RESEARCH Open Access

A self organizing cyber physical system for sustainable agriculture

Michael Batistatos[1,2*], Tomaso De Cola[3], Diogo Oliveira[5], Sergio Aguilar[5], Aubrey Dunne[4], Antonio Ion Mazilu[1], Michail-Alexandros Kourtis[2], George Xilouris[2] and Nikolaos Sagias[1]

*Correspondence:
Michael Batistatos
mbatist@uop.gr
[1]Informatics and Telecommunications, University of the Peloponnese, Tripolis, Greece
[2]Institute of Informatics and Telecommunications, National Centre for Scientific Research"DE MOKRITOS"(NCSRD), Athens, Greece
[3]German Aerospace Center (DLR), Wessling, Germany
[4]Ubotica Technologies Ltd, Dublin, Ireland
[5]Sateliot, Barcelona, Spain

Abstract
The transformation of food production through digital tools and technologies requires not only advanced devices and processing capabilities, but also a holistic agricultural framework that unifies multiple types of stakeholders across the agri-food chain. This paper introduces the AGRARIAN platform, a stakeholder-centric, self-organizing Cyber-Physical System designed to answer the modern and future demands of sustainable agriculture, unifying farmers, developers and policy makers, all in one dynamic and scalable ecosystem. Integrating hybrid communications, lightweight services orchestration, continuous development and cross-farm decision support, the AGRARIAN platform offers a complete solution for the next generation of farm management. For the platform validation, field trials were performed for the use case of livestock monitoring in a remote area, using commercial communication links (5 G, WiFi and satellite) for real-world system evaluation. The measurements and the results analysis confirmed the capacity of the AGRARIAN platform to support the necessary functions and services, opening the way towards modern and sustainable farming.

Keywords Cyber-physical systems (CPS), Smart agriculture, Internet of things (IoT), Edge-oriented orchestration

1 Introduction

The agricultural sector is undergoing a fundamental transformation towards more sustainable, efficient and resilient production systems. Especially with the transition from legacy forms of food production (e.g. Agriculture 4.0) to modern ones (i.e. Agriculture 5.0 [1]), digital technologies such as sensors, robotics, digital twins, computing at the edge and modern communications promote automation, data-driven decision-making and system-wide optimization by embedding intelligence throughout the agricultural chain. However, implementing such solutions in farms that are mostly located in remote/rural areas introduces significant challenges. These include intermittent or low-bandwidth connectivity, limited computing resources, lack of infrastructure on the field and farmers with no (or minimal) technical expertise. In this context, conventional Internet-of-Things (IoT) systems fall short, particularly when real-time responsiveness and autonomous operation are required under constrained and harsh conditions. Most importantly, beyond the new technologies, there is a growing need for advancing the methodology and practices on how agricultural data is collected, processed and shared. Enabling more holistic, cross-farm insights is essential for optimized and efficient decision-making at the policy level, particularly in supporting clusters of farms with common environmental and economic conditions.

AGRARIAN platform, presented in this paper, tries to answer the aforementioned challenges and proposes a solution for enabling both the smooth adoption of digital tools in future farming and the integration of third parties to the agricultural ecosystem. The proposed platform, acting as a self-organizing cyber-physical system (CPS) [2], integrates edge intelligence, hybrid connectivity and Continuous Integration/Continuous Delivery (CI/CD) practices [3] to support sustainable and scalable agricultural applications, especially in rural and infrastructure-limited agricultural environments. The developed platform provides a programmable, container-based environment for deploying and managing applications across edge, fog and cloud layers. Using a lightweight containerized orchestration framework and the CI/CD pipeline, it supports continuous delivery and autonomous updates of services even under limited network connectivity. Its hybrid communication stack, integrating terrestrial (5 G, Wi-Fi, etc.) and non-terrestrial (satellites) communications, ensures resilient connectivity and fault-tolerant operation. Its core strength lies in its ability to autonomously handle the data from the IoT devices on the field, the users' requests, the applications operation and the computational resources, emphasizing on the edge and real time service deployment.

AGRARIAN, identifying the rising need for multiple actors collaboration and interaction into a modern agricultural ecosystem, delivers an integrated platform designed with three primary stakeholder groups in mind: developers, policy makers and farmers. For developers, it provides a programmable, modular environment to deploy and test distributed applications. For policy makers, it offers a replicable reference architecture that aligns with the goals for digital rural development and sustainability, while for farmers, it ensures smooth platform access and service availability even in constrained conditions.

To validate the platform, a representative field trial, acting as a proof-of-concept, was conducted in a livestock monitoring scenario. The trial included autonomous deployment of a requested AI application at the edge of the system, a flying Unmanned Aerial Vehicle (UAV) streaming live video through multiple communication links and performance validation based on metrics of the inference model. The results and overall field demonstration underscore the AGRARIAN platform ability to support reliable CPS services in the agricultural sector and push farming towards the new era.

The core innovation of AGRARIAN lies in its mission to bring technology to the direct service of agriculture, introducing a holistic multilevel approach for main actors' interaction: AGRARIAN not only integrates lightweight orchestration (K3s), hybrid communication fallback and cross-farm decision support systems, but it also builds an ecosystem that brings together farmers, developers and policy makers into a shared operational and innovation space. This ecosystem encourages the continuous development and deployment of digital agricultural solutions through an accessible, dynamic and scalable framework, offering real-world usability and resilience, supporting service continuity even in the most infrastructure-limited areas.

The rest of the paper is structured as follows: Chapter 2 presents related work in CPS for agriculture, edge computing and DevOps for distributed IoT Systems, hybrid communications and existing agricultural platforms, chapter 3 describes the AGRARIAN system architecture, chapter 4 discusses the platform deployment and operational workflow, chapter 5 analyses the field trials, chapter 6 discusses the scalability and economic viability of AGRARIAN platform and chapter 7 concludes the paper summarizing the presented work and discussing future directions.

2 Related work

2.1 Cyber-physical systems in agriculture and smart farming

CPS have been increasingly gaining ground in the agricultural sector in recent years, trying to offer efficient and smart practices in modern farming. Authors in [4] discuss the overarching concept of Cyber-Agricultural Systems (CAS) that utilize ubiquitous sensing, artificial intelligence (AI) and scalable cyber-infrastructure for improving both crop breeding and production agriculture. Key elements like sensing, modeling and actuation act as fundamental components of CAS, while tools like digital twins are emerging technologies in future farming. The authors highlight also the importance of CAS on revolutionizing the agricultural efficiency and sustainability by implementing data-driven decision support systems directly into the agricultural processes, achieving real-time remote management and simulation-driven planning.

The work in [5] discusses the application of digital twins technology in agriculture as a form of CPS. The researchers identify a digital twin as a virtual replica of a farm physical processes, which mirrors the real-time states and behaviors. They develop a taxonomy of different digital twin types and propose a reference framework for implementing them in agriculture, based on general systems theory and IoT architecture principles. This framework has been validated in five use cases (arable farming, dairy, greenhouse, vegetable and livestock farming), demonstrating that digital twins enable remote monitoring and control of farm operations with near real-time data. Thus, enhancing decision making and responsiveness in farming. Their work shows that CPS concepts like digital twins can decouple on-farm physical processes from their management, allowing farmers to “manage by exception” and simulate interventions before applying them in the field.

Data analytics and IoT integration are also basic elements in agricultural CPS. Authors in [6] provide a comprehensive survey of smart farming from a data-centric perspective, discussing the types of big data generated in agriculture, key applications and techniques for analysis. They emphasize how CPS in agriculture should handle diverse data (soil, weather, crop health, etc.) and convert it into actionable intelligence. The authors categorize applications of agricultural CPS, such as precision crop monitoring, automated irrigation and yield prediction, while reviewing techniques ranging from machine learning, for crop modeling, to blockchain, for the transparency and security of the supply chains. They emphasize on how modern smart farming solutions rely on a convergence of technologies, like sensing, communications and AI, in order to support closed-loop control in farming, reflecting core CPS principles of interconnectedness and feedback.

Several works address the overall contribution of IoT-based CPS platforms to agriculture. The research in [7] surveys emerging IoT technologies for smart agriculture in the context of Industry 4.0. It details how modern farms employ multi-layer CPS architectures: low-power sensor networks and drones gather field data, nodes at the edge perform local processing and cloud platforms provide global optimization and big data analytics. The authors highlight the importance of interoperability and defined standards, noting initiatives that integrate IoT with cloud computing, machine learning and even software-defined networking for agriculture. They also discuss challenges, such as security, data management and the importance of real-time analytics in agricultural CPS. By reviewing over 250 recent works, they justify that IoT-driven CPS are becoming the backbone of sustainable agriculture, enabling use cases varying from smart greenhouses to AI-based pest detection with improved productivity and efficiency.

In resource-constrained farming conditions, CPS approaches must be particularly cost-effective and resilient. Authors in [8] present the challenges and opportunities of implementing smart agriculture CPS in developing regions with limited resources. They propose a holistic framework (based on a specific tool for assessing and aligning internal organizational elements) that addresses not only the hard elements of technology, such as IoT sensors, hybrid network communications and edge data management, but also the soft elements like farmer training and supportive policies. They case study a low-cost, open-source IoT system deployed in a greenhouse using an edge system architecture for climate control. The results demonstrate that even with limited hardware and applying CPS principles, such as sensing, automated decision and feedback control, the water consumption and the crop yield were significantly improved. This indicates how CPS platforms, when are carefully designed, can make agriculture smarter and more climate-resilient, even in rural and under-developed areas, as long as they are adapted to local constraints and coupled with capacity-building efforts.

Beyond farm management, CPS in agriculture also extends to supply chains and broader Industry 4.0 integration. Authors in [9] introduce a conceptual Cyber-Physical Agricultural System (CPAS) framework aimed at integrating farm operations with the supply chains and food processing networks. Their framework intelligently combines on-farm CPS (sensors, drones and automation) with cloud services and big data analytics to enable end-to-end visibility from "farm to fork." The authors argue that such integration can enhance productivity and sustainability. For example, crop status data produced by farm CPS can be consumed by logistics and market forecasting systems in real time. They also discuss implementing Industry 4.0 technologies (IoT, cloud computing, big data) in agriculture for predictive analytics and resource optimization. The CPAS concept presented by the authors demonstrates how future farming can operate not in an isolated environment, but rather as part of a connected cyber-physical value chain, where decisions in cultivation, harvesting, distribution and even retail are coordinated through shared data and intelligent systems.

2.2 Edge computing and DevOps for distributed IoT systems

As IoT systems scale across remote and heterogeneous environments, edge computing has emerged as critical enabler for low-latency data processing, reduced bandwidth usage and localized decision-making. The work in [10] provides a comprehensive survey of edge computing frameworks, highlighting how proximity to data sources enables real-time responsiveness in applications such as smart farming and autonomous monitoring. On the same page, authors in [11] highlight fog computing as a broader paradigm, incorporating intermediary layers of computation between cloud and devices to improve scalability and resilience.

Authors in [12] extend this view by analyzing the integration of AI into the edge, forming the basis of what is termed the Artificial Intelligence of Things (AIoT). They emphasize the role of federated learning and distributed inference in enabling intelligent behavior without reliance on continuous cloud connectivity, capabilities particularly vital in agriculture. Researchers in [13] further reinforce that a hybrid cloud, fog and dge approach is essential for maintaining service availability and robustness in real-world IoT deployments.

Building on these computing models, recent research has focused on how edge services can be developed, deployed and maintained reliably, applying Development Operations (DevOps) [14] and CI/CD principles to distributed environments. Authors in [15] introduce GeneSIS, a model-driven deployment system that brings security and automation to IoT orchestration, while authors in [16] explore GitOps-based continuous deployment in edge computing, using Kubernetes [17] and declarative configuration to streamline service updates across diverse hardware.

Platforms like Sophon Edge [18] and DeepRIoT [19] demonstrate how containerized microservices, CI/CD pipelines and automated model updates can bring cloud-grade agility to IoT and robotics applications. These systems validate that with specialized implementation and orchestration, modern software engineering practices can be extended to constrained nodes at the edge of the network.

2.3 Hybrid communication in remote IoT deployments

Reliable and seamless communication is an important challenge in remote areas, where mostly the farms exist, where terrestrial cellular coverage is inconsistent or entirely absent. Hybrid communication systems, integrating both terrestrial (e.g., NB-IoT, Wi-Fi, LoRa) and non-terrestrial (e.g., LEO/GEO satellite) technologies, can offer an alternative solution for ubiquitous connectivity is such areas, acting as essential enablers for resilient and scalable agricultural CPS platforms.

Authors in [20] highlight the use of satellite IoT to extend connectivity in remote farms, proposing systems that combine local LoRaWAN networks with LEO satellites for data backhauling. This dual-layer architecture ensures that data from field sensors can be aggregated locally and transmitted globally when needed. Researchers in [21] examine the evolution of NB-IoT for satellite communication, as formalized in 3GPP Release 17 and they identify optimizations such as time repetition and Doppler resilience that make NB-IoT suitable for low-power and delay-tolerant rural applications.

The work in [22] focuses on Non-Terrestrial Networks (NTNs) as a foundational part of 6 G communication systems. Their relevance to agriculture lies in ensuring global coverage, especially when combined with UAVs or High Altitude Platform Systems (HAPS), to deliver latency-sensitive services such as precision livestock tracking or irrigation control.

Authors in [23] propose a multi-tier architecture for rural deployments, using NB-IoT systems for connectivity on the field, UAVs for edge data relaying and satellite links as cloud backhaul. Simulations demonstrate that dynamic selection of communication tiers based on energy, bandwidth and signal strength, dramatically improves service uptime in sparsely covered zones.

Case studies in hybrid deployments, such as LoRa for local sensing and satellite fallback for reliable data reporting, show that hybrid connectivity is already used in precision irrigation, soil monitoring and environmental sensing [24]. These systems allow local data collection to continue uninterrupted even when remote connectivity is unavailable, enabling queued transmission when link quality improves.

2.4 Existing agricultural platforms

In recent years, several EU-funded projects and open platforms have aimed to digitally transform agriculture by offering interoperable services, data sharing frameworks and digital tools integration. Notable among them are:

SmartAgriHubs [25], was a major EU-funded project (2018–2022) aimed at accelerating digital innovation in European agri-food sector by building a wide network of Digital Innovation Hubs (DIHs). Rather than developing new IoT devices or software platforms itself, SmartAgriHubs innovation lay in its network-building. It improved the innovation services of DIHs and helped align regional funding (including structural funds) to support agri-tech innovation. The project facilitated 28 flagship innovation experiments across 9 regional clusters, demonstrating how digital solutions (IoT, big data, robotics, etc.) can be adopted in farming with support from local DIHs.

DEMETER [26] was a large-scale EU project (2019–2023) focused on deploying interoperable smart farming IoT platforms that are farmer-driven. Its main objective was to lead the digital transformation of agri-food sector in Europe through the widespread adoption of IoT, data science and smart farming technologies. The key feature of DEMETER was its emphasis on interoperability as the main enabler, integrating data and services across different farm management systems, sensor networks, machinery, and software platforms. It established a common "Interoperability Space" for agriculture, built on standard information models and open APIs, to break down data silos between proprietary systems.

FIWARE [27] technology offers a general-purpose open-source framework for building smart applications in various domains, including agriculture. Its architecture supports semantic data models and API standards, aiming for interoperability and scalability. Its approach brings standardization and interoperability to farm software. It defines common smart data models for agricultural concepts (fields, crops, machinery, etc.), so that different applications can exchange information easily. As it is standard-based, it promotes an "app store" or marketplace model where third-party smart farming solutions can plug into any FIWARE-based system without custom integration. FIWARE project doesn't itself provide farm sensors or connectivity, but it ensures that as data is collected (from any IoT device or software), it can be published, shared or monetized through a uniform interface.

OpenAgri [28] is a new ongoing EU Horizon Europe project (2024–2026) that aims to enhance modern farming with open-source digital tools. It focuses on practical accessibility and co-creation: developing free, open-source farming applications that can run even with limited internet connectivity and co-designing these tools with farmers themselves. This approach is a shift from traditional top-down tech deployment, as farmers and developers are equal partners in designing solutions, which builds trust and ensures real needs are met. A key outcome of OpenAgri will be the creation of an open-source ecosystem for digital farming tools, effectively lowering barriers of cost and proprietary lock-in.

NOSTRADAMUS [29] is a new and ongoing EU project (2024–2028) focused on leveraging big data and Earth observation for European food security. Its vision is to create a scalable platform for real-time data harvesting and analysis that provides farmers and decision-makers with actionable insights for sustainable agriculture. A distinguishing technology aspect of NOSTRADAMUS is the integration of Copernicus Earth Observation data (satellite imagery) with on-the-ground IoT sensor data, combined through advanced algorithms (AI/ML) in a unified “agricultural data cube” or data platform. The project emphasizes open-source applications and a multi-actor approach, meaning the tools and models developed will be openly available and designed in collaboration with end-users (farmers, agri-cooperatives, public agencies).

2.5 Research gaps and AGRARIAN platform innovation

While prior research has advanced the individual domains of CPS, edge computing, hybrid communication architectures and CI/CD methodologies, these contributions often remain isolated, narrowly scoped and in many cases confined to laboratory setups or simulated environments. A persistent gap exists in delivering an integrated, modular platform that serves not only end users, but also developers and policy makers, who require scalable, configurable and maintainable IoT infrastructures tailored to the complexities of agriculture.

Existing platforms such as DEMETER, FIWARE for Agriculture and SmartAgriHubs have made important strides in standardization, interoperability and data integration. However, they primarily operate through centralized or cloud-based models and often lack autonomous service orchestration at the edge, real-time fallback mechanisms across heterogeneous networks (e.g., 5 G, Wi-Fi, satellite) or seamless CI/CD pipelines for deploying new applications in infrastructure-limited settings.

The AGRARIAN platform presented in this paper addresses these limitations by offering a self-organizing, edge-native architecture that unifies (i) lightweight application orchestration (via K3s), (ii) dynamic connectivity fallback (via SDN-managed hybrid communication), (iii) continuous deployment pipelines for developer-driven innovation and (iv) a cross-farm support decision system for high-level policy guidance. Furthermore, AGRARIAN places a strong emphasis on building an ecosystem that brings farmers, developers and policy makers together in a shared space for co-creation, configuration and innovation. Although field validation is ongoing, AGRARIAN is explicitly engineered to bridge the gap between controlled environments and real-world agricultural deployments, setting the foundation for resilient and scalable smart farming solutions in clusters of rural and resource-constrained areas. Its open, modular approach encourages community participation and long-term sustainability beyond the lifespan of the project.

It is noted that the last two aforementioned projects, OpenAgri and NOSTRADAMUS, are part of the same European Commission project cluster as AGRARIAN, focusing on advancing digital transformation in agriculture through complementary approaches. This shared ecosystem fosters collaboration, knowledge exchange and alignment of objectives across projects, encouraging interoperability and mutual reinforcement of outcomes. By positioning AGRARIAN within this cluster, the platform not only benefits from synergies with other initiatives, but also distinguishes itself through its focus on edge-native processing, autonomous application deployment and hybrid communication infrastructure. This coordination within the cluster supports a broader vision for scalable, inclusive and policy-relevant agricultural innovation across Europe, while highlighting AGRARIAN unique contributions to infrastructure resilience and multistakeholder engagement.

3 System architecture

The conceptual architecture of the AGRARIAN platform, illustrated in Fig. 1, depicts the interaction between the key system actors and components across the data lifecycle. End users, such as farmers, access the applications via a web portal, while developers contribute and manage AI-driven services through a development dashboard. These applications are deployed over a hybrid infrastructure, comprising IoT sensors, UAVs, edge computing units and cloud services, coordinated by an orchestration and data processing layer. The platform leverages seamless hybrid communications, including terrestrial (e.g. 5 G) and non-terrestrial (e.g. satellites) links, for resilient data transmission and storing of selected data for both immediate feedback and long-term analysis. Selected processed data are consumed by the Agricultural Decision and Support System (ADSS), which combines internal insights with external data sources to generate actionable information for policymakers. Enabling this unified workflow and CPS principles within the agricultural sector, the aim of the AGRARIAN system architecture is to effectively bridge the localized sensing with strategic agricultural planning.

In the context of this paper, we distinguish between edge, fog and cloud computing components based on their physical location and functional role. Edge refers to computing nodes deployed close to the data sources (e.g., on-farm servers) that handle real-time data processing and inference. Fog represents intermediate infrastructure (e.g., regional aggregators or gateways) that can perform localized coordination and buffering. Cloud indicates centralized computing resources used for non-real-time processing, long-term storage or advanced analytics. Throughout the paper, we use these terms consistently to reflect their respective roles within the system architecture.

A more detailed high level system architecture is depicted in Fig. 2. The core of the system is the orchestration mechanism, built on lightweight Kubernetes (K3s) [30] and

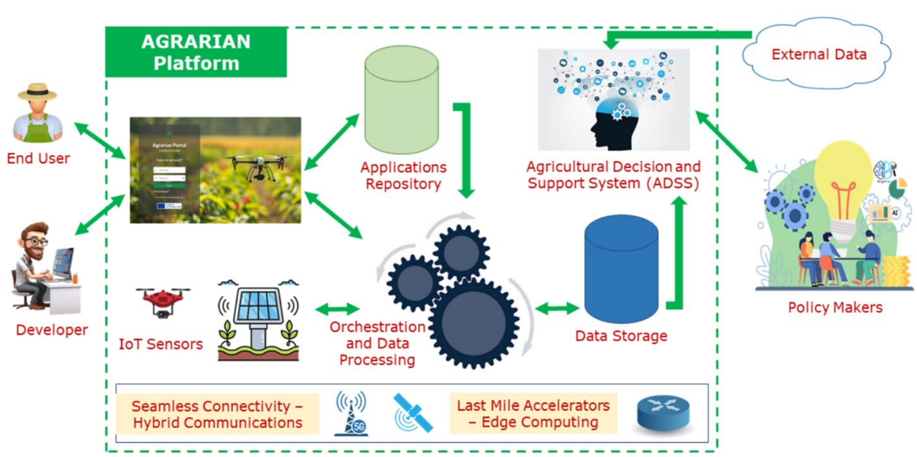

Fig. 1 A high-level conceptual illustration of the AGRARIAN platform, highlighting its core vision of unifying stakeholders and services within a dynamic, distributed agricultural ecosystem

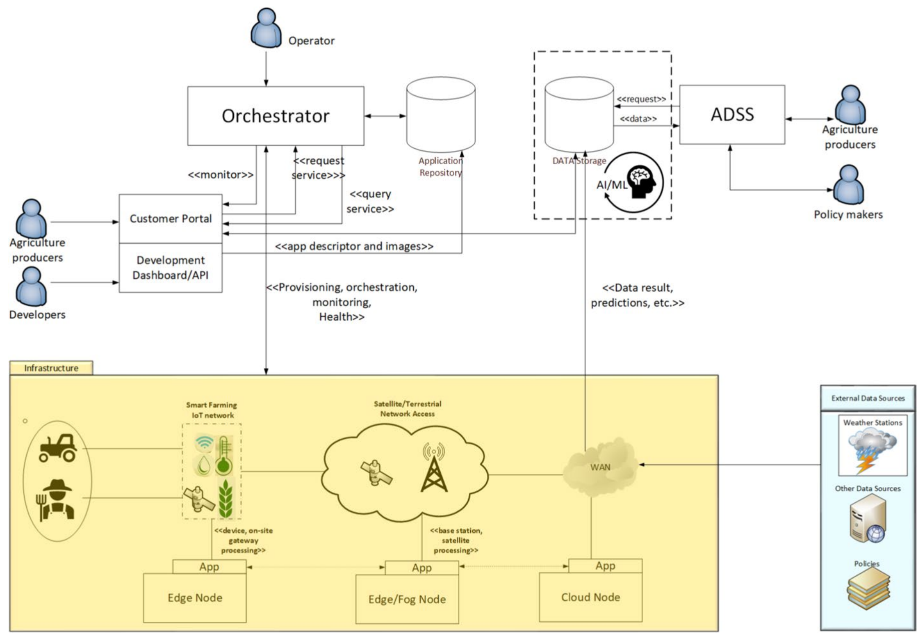

Fig. 2 A detailed architecture diagram of the AGRARIAN platform, showing the integration of edge nodes, orchestrator, hybrid communication layers, data pipelines and the Agricultural Decision Support System (ADSS)

driven by GitOps-based CI/CD workflows [31]. This orchestration layer ensures that containerized services and AI applications can be listed, deployed and updated dynamically across distributed nodes at the edge of the network. Developers interact with this layer using version-controlled repositories and declarative configuration files, enabling reproducibility and automation across the platform.

A hybrid communication infrastructure layer ensures the seamless connectivity across the CPS loop. AGRARIAN integrates terrestrial (5 G, NB-IoT, Wi-Fi) and non-terrestrial (GEO-LEO satellite) networks into a resilient communication fabric. This multi-modal connectivity ensures that edge devices and gateways remain connected even in remote regions with intermittent or no cellular coverage. The system dynamically selects communication paths based on availability and energy constraints, supporting seamless data flow between physical assets and cloud-based services. This capability is crucial for achieving persistent feedback and actuation in the CPS model, even under highly variable network conditions.

Layered on top of this resilient infrastructure is the AGRARIAN distributed data intelligence layer, where sensor inputs are locally analyzed by edge-deployed applications. These applications, onboarded by developers, use AI/ML models or other algorithms to interpret field data and provide real-time feedback to farmers, such as livestock detection, plant diseases or microclimate monitoring. The derived results are transmitted upstream and selectively stored in the AGRARIAN data repository. At this level, the ADSS consumes accumulated insights from across multiple farms, enabling cluster-level analysis that informs regional agricultural policy. This two-tier data allows AGRARIAN to bridge the CPS feedback loop from field-level operations to high-level strategy.

Completing the architecture is the user access and interaction layer. This includes secure, role-specific interfaces for all stakeholders. Farmers access a simplified operational portal for accessing the list of available applications, selecting them and then receiving the service. Developers manage deployments through a Git-integrated dashboard, while policymakers visualize large-scale agricultural patterns via the ADSS interface. Authentication, authorization and data integrity are enforced throughout this layer to maintain trust and system transparency.

4 Platform deployment and operational workflow

AGRARIAN platform is designed to streamline the delivery and management of distributed services, ensuring real-time responsiveness, autonomy and resilience through self-organizing orchestration mechanisms. This chapter outlines the operational workflow that bridges the architectural foundation of the system with its practical field deployment, highlighting the developer-to-field application cycle and the interfaces that enable edge intelligence, including satellite-based AI capabilities.

Figure 3 depicts the detailed operational workflow. The deployment process begins when the developer uploads a new application via the AGRARIAN development dashboard. This action triggers a CI/CD pipeline that automates testing, containerization of the application and pushes it to the AGRARIAN application repository, where it is stored as a containerized workload, accompanied by configuration files with deployment specifications.

AGRARIAN employs the GitOps workflow to synchronize the desired system state with the live deployment across the edge. The K3s periodically polls the repository for changes and when a new application or an update is detected, the orchestration layer automatically pulls the relevant container image and configuration, deploying it to the target node. The orchestration method includes resource-aware scheduling, health checks and rollback mechanisms, enabling self-healing and adaptive reconfiguration.

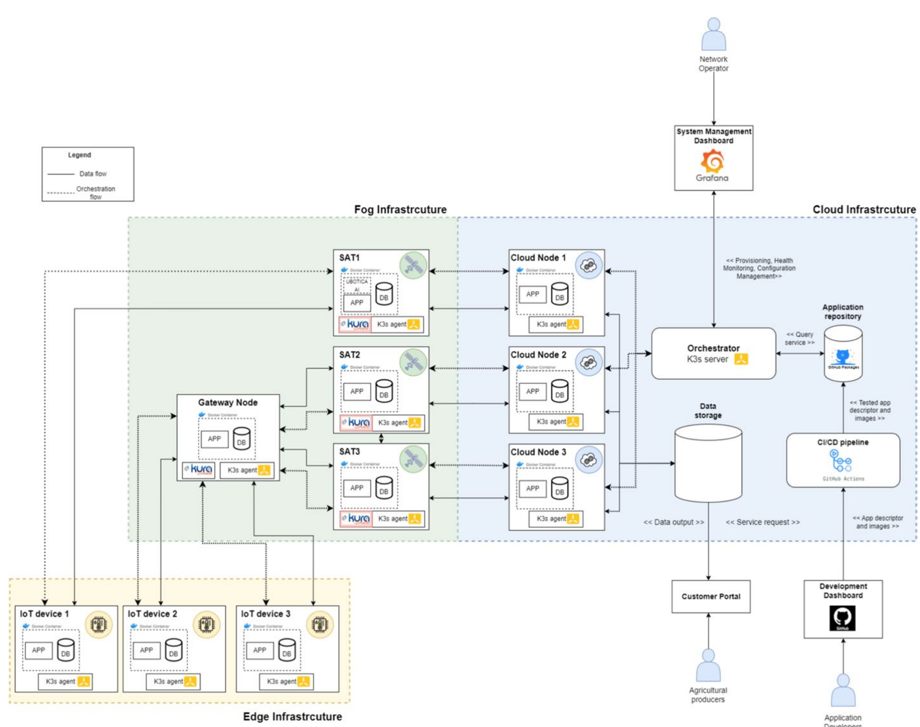

Fig. 3 A step-by-step depiction of the service request and deployment process, from end-user interaction via the AGRARIAN portal to containerized application orchestration, execution and data feedback

When a farmer requests the installation of an application, such as an AI model for livestock detection, via the user portal, the system maps the request to a compatible deployment profile and initiates a lightweight remote installation. Based on node availability and application-specific requirements (e.g. processing needs, delay-tolerance) the system selects the most suitable node across the distributed edge/fog/cloud infrastructure to meet the application performance requirements.

An integral component of AGRARIAN edge intelligence is its support for satellite-based AI processing [32]. The platform integrates satellite edge nodes as part of its fog infrastructure, utilizing them as deployable devices for selected applications. Using containerized frameworks managed by K3s, these satellite nodes run both persistent infrastructure services and application-specific AI containers. Communication between containers for hardware control is handled using the Remote Procedure Calls gRPC [33] framework, while AI models in the Open Neural Network Exchange (ONNX) format [34] are optimized and compiled with the OpenVINO open source toolkit [35] prior to deployment. Eclipse Kura [36], an open-source IoT edge framework is adopted to enable edge device programmability on terrestrial IoT gateways and is being evaluated for satellite-based deployments. It facilitates data aggregation from field sensors and supports the integration of custom functions over the network stack.

Throughout the deployment process, system telemetry and logs are collected and visualized through monitoring dashboards accessible to developers and system operators. This observability supports rapid troubleshooting, performance analysis and application tuning. The continuous interaction between Git-based development workflows, automated orchestration and adaptive edge deployment sets the foundation for robust, dynamic, scalable and intelligent agricultural systems that are not only developer-friendly, but also ready for operational use in demanding field conditions.

5 Field trials and system validation

To validate the operational viability of the AGRARIAN platform in real-world conditions and in rural areas, a field trial was conducted, centered on a livestock monitoring use case. The objective was to assess the system ability to deploy, execute and sustain AI-driven applications in environments characterized by variable or limited connectivity. In this scenario, a farmer utilized the AGRARIAN web portal to request and deploy a sheep detection application, which was installed on an edge node connected to a UAV used for aerial observation. The UAV captured real-time video of grazing livestock, which was then processed locally by an AI model hosted on the edge node. This trial provided a practical demonstration of how AGRARIAN enables seamless application delivery, edge intelligence and hybrid communication support for mission-critical agricultural tasks.

5.1 Use case deployment

In addition, the livestock monitoring trial was also designed to demonstrate how a non-technical end user, such as a farmer, could request, deploy and benefit from an AI-powered application without needing to interact directly with the technical operation of the system. Figure 4 shows the field trial setup. The farmer began by accessing the AGRARIAN portal via a secure login interface. Upon authentication, the farmer was presented with a curated list of available applications, each described by its function,

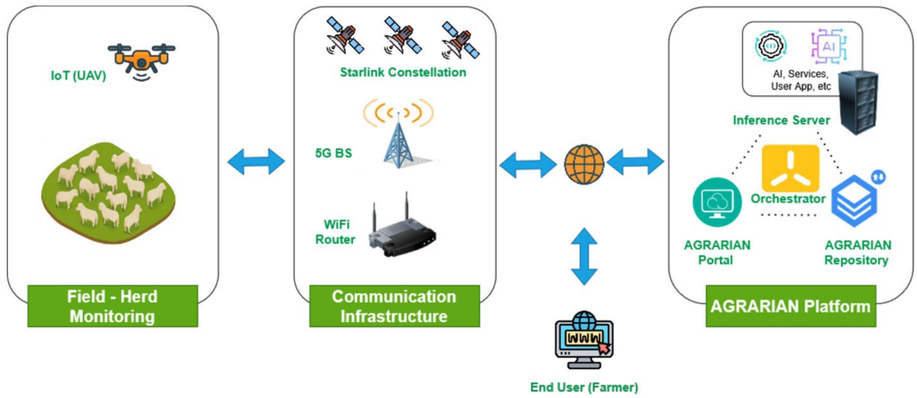

Fig. 4 Schematic of the field trial testbed setup used for validating the AGRARIAN platform, including UAV, edge node, hybrid communication backbone (5 G, Wi-Fi, Ethernet, Starlink) and cloud integration

data requirements, and deployment compatibility. From this list, the farmer selected the "Sheep Detection and Counting" application.

Once the application was selected, the platform orchestration layer automatically processed the request and deployed the application to the nearest available edge node. This process was managed through the GitOps pipeline and K3s, confirming that the system could seamlessly translate user input into an operational edge deployment. The deployment included the AI model for object detection, a lightweight inference engine and a local web server for exposing results. Only after the successful completion of this step did the UAV begin streaming video data. All subsequent performance measurements, including inference confidence, latency and frame loss were measured at the application server hosting the AI model.

The core sensing component of this use case was a high-resolution camera, onboard the UAV, functioning as a mobile IoT sensor. The UAV was manually piloted over the grazing area, capturing continuous video of the sheep herd. The video stream was transmitted to the deployed application node using the best available local connectivity, which could include Wi-Fi, 5 G or fallback to satellite if necessary. The edge node processed the incoming video frames in real time, executing the sheep detection model to identify and count the animals present in each frame.

The inference results, bounding boxes around detected sheep and a running count, were rendered into a web application hosted on the edge node itself. The application exposed a secure web endpoint that the farmer could access using any standard browser, either via local Wi-Fi or a wider network if available. This interface provided real-time feedback and visualization of the detection output, allowing the farmer to confirm the count and observe the distribution of livestock directly from the portal. The left side of Fig. 5 shows a representative screenshot of the inference result in user's end, while the right side shows the UAV that was used for the herd detection.

This use case illustrated not only the simplicity of the AGRARIAN app onboarding process from the user's perspective, but also the robustness of the underlying infrastructure to handle real-time AI workloads at the edge, even in scenarios with dynamic and constrained connectivity. It highlighted the effectiveness of the AGRARIAN CPS model, where sensor data was captured, analyzed and transformed into actionable insight in situ, closing the loop between environment, computation and decision-making.

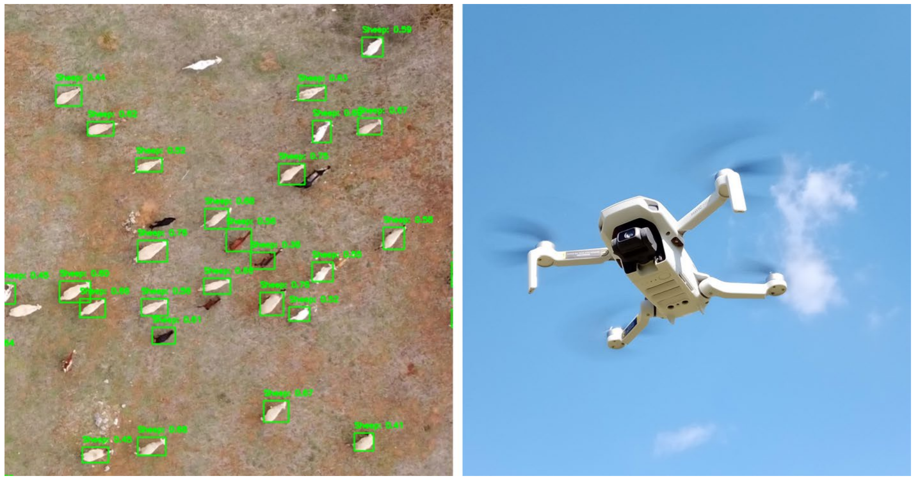

Fig. 5 Real-world example from the AGRARIAN field trial where aerial imagery from a UAV is processed using an edge-deployed YOLOv11 model to detect and count sheep in real time

5.2 App performance measurements

To ensure that the validation of the AGRARIAN platform reflected real-world conditions, all the testing was conducted using the available commercial communication infrastructure of the remote trial area. The 5 G tests were carried out over a commercial public mobile network, simulating typical coverage and load conditions encountered by rural users, while the satellite connectivity tests utilized the Starlink constellation system [37], a commercially available LEO satellite service. By relying on non-specialized, production-grade communication systems, rather than controlled lab environments or proprietary links, the trial aimed to demonstrate the platform's practical viability and performance in real operational settings, where infrastructure is shared, variable and often constrained.

To validate the AGRARIAN platform in remote underserved agricultural areas, the trial location was selected in the 5 G cell edge where four different types of networks (Ethernet, cellular, WiFi, satellite) were used to transmit the same type of video feed and to assess how communication conditions affected the AI model performance. The Ethernet connection provided a baseline measurement for ideal conditions, offering stable bandwidth and negligible frame loss. The weak 5 G link was tested at the edge of a cellular coverage area, where signal degradation introduced higher latency and intermittent bandwidth drops. The Wi-Fi connection simulated long-range rural scenarios by increasing the distance between the drone's transmitter and the receiver, leading to degraded signal strength and increased jitter. The Starlink satellite connection was tested under suboptimal weather conditions (heavy clouds), allowing us to observe the platform performance under non-ideal, non-terrestrial and variable-latency communication links.

During each trial, three key performance indicators were measured: Average confidence level of sheep detection, inference time and the number of dropped or lost frames. The selection of these performance metrics was initiated by their direct relevance to validating the responsiveness, reliability and accuracy of the AGRARIAN platform in edge AI deployments. Average confidence reflects the quality and stability of the AI model detections under varying conditions, indicating how communication quality affects inference certainty. Inference time measures the responsiveness of the system, capturing the total latency from data capture to usable results, critical for real-time decision-making in agricultural applications. Frame loss quantifies the platform ability to handle data transmission under fluctuating network conditions, revealing the robustness of the end-to-end video stream from sensor to edge node. Together, these metrics provide a comprehensive view of how well the system can maintain AI service integrity across hybrid connectivity scenarios.

It is noted that the goal of the field trial was not to perform a formal benchmarking of the specific AI model or to define strict performance thresholds, but rather to validate the ability of the AGRARIAN platform to maintain continuous service under variable connectivity conditions in a real-world agricultural setting. This proof-of-concept demonstration was conducted as a single, representative trial, focusing on the system responsiveness and adaptability in rural areas, when switching to alternative communication options and especially when using satellite commercial links. As such, comparative metrics were collected to highlight the practical impact of existing different communication paths, not to set universal thresholds for application acceptability. Future phases of the project may involve broader validation, with multiple use cases and trial replications under varying conditions.

5.3 Data analysis and validation

The test environment was carefully configured to reflect realistic rural deployment conditions. The Ethernet connection served as the ideal reference line, providing a stable uplink data rate of approximately 800 Mbps with negligible latency or jitter, enabling the AI inference application to operate under optimal conditions. In contrast, the 5 G connection was tested at the edge of the cell coverage area, where the signal was weak and susceptible to variability due to fading and congestion, introducing moderate latency and frame loss. The Wi-Fi link represented a long-range deployment scenario, where attenuation from distance and obstacles significantly degraded the signal strength, leading to lower throughput and higher jitter. The varying link characteristics allowed us to assess the robustness of the AGRARIAN platform and the resilience of its edge inference pipeline under diverse and challenging network environments.

While the terrestrial links (Ethernet, 5 G, and Wi-Fi) exhibited almost consistent bandwidth profiles, with minimum fluctuation in uplink speed rates, depending mainly on signal attenuation due to obstacles or distance from the base station, the satellite link experienced moderate to severe uplink bitrate variabilty. As shown in Fig. 6, the Starlink upload bitrate values fluctuate between 20 Mpbs and 70 Mbps, with a mean value of 45 Mbps. This instability is a result of the dynamic nature of LEO satellite Starlink constellation, where satellites rapidly traverse the sky and frequent handovers occur between them.

For the detailed study of the impact of each communication type on application performance, multiple measurements were collected across repeated trials for all network configurations. For each key performance metric, including inference confidence level, inference latency, and frame loss, we aggregated the recorded values and constructed Cumulative Distribution Functions (CDFs). These CDFs allowed the representation of the full distribution of results over time, rather than relying solely on averages or fixed thresholds. By plotting each communication scenario (Ethernet, 5 G, Wi-Fi, and Starlink) on a common CDF graph, a direct comparison of system behavior was enabled,

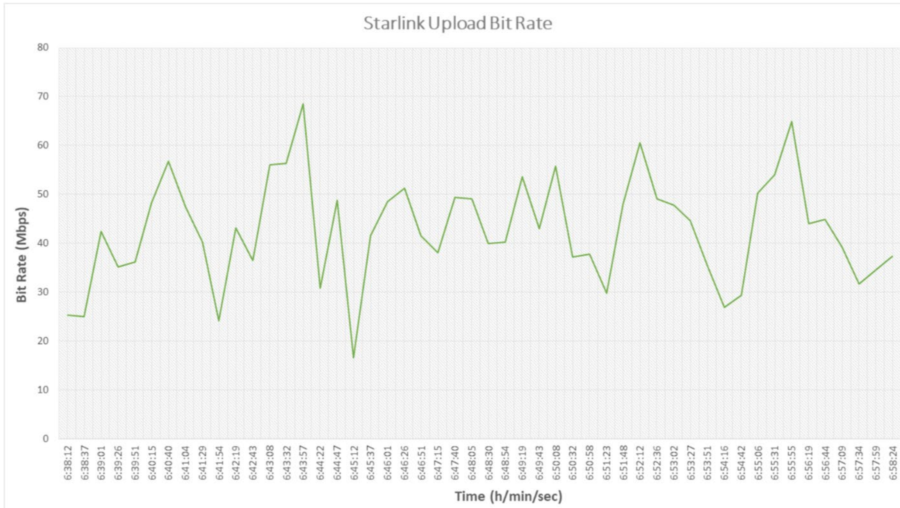

Fig. 6 Recorded fluctuations in Starlink uplink bandwidth during the sheep detection use case, illustrating the impact of satellite conditions on network throughput

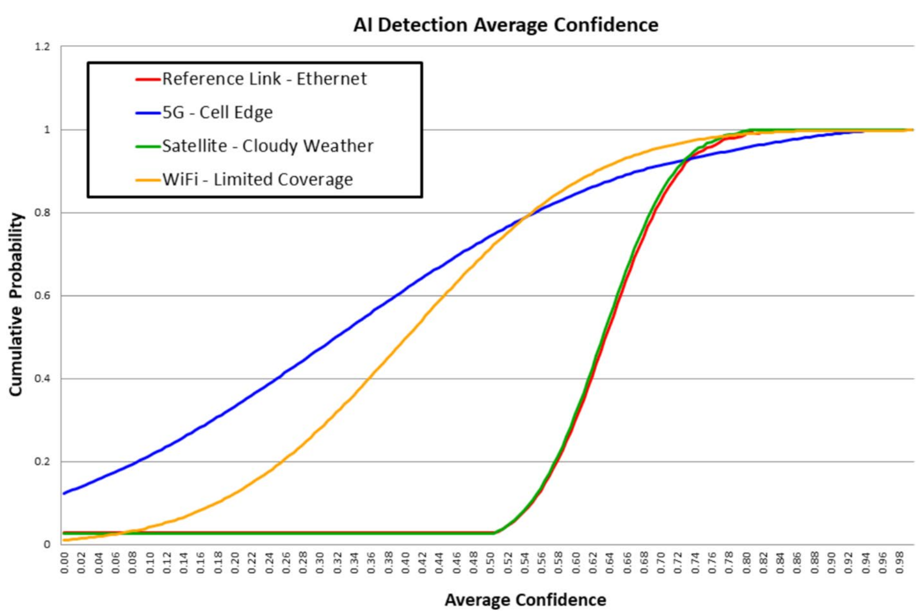

Fig. 7 Cumulative Distribution Functions (CDFs) showing the AI inference confidence levels across different communication links (5 G, Wi-Fi, Ethernet, Starlink), highlighting that satellite link delivers performance comparable to the reference wired Ethernet connection

under varying network conditions, highlighting differences in consistency, tail performance, and robustness.

Figure 7 shows the inference confidence analysis. The model exhibited high average confidence levels under Ethernet and satellite conditions, having almost identical graph curves, with a a probability of 80% for average confidence level higher than 0.58. The 5 G cell edge scenario resulted in a lot lower confidence at early percentiles (minimum 0.08 for 80% probability), likely due to limited bitrate and high frame loss. The weak Wi-Fi link also underperformed in comparison to the reference and the satellite links, with a confidence level of 0.25 for 80% probability, revealing its limitations in remote deployments.

The Inference time analysis, a key metric for real-time responsiveness, showed minimal latency for the Ethernet and satellite links, that again resulted in almost identical graphs. Figure 8 shows that the median values for both of them were at the levels of 40 ms, with no latency higher than 2.2 s. The 5 G link introduced moderate delays of almost 500 ms median, but reaching up to 6.2 s in the worst cases. The Wi-Fi link offered the heavier starting delays, with a median at 700 ms, while its maximum reached almost 3 s. This graph shows that the LEO satellite communications can act as a practical fallback for time-sensitive edge applications.

Frame loss measurement was the most sensitive indicator of communication reliability. As depicted in Fig. 9 the reference link (Ethernet connection) experienced almost zero frame loss (0.5%), while the satellite link also maintained strong performance, with minimal losses (<3%). In contrast, 5 G and Wi-Fi links suffered severe frame losses. The 5 G link experienced a frame loss of 87%, while the WiFi dropped 43% of the frames during the test duration, indicating potential service disruption and AI model performance degradation.

The current trial focused on sheep detection, so a targeted dataset was used. Although the AI model demonstrated high accuracy in this context, its generalization to different species or conditions has not yet been evaluated. This trial served as a proof-of-concept to demonstrate AGRARIAN capability to support real-time AI-based inference in constrained environments. Future work will involve extending training datasets and validating the model across multiple use cases, animal types and environmental scenarios to assess robustness and general applicability.

However, the real-world field trial that was conducted, successfully showcased the operation of the AGRARIAN platform under real environment and commercial communication links. The system autonomously orchestrated the application implementation and the user received the final service, without any technical expertise or action required from the user. As the uninterrupted and reliable connectivity is a critical element for

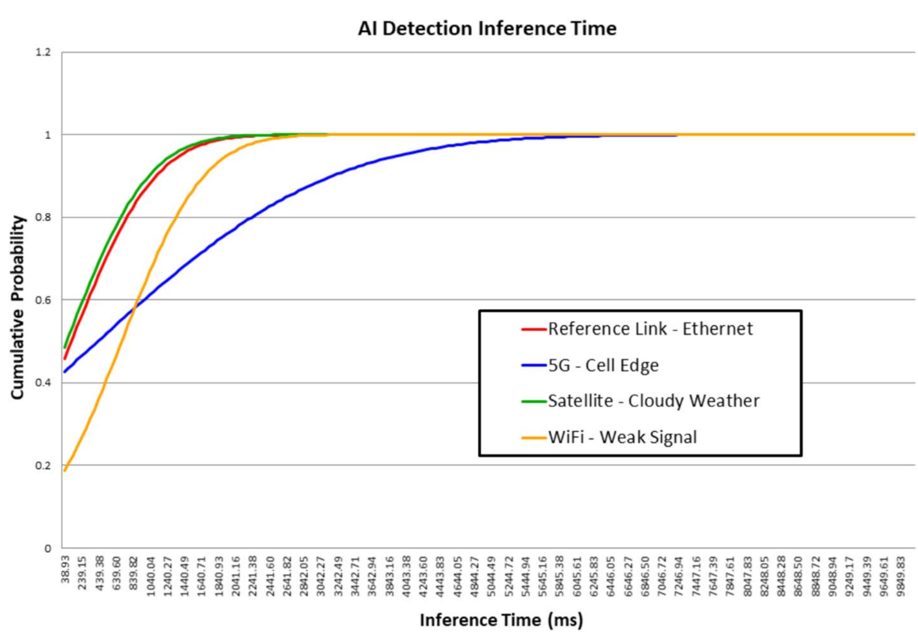

Fig. 8 CDF comparison of the inference processing time required per video frame, evaluated over various network types, to assess system latency performance under real conditions

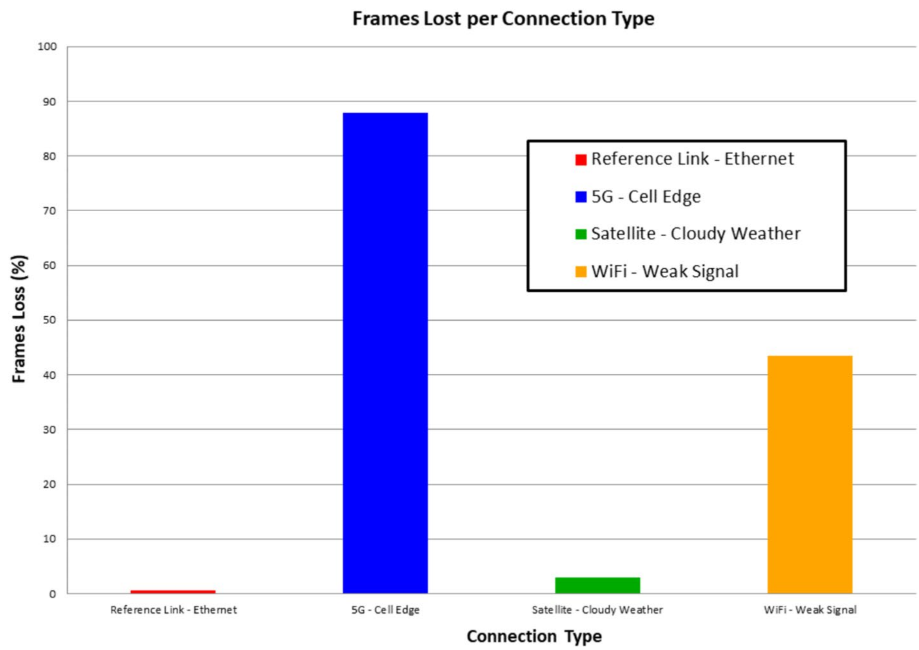

Fig. 9 Bar chart showing the percentage of video frames lost during real-time streaming and inference, categorized by communication link type, to evaluate reliability and quality of service

future farming, the AGRARIAN platform communication infrastructure was evaluated. The AI model performance measurements under different potential choices of communications types in a remote farm, showcased that hybrid communications, having LEO satellites as a fallback alternative, could offer a great solution for smooth service implementation in remote agricultural areas. The satellite link, not only held high quality performance, but produced quite similar results as the ideal reference Ethernet link did. The overall field trial validated the AGRARIAN system architecture as a resilient, self-organizing CPS, capable of offering application continuity in diverse connectivity environments, requiring minimum technical actions from the end users.

5.4 Experimental setup and reproducibility details

To ensure transparency and reproducibility of the field trial, this section provides detailed information and insight about the trial setup, including hardware and software configurations, environmental/terrain conditions and key parameters.

5.4.1 Trial conditions

The performed field trial was conducted in a real farm area with the following characteristics:

- **Type**:Proof-of-concept demonstration in real world scenario
- **Location**:At the wider area of Tripolis, Arcadia, Greece
- **Terrain**: Rural with a mixed terrain, a combination of flat and hilly farmland
- **Weather**: Heavy cloudy with a temperature of around 18 °C.
- **Area Size**: Approximately 10,000 m^2.
- **Obstacles**: Sparse trees and scattered low building structures.

5.4.2 Testbed details

The testbed characteristics of the performed trial were:

- **UAV**: The aerial video source was captured using a DJI Mini drone. The video stream ended up to the UAV controller which was connected to the hybrid communication backbone, enabling streaming through 5 G, Wi-Fi, Ethernet or satellite links. No onboard inference was performed on the UAV; instead, real-time streaming was relayed to the ground inference node.
- **Edge Node**: The edge node was a high-performance PC running Ubuntu 22.04, equipped with an Intel i9 13th-generation processor, 32 GB RAM, and Docker Engine for container orchestration. To support real-time inference and video streaming, the YOLOv11 model was deployed alongside MediaMTX (an RTSP media server) [38] within the same Kubernetes pod or node. Resource requests and limits were carefully configured to ensure that both the AI model and MediaMTX could operate reliably without interfering with each other. CPU and memory resources were allocated using Kubernetes specifications, where resource requests defined the minimum needed for stable operation and limits enforced upper bounds to prevent overuse. For example, the pod may request 4 vCPUs and 8 GB of memory to guarantee baseline performance, with limits set to 12 vCPUs and 16 GB to allow for peak load handling. This setup ensured the YOLOv11 model had sufficient compute for inference while allowing MediaMTX to handle video streams smoothly, maintaining low latency and preventing crashes due to resource exhaustion.
- **Communication Backbone**: The testbed integrated a hybrid communication infrastructure featuring the public 5 G network of the trial area, provided by Cosmote [39] and commercial Starlink constellation connectivity.

5.4.3 AI model details

The use case involved detecting sheep and counting them via video analysis. The AI model characteristics were:

- **Model Dataset**: Two annotated Roboflow [40] datasets were combined, preprocessed for YOLO compatibility with auto-orientation and resizing to 640×640 pixels. The dataset was split into 92% training (16,664 images), 6% validation (1,000 images) and 2% testing (200 images). Data augmentation included mosaic, horizontal flipping (0.5 probability), random rotations (±15○), shear (±10%) and color adjustments (hue, saturation, brightness, exposure)
- **Model Training**: YOLOv11 was trained for 100 epochs using a batch size of 16 with pretrained weights and automatic optimizer selection. The initial learning rate was set at 0.01, momentum at 0.937 and weight decay at 0.0005. Mixed precision (AMP) training was enabled for efficiency, with 8 data loader workers. Validation was deterministic, performed after each epoch, using early stopping with 100 epochs patience, saving the best model checkpoint.

5.4.4 Reproducibility summary

Table 1 summarizes the key configuration parameters used during the field trial to facilitate repeatability and transparency. It includes hardware specifications, software stack configurations, model training setup and trial conditions. These parameters aim

Table 1 Summary of Field Trial Setup and Parameters

Category	Details
Type	Proof-of-concept demonstration in real world scenario
Location	Tripolis, Arcadia, Greece
Terrain	Rural, mixed flat and hilly farmland
Weather	Heavy cloudy, ~ 18°C
Area Size	10,000 m^2
Obstacles	Sparse trees, scattered low buildings
UAV	DJI Mini; video relayed via controller to communication backbone
Edge Node	Ubuntu 22.04, Intel i9 13th Gen, 32 GB RAM, Docker
Media Server	MediaMTX (RTSP server), containerized in K3s
AI Model	YOLOv11, sheep detection
Dataset	Roboflow (16,664 training, 1000 validation, 200 testing)
Training Parameters	100 epochs, batch size 16, learning rate 0.01, AMP enabled
AI Performance Metrics	Average Confidence, Inference time, Frames lost
Communication Backbone	Cosmote 5G, LEO Satellite (Starlink constellation), WiFi and Ethernet

to provide a holistic view of the experimental setup so that future experiments, either for validation, comparison or extension, can be accurately replicated. However, as it has been already mentioned, the trial was designed as a real-world, proof-of-concept demonstration rather than a controlled benchmark. Therefore, while detailed, the configuration reflects a practical deployment scenario and may be adapted and modified in future trials.

5.5 Farmers–experience and feedback

To evaluate the platform's usability and accessibility by non-technical users, a group of local farmers was invited to participate during the field trial. Using their mobile phones and the local public internet connection they accessed the AGRARIAN portal and were able to observe the real-time sheep detection and counting service directly from the platform. Although their role was observational, they acted as end-users receiving the output of the deployed application under actual field conditions. Informal feedback collected on-site indicated high satisfaction and strong interest in the technology. Farmers expressed enthusiasm about the potential of AGRARIAN to support day-to-day operations without requiring technical expertise, validating the system aim of being intuitive, accessible and practical for real-world agricultural stakeholders.

6 Scalability and economic viability

To assess the viability of the developed platform beyond the proof-of-concept trial, both the economic feasibility and the scalability of the system deployment are discussed. While AGRARIAN integrates advanced components such as UAVs, edge processing units, satellite links and lightweight orchestration frameworks, the platform has been designed with modularity and cost-effectiveness in mind, specifically to support scalability toward small and medium size farms. One of the core objectives of AGRARIAN is to lower the entry barriers for digital farming tools by adopting open-source technologies (e.g., K3s, Docker, GitOps), cost-efficient edge nodes (e.g., Raspberry Pi 5, Jetson Nano or commodity PCs) and consumer-grade UAVs, such as the DJI Mini series, whose cost is quite affordable.

Moreover, AGRARIAN leverages dynamic deployment and resource-efficient orchestration, meaning that applications are only installed and run when needed, significantly optimizing resource consumption. The use of K3s, as a lightweight orchestrator, further reduces infrastructure overhead, making it suitable even for edge devices with limited compute capacity.

Regarding connectivity, although satellite communication is generally considered cost-intensive, AGRARIAN supports hybrid networking, which prioritizes terrestrial links like 5 G, Wi-Fi or Ethernet. Satellite fallback is only triggered under specific conditions (e.g., outage or low signal quality), reducing bandwidth usage and associated costs. The use of commercial low-orbit services (e.g., Starlink), which are increasingly affordable and accessible, reinforces the platform feasibility in rural and underserved areas.

At scale, the economic model is adaptable: edge hardware can be shared among cooperative farms or agricultural unions and UAVs can be managed by service providers or local municipalities. The platform is also structured to support third-party services and open call extensions, which can lower costs through innovation and competition.

Therefore, AGRARIAN is not only suited for large agribusinesses but also strategically designed to scale across clusters of small and medium size farms, making advanced precision agriculture tools more accessible, sustainable and economically viable.

7 Conclusions

In this paper the AGRARIAN platform was presented and validated through a real world field trial. The presented platform distinguishes itself through a clear innovation-driven approach: it places technology in the service of agriculture by building a flexible, inclusive ecosystem that brings together farmers, developers and policy makers. This collaborative foundation enables stakeholders to co-develop and deploy smart agricultural solutions in a way that is scalable, sustainable and aligned with real-world needs. This multilevel framework promotes the co-development and seamless deployment of digital farming applications, ensuring that innovations are not only technically robust but also usable, accessible, and resilient in real agricultural settings. By enabling service continuity in remote areas and supporting both local farm operations and broader policy formulation, AGRARIAN stands as a scalable, inclusive and future-proof model for sustainable agriculture.

The AGRARIAN system architecture enables real-time, distributed applications by utilizing containerized service orchestration over heterogeneous networks (e.g. 5 G, WiFi, LoRa etc.). The LEO satellite communications have a key role in maintaining connectivity continuity, as they offer global coverage with high bandwidth and low latency. The developed platform, implementing a lightweight Kubernetes (K3s) orchestration for autonomous applications management, supports dynamic deployment of digital tools and modern agricultural services directly at the network edge, while providing dedicated interfaces adapted to farmers, developers and policymakers for smoothly facilitating their inclusive participation in the digital agricultural ecosystem. The AGRARIAN platform allows farmers to request services on-demand, developers to deploy, test and continuously update applications via CI/CD workflows, while offers to policy makers access to aggregated data and insights for macro-level decision making.

The purpose of the conducted field trial was to validate AGRARIAN platform operational capability in a realistic livestock monitoring scenario. The system successfully installed the requested application on the edge node, while the commercial satellite link, the Starlink constellation, successfully supported the communications. Field measurements showed that, in spite of satellite uplink speed fluctuation (due to frequent satellite handovers being required because of the high orbital speed of the LEO constellation), the service was not affected, even when the weather was heavily cloudy. Thus, it was demonstrated that the hybrid communications infrastructure of the AGRARIAN platform ensures the operation of scalable and sustainable farming solutions, in real world scenarios in remote agricultural areas.

The AGRARIAN platform is currently being further developed and tested. Future work includes statistical and quantitative analysis of field results and platform validation with broader field trials in diverse areas and agricultural scenarios, with active participation of more end users. Feedback will be collected from all types of stakeholders (i.e. farmers, developers and policy makers) for system improvement and fine tuning.

A future key element of the AGRARIAN ecosystem is the support of federated learning mechanisms [41]. As the AGRARIAN platform is still under active development, current deployment supports localized edge AI processing with model inference executed independently at each node. Federated learning will enable distributed knowledge sharing across farms without exchanging raw data, improving model accuracy. In particular, enabling collaborative AI model training across multiple farms, while preserving data privacy and minimizing bandwidth consumption, could significantly enhance AGRARIAN intelligence and adaptability. Integrating federated learning would allow decentralized edge nodes to contribute to shared AI improvements without transferring sensitive raw data, aligning with the platform goals of security, scalability and stakeholder empowerment in rural environments. This remains a key research objective for the next stages of development and evaluation.

In addition to future technical planning, widespread adoption of the system is also a target. While the AGRARIAN platform has strong potential for offering an integrated solution to smart farming, the adoption among farmers, developers and policymakers faces significant challenges. For farmers, especially the non-technical ones, system elements (e.g. user interfaces) should be quite simple, while the hardware cost must be minimized to suit small and medium-sized farms. Developers need efficient Application Programming Interfaces (APIs), well-documented and suitable to their needs, so that they can build, deploy and test applications quickly and efficiently. Most importantly, the system should offer to policymakers trust in data integrity, interoperability with external, national systems and usable regional insights, without overwhelming them with unnecessary and complex information. To answer the aforementioned challenges, future actions should include dedicated training programs for all types of stakeholders, targeted farmers' funding for covering the cost for multiple edge devices, mechanisms (e.g. benefits, rewards etc.) to encourage developers to build, deploy and maintain applications in the platform and standardized protocols for data sharing, security and governance. By combing all the aforementioned discussed actions with the principles of Cyber-Physical Systems and self-organizing networks, the AGRARIAN platform contributes strongly to the new era of smart farming and sustainable food production.

Abbreviations

ADSS	Agricultural Decision Support System
AI	Artificial Intelligence
ML	Machine Learning
AMP	Automatic Mixed Precision
API	Application Programming Interface
CAS	Cyber-Agricultural System

CDF	Cumulative Distribution Function
CI/CD	Continuous Integration/Continuous Deployment
CPAS	Cyber-Physical Agricultural Systems
CPS	Cyber-Physical System
CPU	Central Processing Unit
EU	European Union
GB	Gigabyte
GEO	Geostationary Earth Orbit
HAPS	High Altitude Platform Systems
IT	Information Technology
LEO	Low Earth Orbit
NB	Narrowband
ONNX	Open Neural Network Exchange
PC	Personal Computer
RAM	Random Access Memory
RPC	Remote Procedure Call
RTMP	Real-Time Messaging Protocol
RTSP	Real-Time Streaming Protocol
SDN	Software-Defined Networking
UAV	Unmanned Aerial Vehicle
YOLO	You Only Look Once

Author contributions

M.B.: Manuscript writing and implementation of measurements and field trials M-A.K. and G.X.: Overall system conceptualization T.C.: System architecture D.O. and S.A.: System workflow orchestration A.D.: Edge computation A.M.: AI models N.S.: Manuscript final review and editing All authors reviewed the manuscript.

Funding

This research is funded by the Horizon Europe project AGRARIAN (Grant Agreement No. 101134128). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union.

Data availability

The datasets used and analyzed during the current study available from the corresponding author on reasonable request

Declarations

Ethics approval and consent to participate

Not applicable.

Consent to participate

Not applicable.

Consent for publication

Not applicable.

Competing interests

The authors declare no competing interests.

Received: 9 June 2025 / Accepted: 25 August 2025

Published online: 02 September 2025

References

1. Fountas S, Espejo-García B, Kasimati A, Gemtou M, Panoutsopoulos H, Anastasiou E. Agriculture 5.0: cutting-edge technologies, trends, and challenges. IT Prof. 2024;26(1):40–7.
2. Berger B, StephanHäckel Häfner L. Organizing self-organizing systems: a terminology, taxonomy, and reference model for entities in cyber-physical production systems. Inf Syst Front. 2021;23(23):391–414.
3. Zampetti F, Geremia S, Bavota G, Di Penta M. Ci/cd pipelines evolution and restructuring: A qualitative and quantitative study. In: 2021 IEEE international conference on Software Maintenance and Evolution (ICSME). 2021. pp. 471–482
4. Sarkar S, Ganapathysubramanian B, Singh A, Fotouhi F, Kar S, Nagasubramanian K, et al. Cyber-agricultural systems for crop breeding and sustainable production. Trends Plant Sci. 2024;29(2):130–49.
5. Verdouw C, Tekinerdogan B, Beulens A, Wolfert S. Digital twins in smart farming. Agric Syst. 2021;189:103046.
6. Alwis SD, Hou Z, Zhang Y, Na MH, Ofoghi B, Sajjanhar A. A survey on smart farming data, applications and techniques. Comput Ind. 2022;138:103624.
7. Friha O, Ferrag MA, Shu L, Maglaras L, Wang X. Internet of things for the future of smart agriculture: a comprehensive survey of emerging technologies. IEEE/CAA J Automat Sinica. 2021;8(4):718–52.
8. Nawaz M, Babar MIK. Iot and ai for smart agriculture in resource-constrained environments: challenges, opportunities and solutions. Discov Internet Things. 2025;5:24.
9. Sharma R, Parhi S, Shishodia A. Industry 4.0 applications in agriculture: Cyber-physical agricultural systems (cpass). In: Advances in Mechanical Engineering (Proc. ICAME 2020). Singapore: Springer; 2021. p. 807–13.
10. Kong L, Tan J, Huang J, Chen G, Liu C, Jia W, et al. Edge-computing-driven internet of things: a survey. ACM Comput Surv. 2022;55(8):1–41.

11. Habibi P, Farhoudi M, Kazemian S, Khorsandi S, Leon-Garcia A. Fog computing: a comprehensive architectural survey. IEEE Access. 2020;8:69105–33.
12. Tu M, Wang X, Ren S, et al. A survey of recent advances in edge-computing-powered artificial intelligence of things. IEEE Internet Things J. 2021;8(18):13849–75.
13. Yousefpour A, Vu T, et al. All one needs to know about fog computing and related edge computing paradigms: a complete survey. J Syst Architect. 2019;98:289–330.
14. Leite L, Rocha C, Kon F, Milojicic D, Meirelles P. A survey of devops concepts and challenges. ACM Comput Surv. 2019;52(6):1–35.
15. Ferry N, Nguyen PH, Song H, Rios E, Iturbe E, Martinez S, et al. Continuous deployment of trustworthy smart iot systems. J Obj Technol. 2020;19(2):16–123.
16. L'opez-Viana R, D'iaz J, P'erez JE. Continuous deployment in iot edge computing: a gitops implementation. In: Proc. 8th iberian conference on information systems and technologies (CISTI). 2022.
17. Kubernetes Authors: Kubernetes: production-Grade Container Orchestration. 2014. https://kubernetes.io/ Accessed 2025-06-02.
18. Rong G, Xu Y, Tong X, Fan H. An edge-cloud collaborative computing platform for building aiot applications efficiently. J Cloud Comput. 2021;10:36.
19. Kim HS, Alizai MH, Duquennoy S, Landsiedel O. Continuous integration and deployment of robotic iot applications. In: Proc. 2023 IEEE international conference on embedded and ubiquitous uomputing (EUC). 2023.
20. Centenaro M, Costa CE, Granelli F, Sacchi C, Vangelista L. A survey on technologies, standards and open challenges in satellite iot. IEEE Commun Surv Tutor. 2021;23(3):1693–720.
21. Kodheli O, Lagunas E, Maturo N, Sharma SK, Shankar B, Montoya JFM, et al. Satellite communications in the new space era: a survey and future challenges. IEEE Commun Surv Tutor. 2021;23(1):70–109.
22. Giordani M, Zorzi M. Non-terrestrial networks in the 6g era: challenges and opportunities. IEEE Network. 2021;35(2):244–51.
23. Wong OS, Gregory MA, Li S. Integration of non-terrestrial network for 5g iot and future 6g. J Telecommun Dig Econ. 2025;13(1):384–405.
24. Patel S, Joshi R, Nguyen T. Hybrid iot communication architectures for rural smart farming: a case study. Comput Electron Agric. 2023;206:107537.
25. SmartAgriHubs Authors: SmartAgriHubs. 2018. https://smartagrihubs.eu Accessed 2025-07-20.
26. DEMETER Authors: DEMETER-Building an Interoperable, Data-Driven, Innovative and Sustainable European Agri-Food Sector. 2019. https://h2020-demeter.eu Accessed 2025-07-20.
27. FIWARE Authors: FIWARE—a Curated Framework of Open Source Software Platform Components. 2016. https://www.fiware.org/ Accessed 2025-07-20.
28. OpenAgri Authors: OpenAgri. 2024. https://horizon-openagri.eu/ Accessed 2025-07-20.
29. Nostradamus Authors: Nostradamus–Data Cube and Copernicus data for Food Security and European Independence. 2024. https://nostradamus-project.eu/ Accessed 2025-07-20.
30. K3s Project: lightweight kubernetes. 2019. https://k3s.io/ Accessed 2025-06-02.
31. Beetz F, Harrer S. Gitops: The evolution of devops? IEEE Softw. 2022;39(4):70–5.
32. Ghiglione M, Serra V. Opportunities and challenges of ai on satellite processing units. In: Proceedings of the 19th ACM International Conference on Computing Frontiers 2022. pp. 221–224.
33. Google: gRPC: a high performance, open-source universal RPC framework. 2016. https://grpc.io/ Accessed 2025-06-02.
34. ONNX Community: open Neural Network Exchange (ONNX). 2017. https://onnx.ai/ Accessed 2025-06-02.
35. Intel Corporation: OpenVINO Toolkit. 2018. https://github.com/openvinotoolkit/openvino Accessed 2025-06-02.
36. Eclipse Foundation: eclipse Kura: an Open Source IoT Edge Framework. 2015. https://eclipse.dev/kura/ Accessed 2025-06-02.
37. SpaceX: Starlink: high-Speed Internet from Space. 2020. https://www.starlink.com/ Accessed 2025-06-02.
38. Bluneviron Authors: MediaMTX - Media Server for RTSP, RTMP, HLS, WebRTC and SRT. 2024. https://github.com/bluenviron/mediamtx Accessed 2025-07-19.
39. Cosmote Authors: Cosmote Telekom. 1997. https://www.cosmote.gr/hub/ Accessed 2025-07-19.
40. Roboflow Authors: Roboflow: Collect, Annotate, Train, Deploy. 2024. https://roboflow.com/ Accessed 2025-07-19.
41. Elavarasan RM, Sankar P, Sanjeevikumar R, Holm-Nielsen JB, Hossain M, Kumaravelan N. Internet of things (iot)-enabled architectures for smart agriculture and farming—a review. Soft Comput. 2023;27(5):1501–27.

Publisher's Note