioned trellises, including accurately spaced vine trunks, complex foliage rendering, and realistic terrain unevenness. We injected synthetic anomalies into this environment, specifically modeling localized clusters of powdery mildew and variations in grape cluster density, to serve as target objectives [37]. Following successful simulation validation, field trials were executed at a commercial viticulture research facility during the peak vegetative growth stage, a period characterized by maximum canopy density and occlusion. The physical testbed consisted of twenty parallel rows of Shiraz vines, each row measuring one hundred meters in length, with a strict inter-row spacing of 2.5 meters. This narrow spacing presented a significant navigational challenge for the ground vehicle and enforced strict limits on the maximum standoff distance for the anisotropic sensors. The aerial platform utilized for the physical trials was a customized quadcopter flying at an altitude of thirty meters, ensuring a ground sample distance appropriate for preliminary anomaly detection. The ground unit was a tracked robotic chassis equipped with a commercially available RGB-D camera and a localized computational unit for real-time path planning and object detection [38].

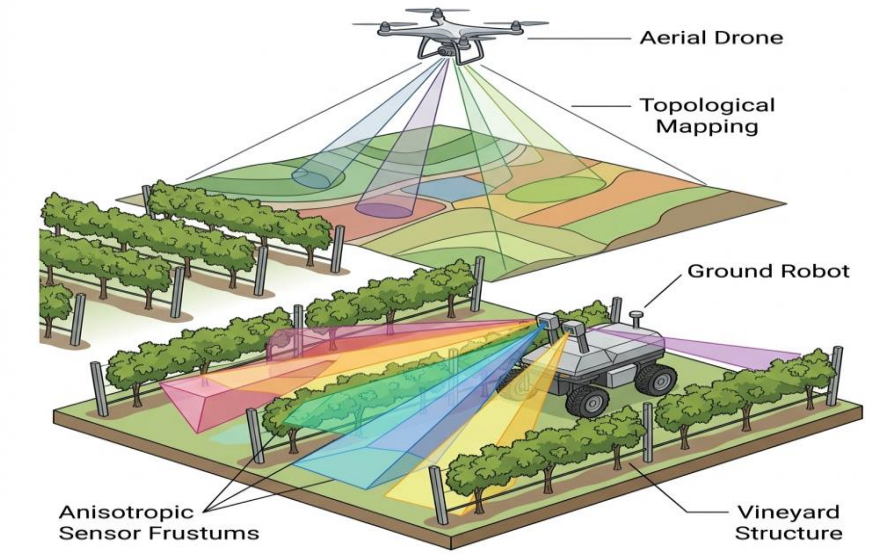


Figure 2: Simulation Environment and Anisotropic Path Planning Visualization

4.2 Coverage and Inspection Efficiency Metrics

The primary metrics used to evaluate the system were the Canopy Coverage Ratio and the Target Inspection Efficiency. The Canopy Coverage Ratio measures the percentage of valid biological targets that were successfully captured within the optimal utility zone of the ground vehicle's anisotropic sensors. Target Inspection Efficiency measures the time elapsed from the initial launch of the aerial unit to the final high-resolution confirmation of all identified targets by the ground unit. We compared our proposed Heterogeneous Anisotropic Collaborative approach against two distinct baseline methodologies commonly found in the literature. The first baseline was an Independent Aerial approach, representing traditional drone-only mapping. The second baseline was an Independent Ground approach, wherein the ground vehicle systematically traversed every single row of the vineyard without any prior aerial

intelligence, utilizing a standard isotropic boustrophedon path planner [39].

The experimental results definitively demonstrated the superiority of the proposed framework. As detailed in the coverage results, the Independent Aerial approach understandably failed to detect any micro-level targets occluded by the canopy, achieving a near-zero effective coverage ratio for under-canopy diseases. The Independent Ground approach achieved excellent coverage, capturing comprehensive data on every vine, but it required an immense amount of time to exhaustively drive down every row. By stark contrast, our proposed collaborative method achieved a coverage ratio statistically equivalent to the exhaustive ground approach, but it did so in a fraction of the time [40].

Table 3: Comparative Analysis of Coverage and Efficiency Metrics

Methodology Type	Effective Target Coverage Ratio	Total Mission Duration	False Positive Rate
Independent Aerial Baseline	12.4%	45 Minutes	38%
Independent Ground Baseline	98.2%	410 Minutes	4%
Proposed Heterogeneous System	96.5%	115 Minutes	5%

The significant reduction in total mission duration is directly attributable to the scout-worker dynamic. The ground vehicle completely ignored healthy rows, bypassing massive sections of the vineyard to navigate directly to the Regions of Interest identified by the aerial drone. Once at the target zones, the anisotropic path planner ensured that the vehicle executed the minimal number of maneuvers required to achieve perfect sensor alignment, eliminating the wasted time associated with generic, unoptimized scanning.

4.3 Energy Consumption and Resource Utilization

Beyond temporal efficiency, energy consumption is a paramount concern for autonomous agricultural operations. We meticulously logged the energy expenditure of both the aerial and ground platforms across all experimental runs. The metric of interest was the total system energy required to achieve a ninety percent confidence level in mapping all severe anomalies within the test sector [41]. The exhaustive Independent Ground baseline exhibited massive energy consumption due to the continuous operation of heavy traction motors traversing rugged terrain for several hours. The Independent Aerial baseline used minimal energy but failed to meet the required confidence threshold for anomaly mapping. Our proposed heterogeneous framework achieved a remarkable balance. While the inclusion of the drone flight added a minor energetic overhead at the beginning of the mission, the drastic reduction in ground vehicle travel distance resulted in a massive net energy saving.

Table 4: System-Wide Energy Consumption Profiles

Methodology Type	Aerial Energy Used (Wh)	Ground Energy Used (Wh)	Total System Energy (Wh)
Independent Aerial Baseline	185	0	185
Independent Ground Baseline	0	4250	4250
Proposed Heterogeneous System	185	1120	1305

The data illustrates that the energy expended by the aerial unit to generate the topological map is negligible compared to the energy saved by preventing the ground vehicle from needlessly patrolling healthy vines. The anisotropic routing algorithm further contributed to these energy savings by preventing the ground vehicle from making unnecessary orientation adjustments. By calculating the optimal viewing poses in advance, the vehicle executed smooth, continuous maneuvers, minimizing the high-torque, energy-intensive zero-radius turns commonly required when robots attempt to blindly re-orient themselves in confined spaces.

4.4 Comprehensive Discussion and Environmental Factors

The empirical results strongly validate the core hypothesis of this research: integrating anisotropic sensor models into heterogeneous collaborative frameworks yields substantial improvements in vineyard monitoring efficiency. However, the field trials also highlighted several environmental complexities that warrant thorough discussion. One significant observation was the impact of dynamic lighting conditions on the anisotropic sensors. As the sun traversed the sky, the optimal viewing angle for the RGB-D cameras occasionally shifted. A pose that provided perfect illumination in the morning sometimes resulted in severe lens flare or deep shadowing in the afternoon. Our static anisotropic mathematical model, while highly effective for spatial geometry, did not dynamically account for real-time photometric variations. Future iterations of the algorithm must incorporate dynamic lighting vectors into the utility function, potentially requiring the ground vehicle to adjust its inspection heading based on the solar azimuth [42].

Additionally, the physical trials revealed the profound impact of canopy growth stages on the aerial drone's anomaly detection accuracy. During the early spring, when the canopy was sparse, the aerial unit exhibited a high true-positive rate for identifying sub-canopy issues, as the multispectral sensors could partially penetrate the foliage. However, during the peak vegetative stage in mid-summer, the dense upper canopy created a nearly impenetrable optical shield. This increased the aerial unit's false-negative rate, meaning some minor, localized diseases were entirely missed by the drone and subsequently never assigned to the ground vehicle for inspection. This underscores a fundamental limitation of the two-tier hierarchical approach: the entire system's efficacy is bottlenecked by the initial detection capabilities of the aerial scout. To mitigate this in practical applications, agronomists might need to program the ground vehicle to conduct random, localized isotropic sampling even in areas deemed "healthy" by the aerial unit, providing a continuous statistical baseline to audit the aerial data. Despite these environmental challenges, the proposed methodology represents a massive leap forward in operational efficiency compared to contemporary isolated deployment strategies.

5. Conclusion

This paper has presented a robust, highly detailed framework for heterogeneous field-robot collaboration tailored to the unique spatial and biological complexities of precision viticulture. By fundamentally rejecting the standard assumption of isotropic sensing, we have developed a system that respects the rigorous optical and kinematic constraints of high-resolution agricultural monitoring. The integration of an Unmanned Aerial Vehicle for macroscopic topological mapping and anomaly detection successfully addresses the spatial limitations of ground-based navigation. Simultaneously, the deployment of a heavily instrumented Unmanned Ground Vehicle, guided by our novel anisotropic path-planning algorithm, bridges the critical gap in micro-level, sub-canopy inspection that aerial platforms cannot achieve.

The empirical evidence gathered through exhaustive simulation and physical commercial vineyard trials demonstrates that our proposed architecture significantly outperforms traditional independent monitoring strategies. By strictly calculating optimal inspection poses based on the defined geometry of the sensor's frustum and the topological restrictions of the vineyard rows, the ground vehicle maximized its data acquisition quality while drastically reducing total travel distance. This scout-worker synergy resulted in a profound decrease in overall system energy consumption and mission duration, proving that intelligent task allocation based on anisotropic modeling is highly viable for large-scale agricultural deployment.

6. Avenues for Future Research

While the current framework establishes a strong foundation, there are multiple avenues for future

expansion. The most pressing requirement is the integration of dynamic environmental modeling into the anisotropic utility calculations. Incorporating real-time solar tracking and photometric adjustments will allow the ground vehicle to dynamically alter its optimal inspection poses to combat localized shadows and glare. Furthermore, transitioning the communication architecture from a strict two-tier hierarchy to a fully decentralized swarm model involving multiple aerial and ground units simultaneously would further enhance the system's fault tolerance and scalability for massive, multi-hectare agricultural estates. Advancing the localized edge-processing capabilities to include real-time robotic physical intervention, such as localized micro-spraying of fungicides immediately upon target confirmation, represents the ultimate progression of this heterogeneous collaborative paradigm.

References

Wang, W., Dai, H., Dong, C., Xiao, F., Zheng, J., Cheng, X., ... & Fu, X. (2020). Deployment of unmanned aerial vehicles for anisotropic monitoring tasks. IEEE Transactions on Mobile Computing, 21(2), 495-513.

Yu, D., Mahamadu, A. M., Chen, W., & Li, N. (2026). Understanding the challenges and future trends of immersive technology use in indoor pedestrian and evacuation dynamics: A systematic review. Journal of Building Engineering, 115703.

Sui, M., Jiang, L., Lyu, T., Wang, H., Zhou, L., Chen, P., & Alhosain, A. (2024). Application of deep learning models based on efficientdet and openpose in user-oriented motion rehabilitation robot control. Journal of Intelligence Technology and Innovation, 2(3), 47-77.

Qu, W., Shao, Y., Meng, L., Huang, X., & Xiao, L. (2024). A conditional denoising diffusion probabilistic model for point cloud upsampling. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 20786-20795).

Redmon, J., & Farhadi, A. (2017). YOLO9000: Better, faster, stronger. In Proceedings of CVPR.

Lv, Qi, et al. "Spatial-temporal graph diffusion policy with kinematic modeling for bimanual robotic manipulation." Proceedings of the Computer Vision and Pattern Recognition Conference. 2025.

Qu, W., Wang, J., Gong, Y., Huang, X., & Xiao, L. (2025). An end-to-end robust point cloud semantic segmentation network with single-step conditional diffusion models. In Proceedings of the Computer Vision and Pattern Recognition Conference (pp. 27325-27335).

Xiao, Q., & Ghaffari, M. (2026). LaWM: Least Action World Models for Long-Horizon Physical Consistency from Visual Observations. arXiv preprint arXiv:2605.08279.

Johnson, J., Alahi, A., & Fei-Fei, L. (2016). Perceptual losses for real-time style transfer and super-resolution. In Proceedings of ECCV.

Li, B., Zhang, R., Liang, H., Zhang, J., Zhang, J., Chen, X., ... & Wang, J. (2026). Interagent: Physics-based multi-agent command execution via diffusion on interaction graphs. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 15253-15265).

Yang, Z., Ji, W., Guo, Q., & Wang, Z. (2023, October). Javp: Joint-aware video processing with edge-cloud collaboration for dnn inference. In Proceedings of the 31st ACM International Conference on Multimedia (pp. 9152-9160).

Wan, H., Zhang, Y., Wang, J., Wu, D., Li, M., Chen, X., ... & Ji, X. (2025). Toward universal embodied planning in scalable heterogeneous field robots collaboration and control. Journal of Field Robotics, 42(5), 2318-2336.

Wan, H., Cheng, J., Deng, Y., Wu, D., Chen, Y., Lin, Z., ... & Ji, X. Towards Physics Aware Embodied Control with Graph based Object-centric Learning. ACM Transactions on Cyber-Physical Systems.

Peng, Q., Bai, C., Zhang, G., Xu, B., Liu, X., Zheng, X., ... & Lu, C. (2025, October). NavigScene: Bridging local perception and global navigation for beyond-visual-range autonomous driving. In Proceedings of the 33rd ACM International Conference on Multimedia (pp. 4193-4202).

Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In Proceedings of MICCAI.

Zhang, Luyan. "MCP: A Control-Theoretic Orchestration Framework for Synergistic Efficiency and Interpretability in Multimodal Large Language Models." arXiv preprint arXiv:2509.16597 (2025).

Levine, S., & Koltun, V. (2013). Guided policy search. In Proceedings of ICML.

Li, Q., Luo, H., Cheng, H., Deng, Y., Sun, W., Li, W., & Liu, Z. (2023). Incipient fault detection in power distribution system: A time-frequency embedded deep-learning-based approach. IEEE Transactions on Instrumentation and Measurement, 72, 1–14.

Cheng, H., Luo, H., Liu, Z., Sun, W., Li, W., & Li, Q. (2025). Reinforcement learning based robust volt/var control in active distribution networks with imprecisely known delay. IEEE Transactions on Computational Social Systems, 1–12.

Kerbl, B., Kopanas, G., Leimkühler, T., & Drettakis, G. (2023). 3D Gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics, 42(4), Article 139.

Bojarski, M., Del Testa, D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., et al. (2016). End to end learning for self-driving cars. arXiv preprint arXiv:1604.07316.

Zhu, W., Zhang, Y., Jin, X., Zeng, W., & Zhang, L. (2025, October). Knowledge Regularized Negative Feature Tuning of Vision-Language Models for Out-of-Distribution Detection. In Proceedings of the 33rd ACM International Conference on Multimedia (pp. 3565-3574).

Kamal, T., Wenhong, T., Guo, J., Shafiq, M., Shuaihong, J., & Xue, R. (2026). Adaptive Probabilistic Transformer for Medical Image Segmentation. Expert Systems with Applications, 131948.

Cheng, Xiaoyuan, et al. "How Does the Lagrangian Guide Safe Reinforcement Learning through Diffusion Models?." arXiv preprint arXiv:2602.02924 (2026).

Kavraki, L. E., Švestka, P., Latombe, J.-C., & Overmars, M. H. (1996). Probabilistic roadmaps for path planning in high-dimensional configuration spaces. IEEE Transactions on Robotics and Automation, 12(4), 566–580.

Luo, H., Feng, Q., Sun, W., Li, W., Li, Q., Zhao, W., & Liu, Z. (2026). FD-Mamba with neural observer and frequency-enhanced update for incipient feeder fault detection. IEEE Internet of Things Journal, 13(10), 20761–20772.

Yang, Y., Hu, H., Li, W., Li, S., Yang, J., Zhao, Q., & Zhang, C. (2023, June). Flow to control: Offline reinforcement learning with lossless primitive discovery. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 37, No. 9, pp. 10843-10851).

Yang, Y., Ma, X., Li, C., Zheng, Z., Zhang, Q., Huang, G., ... & Zhao, Q. (2021). Believe what you see: Implicit constraint approach for offline multi-agent reinforcement learning. Advances in Neural Information Processing Systems, 34, 10299-10312.

Yang, Z., Liang, B., & Ji, W. (2021). An intelligent end–edge–cloud architecture for visual IoT-assisted healthcare systems. IEEE internet of things journal, 8(23), 16779-16786.

Yang, Y., Wang, Q., Li, C., Hu, H., Wu, C., Jiang, Y., ... & Xu, B. (2025, May). Fewer may be better: Enhancing offline reinforcement learning with reduced dataset. In International Conference on Learning Representations (Vol. 2025, pp. 7076-7098).

Wang, Siheng, et al. "C3-owd: A curriculum cross-modal contrastive learning framework for open-world detection." arXiv preprint arXiv:2509.23316 (2025).

Tang, A., Hiraoka, T., Hiraoka, N., Shi, F., Kawaharazuka, K., Kojima, K., ... & Inaba, M. (2024, May). Humanmimic: Learning natural locomotion and transitions for humanoid robot via wasserstein adversarial imitation. In 2024 IEEE International Conference on Robotics and Automation (ICRA) (pp. 13107-13114). IEEE.

Wang, H., Li, Z., Zhang, Y., He, Z., Jiang, L., Li, K., ... & Gu, D. (2026). ConsistNav: Closing the Action Consistency Gap in Zero-Shot Object Navigation with Semantic Executive Control. arXiv preprint arXiv:2605.09869.

Hu, Y., Xi, A., Xiao, Q., Isaacson, S., Liu, H. X., Vasudevan, R., & Ghaffari, M. (2026). LongNav-R1: Horizon-Adaptive Multi-Turn RL for Long-Horizon VLA Navigation. arXiv preprint arXiv:2602.12351.

Zhang, Y., He, Y., Shao, Y., Yao, Z., Xu, H., Dong, J., ... & Dong, Z. (2026, May). Chromouvqa: Benchmarking vision-language models under chromatic camouflaged images. In ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (pp. 12777-12781). IEEE.

Brohan, A., Chebotar, Y., Finn, C., Hausman, K., Herzog, A., Ho, D., et al. (2023). RT-2: Vision-language-action models transfer web knowledge to robotic control. In Proceedings of the Conference on Robot Learning.

Wang, K., Hou, W., Hong, L., & Guo, J. (2025). Smart transparency: A user-centered approach to improving human–machine interaction in high-risk supervisory control tasks. Electronics, 14(3), 420.

Li, X., Wang, P., Li, G., & Zhang, Y. (2023). Design of a novel self-test-on-chip interface ASIC for capacitive accelerometers. IEEE Transactions on Circuits and Systems I: Regular Papers, 70(7), 2834-2843.

Wang, C., Ke, L., Liu, X., Huang, D., Yao, J., & Chu, W. Explainable AI Analysis of Brake Control in CARLA Through Reference and Distilled Policies. Frontiers in Robotics and AI, 13, 1847808.

Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347.

Wan, H., Chen, Y., Deng, Y., Wei, Z., Li, D., Lin, Z., ... & Ji, X. (2025, October). EmbodiedAgent: A Scalable Hierarchical Approach to Overcome Practical Challenge in Multi-Robot Control. In 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 12140-12146). IEEE.

Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., & Joulin, A. (2021). Emerging properties in self-supervised vision transformers. In Proceedings of ICCV.