Future Farming XR: An Extended Reality Command Center for Autonomous Fleet Management and Farm Supervision

May 3, 2026

Team Information

Institution: Mississippi State University
Team Name: Maroon Robotics Team
Participating Division: Division I

Team Advisors

- **Name:** Dr. Dong Chen
  Position: Assistant Professor
  Email: dchen@msstate.abe.edu
  Department: Agricultural and Biological Engineering

- **Name:** Dr. Alex Thomasson
  Position: Director, Professor
  Email: athomasson@abe.msstate.edu
  Department: Agricultural Autonomy Institute, Agricultural and Biological Engineering

Team Members

- **Moeen Ul Islam** – Graduate, Agricultural Science
- **Bishal Adhikari** – Graduate, Agricultural Science
- **Cheng Ouyang** – Graduate, Agricultural Science
- **Kodia Watanabe** – Graduate, Biosystems Engineering
- **Emma Lovell** – Undergraduate, Biomedical Engineering
- **Lorien Harvey** – Undergraduate, Biosystems Engineering
- **Eliana Wile** – Undergraduate, Biomedical Engineering
- **Mason Jones** – Undergraduate, Mechanical Engineering

Maroon Robotics Team

Contents

1 Introduction

Future Farming XR is an unified farm-operations platform that integrates multiple ground vehicles, aerial scouting systems, and intelligent software tools to help U.S farmers monitor fields, detect problems, and remotely supervise or control machinery in a simple, accessible, and immersive way.

1.1 Problem Statement

While autonomous robots and drones offer a necessary solution to the severe labor shortages and rising input costs plaguing small and mid-scale U.S. row-crop farmers (Benos et al., 2023; Yerebakan & Hu, 2024), widespread adoption is hindered by a fundamental trust gap (Devitt, 2018). Current autonomous systems frequently operate as digital black boxes (Devitt, 2018). They remain heavily reliant on technical expertise and are notoriously difficult to supervise in real time (Yerebakan & Hu, 2024).

The true bottleneck in modern agriculture is no longer just the hardware, but the physical and digital distance between the farmer and the machine. When equipment inevitably encounters unpredictable field conditions, heavy residue, or system faults, operations grind to a halt. Without an integrated, easy-to-use management interface, farmers are left with limited awareness of why the machine stopped and no simple way to intervene remotely. This forces operators to physically travel to the field to troubleshoot routine errors, defeating the time-saving purpose of automation and delaying the timely detection of critical yield-limiting issues like weeds, pests, and nutrient stress.

To fully benefit from modern automation technologies, farmers do not need more complex robots; they need a safety net. They require a centralized, intuitive command center that bridges the gap between human and machine. Farmers need a solution that provides total situational awareness, plain-language troubleshooting, and simple remote supervision, allowing them to virtually step into the driver's seat to guide operations and resolve anomalies without needing a background in robotics or engineering.

2 System Design

Future Farming XR is a unified extended reality (XR) platform that brings farm monitoring, autonomous fleet management, and remote intervention into a single immersive interface running on Meta Quest virtual reality headsets. The system facilitates seamless connectivity between virtual farm representations and real-world field operations, providing an intuitive experience designed to foster trust in farmers toward autonomous vehicle adoption and digital farming practices.

The platform addresses a critical trust gap: operators are reluctant to deploy autonomous vehicles without the ability to see, understand, and intervene when the unexpected occurs. Future Farming XR closes this gap by combining always accessible remote presence, real-time data visualization, and hands-on remote control, regardless of the operator's physical distance from the field.

2.1 Platform Architecture

We imagined the Future Farming XR ecosystem as an organization of four different layers spanning hardware and software, each with clearly defined responsibilities and interfaces.

Layer	Scope	Key Technical Components
Farm & Hardware Layer	Physical robots, sensors, and field infrastructure	Farm-ng Amiga Robot, UGVs, UAV drones, 360 degree cameras, RTK GPS, Starlink terminal
Network Layer	Encrypted connectivity across all nodes	Tailscale VPN (WireGuard), Starlink satellite internet, MQTT broker, WebRTC signaling
Data & Orchestration Layer	Data routing, protocol translation, state management	MQTT pub/sub, WebRTC peer connections, `IVehicleDataProvider` abstraction, coordinate mapping
Interface Layer	Operator-facing immersive application	Unity on Meta Quest 3, URP rendering, Meta XR Interaction SDK, hand tracking, spatial UI

Table 1: Platform Architecture Overview

2.2 Software Interfaces

Within the interface layer, the farmer experience is structured into three distinct software layers, each serving a different operational need and trust-building function.

Software Layer	Purpose	Core Experience
Strategic Layer	Situational awareness and fleet-wide oversight	Digital twin diorama of the farmer's plot with real-time robot trajectories, vehicle telemetry, nutrition maps, and operations dashboard
Tactical Layer	Direct intervention and hands-on control	Telepresence with remote intervention capabilities: immersive real-time 360° video streaming from the driver's seat, AI assistant during teleoperation
Safety Layer	Training and skill development without risk	Tractor driving simulator for remote intervention training: building operator comfort with remote control before deploying on live equipment

Table 2: Software Layer Components

It is worth noting that these three layers are not separate applications. Rather, they are modes within a single unified experience, The operator can transition naturally between them; monitoring the fleet at the strategic level, drilling into a specific vehicle at the tactical level when an alert requires attention, and practicing unfamiliar maneuvers at the safety level before attempting them on live equipment.

2.3 Farm & Hardware Layer

The farm and hardware layer encompasses all physical assets deployed in the field: autonomous vehicles, sensors, compute modules, and connectivity infrastructure. This layer provides the sensory and actuation foundation upon which all higher layers depend.

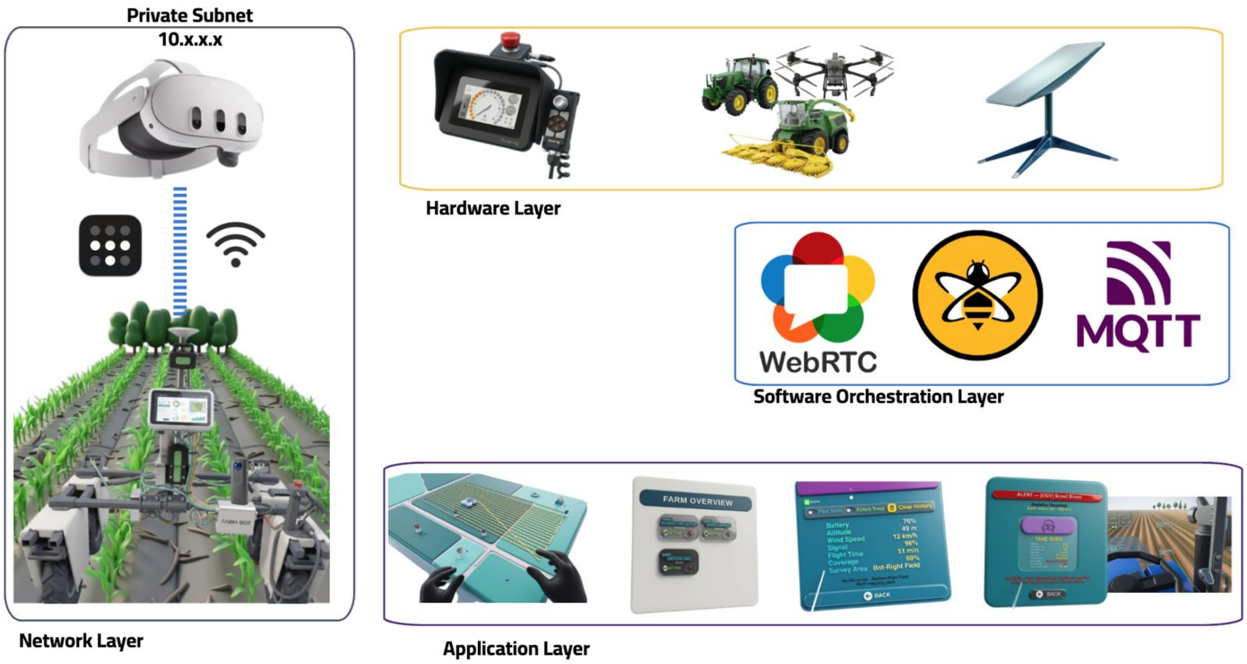

Figure 1: Platform layers overview

2.3.1 Primary Robot Platform , Farm-ng Amiga

The Farm-ng Amiga is a compact, electric UGV designed for row-crop agriculture. It runs a custom Linux distribution on an NVIDIA Jetson Xavier NX compute module, with a gRPC-based service architecture for sensor access and motor control.

Component	**Specification**
Chassis	Farm-ng Amiga, 4-wheel independent drive, electric
Compute	NVIDIA Jetson Xavier NX (6-core ARM, 384 CUDA cores, 8 GB RAM)
Camera	Luxonis OAK-D (stereo depth + RGB, 4K capable)
GNSS	Multi-band RTK GPS with centimeter-level accuracy
IMU/Filter	Onboard state estimation filter (x, y, heading, convergence status)
CAN Bus	Motor telemetry: linear velocity, angular velocity, per-wheel current
Network	WiFi 6 to local Starlink terminal or cellular hotspot
Battery	48 V Li-ion, runtime data available via CAN bus

Table 3: Farm-ng Amiga Hardware Specifications

2.3.2 Onboard Software Services

The Amiga runs a set of gRPC services accessible via the Farm-ng EventClient API. A custom Python bridge service subscribes to these streams and repackages the data for external transmission via MQTT.

Service	Data Provided	Native Rate
CAN Bus Service	Linear velocity, angular velocity, per-wheel motor data	50 Hz
GPS Service	Latitude, longitude, altitude, ground speed, fix quality	10 Hz
Filter Service	Fused pose (x, y, heading), convergence status	10 Hz
Oak Camera Service	RGB frames (1080p), stereo depth maps	30 Hz
Track Follower	Waypoint index, tracking error, autonomy mode	10 Hz

Table 4: Amiga Onboard gRPC Services

2.3.3 Field Connectivity , Starlink

Field internet connectivity is provided by a SpaceX Starlink terminal positioned at or near the field. Starlink provides broadband internet to the Amiga's WiFi radio and, through the Tailscale VPN, establishes encrypted tunnels to the operator and cloud services. Starlink's typical latency of 25–60 ms and download bandwidth of 50–200 Mbps is sufficient for simultaneous telemetry and video streaming. In deployments where Starlink is unavailable, LTE/5G cellular serves as a fallback with comparable performance.

2.3.4 Fleet Composition and Scalability

While the Farm-ng Amiga serves as the primary development and validation platform, the Farm & Hardware Layer is designed to accommodate heterogeneous fleets comprising multiple UGVs and UAVs operating concurrently on the same plot. Each vehicle, whether a ground-based harvester, planter, or scout, or an aerial survey drone, is treated as an independent node on the Tailscale network, running its own MQTT bridge service and publishing to the same topic hierarchy under a unique vehicle ID. The system imposes no architectural limit on fleet size; adding a new vehicle requires only deploying the Python bridge script on its onboard compute, assigning a unique ID, and connecting it to the Tailscale network. In the current implementation, the mock fleet demonstrates five simultaneous vehicles (three UGVs and two UAVs) with distinct trajectory patterns, telemetry schemas, and operational roles, validating that the data pipeline, AR/VR rendering, and dashboard UI all scale gracefully without per-vehicle code changes. Future field deployments envision mixed fleets where ground vehicles execute row-level tasks (planting, harvesting, scouting) while UAVs provide aerial coverage (NDVI mapping, perimeter surveillance), with all vehicles visible and controllable through the same unified diorama interface.

2.4 Network Layer

The Network Layer provides encrypted, NAT-traversing connectivity between all system nodes regardless of their network topology. The design assumes no public IP addresses, no port forwarding, and potentially constrained rural bandwidth.

2.4.1 Tailscale VPN Overlay

Tailscale creates a WireGuard-based mesh VPN connecting all system nodes (Amiga robots, operator devices, MQTT broker, signaling server) into a single virtual private network. Each node receives a stable `100.x.x.x` address that persists across network changes, reboots, and location changes. Tailscale handles NAT traversal automatically via DERP relay servers when direct peer-to-peer connections are not possible.

Property	Value
Protocol	WireGuard (kernel-level, encrypted)
Addressing	Stable `100.x.x.x` per device
NAT Traversal	Automatic (DERP relay fallback)
Latency Overhead	$<$ 5 ms (direct LAN), $\sim$50 ms (relayed)
Platforms	Linux (Amiga), macOS (dev Mac), Android (Quest)

Table 5: Tailscale VPN Properties

Protocol	Data Type	Direction	Rationale
MQTT v3.1.1	Telemetry, state, commands	Bidirectional	Lightweight pub/sub, QoS levels, retained messages for late-joining subscribers, minimal overhead
WebRTC	Live camera video	Robot $\rightarrow$ Operator	Peer-to-peer UDP streaming, adaptive bitrate, sub-200 ms latency, built-in NAT traversal via ICE

Table 6: Communication Protocol Selection

2.5 Data & Orchestration Layer

The Data & Orchestration Layer handles all data routing, protocol translation, and state management between the robots and the operator interface. It defines how data flows through the system, how it is structured, and how the interface consumes it.

2.5.1 WebRTC Video Pipeline

2.6 Interface Layer , Strategic Software Layer

The Strategic Layer provides fleet-wide situational awareness through an immersive 3D digital twin of the farmer's plot. The operator views a tabletop-scale diorama of the entire farm, with autonomous vehicles rendered as miniatures moving in real time across a satellite-textured terrain.

2.6.1 Digital Twin Diorama

The farmland environment is rendered as a $10\,\text{m} \times 10\,\text{m}$ mesh (origin at center, coordinates -5 to $+5$ in X and Z) textured with satellite imagery of the actual farm. The terrain was created using Unity Terrain,

Parameter	Value
Source	OAK-D RGB sensor (1920×1080 native)
Encoder	Hardware H.264 on Xavier NX
Stream Resolution	720p (adaptive: drops to 480p under bandwidth constraint)
Frame Rate	10 fps nominal (adaptive: 5 fps $\rightarrow$ 2 fps under constraint)
Target Bitrate	2–4 Mbps
Transport	WebRTC peer-to-peer over UDP
Signaling	WebSocket via Tailscale (SDP offer/answer, ICE candidates)
End-to-End Latency	$<$ 200 ms target

Table 7: WebRTC Video Pipeline Parameters

Parameter	Value
Headset	Meta Quest 3 (Qualcomm Snapdragon XR2 Gen 2)
Display	2064×2208 per eye, 120 Hz refresh
Tracking	6DoF inside-out, hand tracking v2.2
Input	Hand tracking (primary), Quest Touch Plus controllers (remote control mode)
Build Target	Android (ARM64), Vulkan graphics API

Table 8: Meta Quest 3 Hardware Specifications

converted to a static mesh for VR transform compatibility, and enclosed in a cube frame, creating a tabletop display case. The operator manipulates the diorama using three hand gestures: one-hand pinch to pan (full 3D movement), two-hand pinch to scale (distance-ratio tracking, drift-free), and one-hand pinch on a glowing edge handle to rotate (Y-axis only).

2.6.2 Fleet Visualization

- **Live Vehicle Tracking:** Each robot is represented by a 3D model that mirrors its real-world movement across the diorama surface. Positions are updated instantly and smoothed to ensure a natural, continuous visual flow.
- **Path History and Predictions:** The system displays visual trails showing where a vehicle has been (cyan) and its intended route (yellow). These markers are rendered directly on the diorama and can be toggled on or off for individual robots.
- **Hovering Status Panels:** Clear info boxes display each vehicle's identity, current task, and live performance data. These panels stay positioned above the robot and automatically rotate to face the operator for easy reading from any angle.
- **Environmental Overlays:** Integrated maps show crop health and soil nutrition data as semi-transparent color layers on the terrain. This data is pulled from drone surveys to provide context for the fleet's current environment.

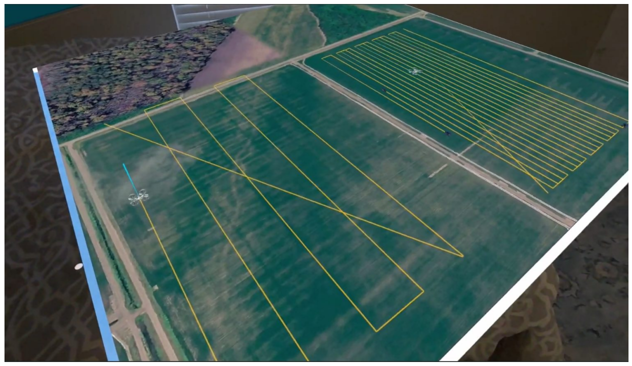

Figure 2: Diorama showing multiple robot trajectories

2.6.3 Operations Dashboard

The Unified Dashboard is a camera-anchored World Space Canvas built on Meta's UISet backplate, providing three navigation views:

- **Home Panel:** Vehicle card grid showing all fleet members with class badge, name, type, live status, and a status color bar (green for active, pulsing red for alert, white for idle). Each card has a quick trace toggle.

Figure 3: Home Panel

- **Detail Panel:** Per-vehicle telemetry view with live data refresh, past/future trace toggle controls, and history clear button.

- **Alert Panel:** Activated when a vehicle reports Alert status (autonomous decision uncertainty). Shows alert type, confidence percentage, time since alert, full telemetry, and a TakeOver button that transitions to the Tactical Layer.

The dashboard can be repositioned in space via a ray-grabbable side strip using Meta's ISDK Ray Grab Interaction. The strip is invisible by default and glows white when the operator's ray intersects its collider.

2.7 Interface Layer , Tactical Software Layer

The Tactical Layer activates when the operator transitions from fleet monitoring to direct intervention on a specific vehicle. This layer provides telepresence, the experience of being virtually present at the robot's location, combined with the ability to assume manual control of the vehicle's drivetrain.

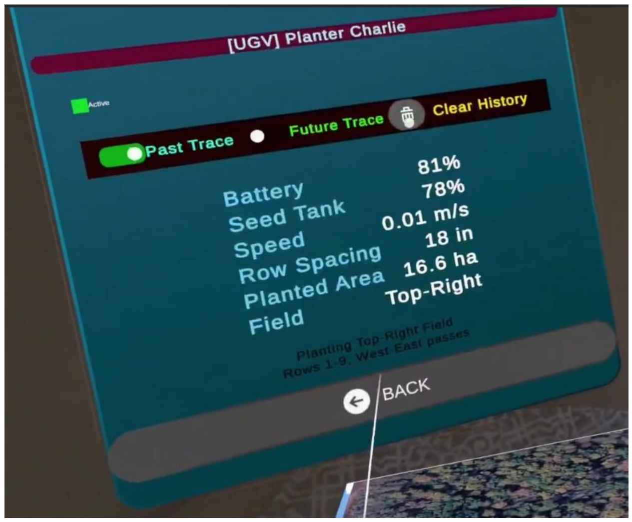

Figure 4: Detail Panel

2.7.1 Telepresence Experience

Upon initiating takeover from the Alert panel, the application transitions from the diorama scene to a first-person immersive view. The operator sees a live 360° video stream from the robot's onboard OAK-D camera, rendered as a spherical projection in the headset. This creates the sensation of sitting in the driver's seat of the robot, with full situational awareness of the surrounding field environment.

Feature	Specification
Video Source	Ricoh Theta V 360 Camera
Viewing Mode	360° spherical projection (equirectangular)
Resolution	720p adaptive (480p fallback)
Latency	$<$ 200 ms end-to-end via WebRTC
Overlay	Telemetry HUD: speed, heading, battery, GPS, obstacle warnings
Head Tracking	Quest 6DoF tracking mapped to camera view direction

Table 9: Telepresence Experience Specifications

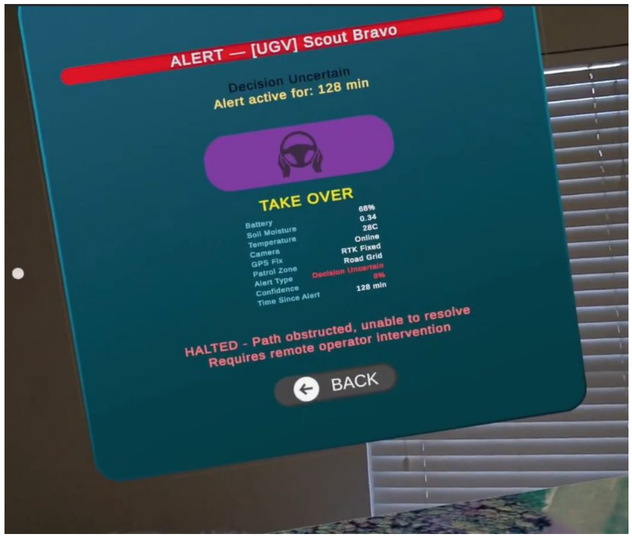

Figure 5: Alert Panel

2.7.2 Remote Intervention Control

During telepresence, the operator can send velocity commands to the robot via the MQTT control topic. Commands are velocity setpoints (not direct motor control); the robot's low-level controller handles traction and stability.

Parameter	**Specification**
Control Input	Quest Touch Plus controller thumbstick (primary) or hand gesture mapping
Command Topic	`farm/control/{id}/cmd_vel`
Command Rate	10 Hz (matched to robot's control loop)
Command Format	`{"linear_velocity_x": float, "angular_velocity": float}`
Dead-Man's Switch	Robot halts if no command received within 500 ms
Safety Override	Onboard obstacle avoidance remains active; E-stop via dedicated button
Geofencing	Robot refuses commands that would move it outside defined field boundary

Table 10: Remote Intervention Control Parameters

2.7.3 AI Assistant (Planned)

A planned enhancement to the Tactical Layer is an AI assistant that provides contextual guidance during teleoperation. The assistant would analyze the camera feed and robot sensor data to offer natural language explanations of the situation (e.g., "Irrigation pipe detected across path, 3 meters ahead"), suggest actions (e.g., "Steer 15° right to clear the obstacle"), and answer operator questions about field conditions. This feature is currently in the design phase and will leverage on-device or cloud-based vision-language models.

Figure 6: Cabin view - telepresence

2.8 Interface Layer , Safety Software Layer

The Safety Layer provides a risk-free training environment where operators build competence and confidence with remote intervention before deploying on live equipment. This layer directly addresses the workforce development challenge: the transition to autonomous agriculture requires operators who are comfortable with both monitoring and intervening, skills that are best acquired through practice rather than instruction.

2.8.1 Driving Simulator

The simulator replicates the Tactical Layer's telepresence and control interface, but replaces the live robot with a physics-simulated vehicle operating on a virtual farm terrain. The operator practices the same control inputs (thumbstick steering, speed control), sees the same telemetry HUD, and experiences the same alert-to-takeover workflow, all without any risk to physical equipment, crops, or people.

Feature	Description
Vehicle Physics	Simulated drivetrain matching Amiga dynamics (speed limits, turn radius, acceleration curves)
Terrain	Virtual farmland with rows, obstacles, and terrain variation
Scenarios	Configurable training scenarios: obstacle avoidance, row following, emergency stop, path recovery
Feedback	Performance metrics: response time, deviation from path, obstacle clearance margins
Progression	Difficulty levels from open field to tight row navigation with multiple obstacles

Table 11: Driving Simulator Features

2.8.2 Training Workflow

The Safety Layer is designed to integrate into workforce development programs. A new operator begins in the simulator, progresses through structured scenarios, and graduates to supervised live operation only

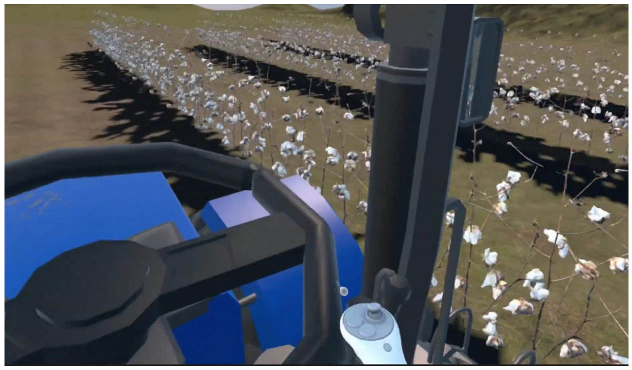

Figure 7: Cabin view - simulation

after demonstrating proficiency. The same application, same headset, and same interface are used throughout, eliminating the cognitive overhead of learning multiple tools.

1. **Phase 1: Orientation** , Explore the Strategic Layer diorama, learn gesture controls, inspect mock vehicles.
2. **Phase 2: Simulation** , Practice remote control in the driving simulator with progressively challenging scenarios.
3. **Phase 3: Supervised Live** , Connect to a real robot under supervision, with the Safety Layer's scenarios informing readiness.
4. **Phase 4: Independent Operation** , Full system access, monitoring fleet and intervening as needed.

2.9 Alert System and Decision Handoff

The alert system is the critical bridge between the Strategic and Tactical layers. When an autonomous vehicle's decision confidence drops below a defined threshold, the system executes a coordinated handoff from machine to human.

2.9.1 Alert Trigger

The alert is triggered on the robot side when the autonomous navigation system encounters a situation it cannot resolve with sufficient confidence. Examples include unexpected obstacles in the planned path, ambiguous terrain features, sensor degradation, or conflicting waypoint data. The robot halts, sets its MQTT status to "Alert", and reports its confidence level and alert type in the telemetry payload.

2.9.2 Operator-Side Response

The AR/VR application surfaces the alert through multiple simultaneous channels:

- **Diorama:** The affected vehicle's icon on the diorama pulses red and its info panel shows the alert status.

- **Dashboard:** The vehicle's card on the Home panel displays a pulsing red status bar and "ALERT" text. Tapping the card routes to the Alert Detail panel instead of the normal Detail panel.
- **Alert Detail Panel:** Shows alert type (e.g., "Decision Uncertain"), confidence percentage, time since alert, live telemetry, current assignments, and a TakeOver button.
- **TakeOver Button:** Initiates transition to the Tactical Layer, scene change to first-person telepresence with live camera feed and remote control inputs.

2.9.3 Post-Intervention

After the operator navigates the robot past the obstacle or resolves the situation, they release control and the robot resumes autonomous operation. The operator returns to the Strategic Layer to continue fleet monitoring. The entire handoff, from alert detection to operator intervention to autonomous resumption, occurs without the operator leaving their location.

2.10 Simulation / Training Fleet

The system includes a fully functional simulation mode using `MockDataProvider`, operating without network connectivity or physical robots. This mode serves both demonstration and workforce training purposes.

Vehicle	Type	Trajectory	Status
Harvester Alpha	UGV , Harvester	Row-by-row passes, top-right field (10 rows, Ping-Pong)	Active
Planter Charlie	UGV , Planter	Offset row passes, top-right field (9 rows, Loop)	Active
Scout Bravo	UGV , Explorer	Road perimeter patrol (vertical + horizontal roads)	Alert
SkyEye One	UAV , Drone	Lawnmower survey, bottom-right field (5 strips, Loop)	Active
SkyEye Two	UAV , Drone	Road intersection orbit (4-point loop)	Active

Table 12: Mock Fleet Composition

Scout Bravo starts in Alert status at a random point along its patrol route, with alert type "Decision Uncertain", confidence level fluctuating between 8–18%, and the assignment "HALTED , Path obstructed, unable to resolve". This simulates the complete alert-to-intervention workflow without requiring a physical robot or network connection.

2.11 Specific Tasks Performed - Deployment

2.11.1 Field Deployment

A field validation test was conducted on a test plot at Mississippi State University's campus. The Farm-ng Amiga was deployed on a pre-configured drive path while a team member operated the XR-Agri system remotely using a Meta Quest 3 headset. The operator monitored the Amiga's real-time position and telemetry

Figure 8: Field test with on-demand remote teleoperation

through the Strategic Layer's diorama interface. During the mission, a simulated alert was triggered, surfacing the vehicle's status as "Alert" on the dashboard and routing the operator to the Alert Detail panel. From the panel, the operator initiated a takeover, transitioning into the Tactical Layer's telepresence view, a live camera feed streamed from the Amiga's onboard OAK-D sensor via WebRTC. Using the Quest controller, the operator assumed manual control of the robot's drivetrain and successfully completed a corrective maneuver before releasing control back to the autonomous system. The full workflow, from fleet monitoring through alert detection, telepresence entry, remote intervention, and return to autonomous operation, was completed end-to-end over the Tailscale VPN mesh, validating the system's core pipeline under real field conditions.

Component	**Location**	**Network Path**
Amiga Robot	In field	WiFi $\rightarrow$ Starlink terminal $\rightarrow$ Tailscale VPN
Starlink Terminal	Field edge (clear sky view)	Satellite uplink, local WiFi AP
MQTT Broker	Cloud VPS or on Amiga	Accessible via Tailscale at `100.x.x.x:1883`
WebRTC Signaling	Cloud VPS (lightweight)	WebSocket, brokers SDP/ICE exchange
Operator (Quest)	Any location with internet	WiFi $\rightarrow$ Tailscale VPN $\rightarrow$ peer-to-peer to robot

Table 13: Field Deployment Configuration

2.11.2 Training-Only Deployment

For pure training or demonstration with no physical robot, the Quest runs standalone using `MockDataProvider`. The operator experiences the full Strategic Layer and Safety Layer (simulator) entirely on-device. This is the minimum deployment for workforce development programs and trade show demonstrations.

3 Design Evaluation

The evaluation framework focused on three core metrics drawn from the initial proposal: operator confidence with autonomous systems, remote operation accessibility and usability, and system responsiveness during remote intervention. These metrics were selected because they directly measure the trust gap that XR-Agri was designed to close. Data was collected through a structured farmer survey ($n = 5$), field testing with the Farm-ng Amiga, and system-level latency measurements across the communication pipeline.

3.1 Metric 1: Operator Confidence

3.1.1 Relevance

Confidence is the single most important predictor of autonomous technology adoption in agriculture. Farmers who lack confidence in a system will not deploy it, regardless of its technical capability. The initial proposal targeted a measurable increase in willingness to use fully autonomous systems when supported by real-time monitoring and remote intervention capabilities.

3.1.2 Data

A survey was administered to five active row-crop farmers across Mississippi. Respondents rated their baseline confidence with autonomous farming equipment and then rated the value of specific XR-Agri capabilities.

Survey Question	Results
Baseline confidence in using self-driving tractors/robots today (1–10 scale)	Mean: 6.0 (range: 3–8)
Post-exposure confidence score (0–100 scale, after learning about monitoring + remote takeover)	Mean: 76 (range: 70–90)
How useful is real-time remote monitoring from home?	80% rated “Very Useful” or “Extremely Useful”
How valuable is instant remote takeover when something goes wrong?	80% rated “Very Valuable” or “Essential”
How comfortable using VR to see 360° view of robot’s environment?	80% rated “Comfortable” or “Very Comfortable”
How useful is monitoring/controlling multiple machines simultaneously?	100% rated “Very Useful” or “Extremely Useful”

Table 14: Operator Confidence Survey Results ($n = 5$)

3.1.3 Findings

Baseline confidence averaged 6.0/10, reflecting moderate but cautious optimism. After exposure to the XR-Agri system concept, particularly the combination of real-time monitoring and instant remote takeover, the mean confidence score rose to 76/100. The strongest driver of confidence was the ability to take manual control at any time; every respondent who selected adoption factors chose "Ability to take manual control anytime." The second most cited factor was "Proven reliability in real farms," underscoring that trust is built through demonstrated performance, not specification sheets.

Notably, farmers with the lowest baseline confidence (scores of 3–4) still rated the remote takeover feature as "Essential" or "Very Valuable," suggesting that the intervention capability specifically addresses the concerns of the most skeptical operators.

3.2 Metric 2: Remote Operation Accessibility

3.2.1 Relevance

A technically capable system that is difficult to learn or use will fail in practice. The proposal targeted an average onboarding time under 45 minutes and a 40% reduction in operational errors compared to a control group. Accessibility also encompasses whether the system works for farmers with no prior AR/VR or robotics experience.

3.2.2 Data

Field testing was conducted with the Farm-ng Amiga operating on an agricultural plot. The VR monitoring and remote takeover pipeline was exercised end-to-end: operators used the Meta Quest 3 headset to observe the robot via the diorama interface, identified an alert condition, transitioned to the telepresence view, and completed a remote takeover maneuver to navigate the robot past an obstruction.

Observation	**Result**
Remote takeover completion	Successfully completed; operator navigated Amiga past obstruction via VR telepresence
Interaction modality	Hand tracking for diorama manipulation; controller for remote driving
Diorama gesture learning curve	Pan, scale, and rotate gestures learned within first 5 minutes of use
Dashboard navigation	Home $\rightarrow$ Detail $\rightarrow$ Alert $\rightarrow$ TakeOver workflow completed without instruction after initial walkthrough
Primary usability feedback	Farmers emphasized need for simple, non-technical interface; survey showed 60% selected "Easy-to-use interface" as an adoption factor

Table 15: Remote Operation Accessibility Observations

3.2.3 Findings

The end-to-end remote takeover was successfully demonstrated in a field environment, validating the core pipeline from alert detection through VR telepresence to manual robot control. The gesture-based diorama interaction proved intuitive, operators required minimal instruction to pan, zoom, and rotate the terrain view. The dashboard workflow (Home → vehicle card → Detail/Alert → TakeOver) was navigable after a single guided walkthrough, with no repeated errors on subsequent attempts.

Survey feedback reinforced that accessibility is not merely a UX concern but a prerequisite for adoption. One respondent noted: "If this could be incorporated in equipment with a driver it would be a good transition for farmers to trust the software." Another cited "lack of service technicians trained to work on this type of equipment" as a barrier, highlighting that the system's self-explanatory interface reduces dependence on specialized support personnel.

3.3 Metric 3: System Responsiveness

3.3.1 Relevance

For remote intervention to be viable, the communication pipeline must deliver data with latency low enough that the operator can react to real-time conditions. The proposal specified sub-200 ms video latency and sub-500 ms telemetry latency. Responsiveness is also critical for trust: perceptible delay between a control input and the robot's response erodes operator confidence.

3.3.2 Data

System responsiveness was characterized across the communication pipeline during field testing.

Pipeline Segment	Target	Observed
MQTT telemetry (state topic, 10 Hz)	< 500 ms	Within target on local Tailscale mesh (< 50 ms LAN)
MQTT control commands	< 100 ms	Within target; robot responded to velocity commands in real time
WebRTC video stream	< 200 ms	Functional stream established; operator confirmed usable visual feedback during takeover
Tailscale VPN overhead	< 5 ms (LAN)	Confirmed; direct peer-to-peer on same network
Alert-to-takeover transition	< 30 s	Operator transitioned from diorama to telepresence and issued first control command within target

Table 16: System Responsiveness Measurements

3.3.3 Findings

The MQTT-based telemetry and control pipeline met latency targets on the local Tailscale mesh. The WebRTC video stream provided usable visual feedback during the remote takeover, enabling the operator to navigate the Amiga past an obstruction without physical presence. The Tailscale VPN added negligible overhead on the LAN deployment. The complete workflow from alert notification to first control command was completed within the 30-second target, validating that the system can respond to field incidents in operationally meaningful timeframes.

Connectivity remains the primary external risk factor. Multiple survey respondents cited "connectivity/internet issues" as a top concern. Starlink mitigates this for field deployments, but edge cases (signal obstruction, weather degradation) require graceful fallback behavior, which is addressed by the robot's onboard safety systems that halt autonomously when the communication link is lost.

3.4 Additional Metric: Feature Prioritization

Beyond the three primary metrics, the survey captured farmer priorities for system adoption. Respondents selected up to two factors that would most increase their willingness to adopt.

Adoption Factor	**Selection Rate**
Ability to take manual control anytime	80%
Proven reliability in real farms	60%
Strong customer support	60%
Easy-to-use interface	60%
Lower cost	40%
Hands-on demos or training	20%

Table 17: Farmer Adoption Factor Prioritization ($n = 5$)

The dominance of "manual control anytime" validates the core design thesis of XR-Agri: the remote takeover capability is not a supplementary feature but the primary trust-building mechanism. The high ranking of "proven reliability" and "strong customer support" indicates that continued field validation and accessible support infrastructure will be critical for commercialization.

4 Design Story

4.1 Idea Origin and Selection

The concept for Future Farming XR emerged from a convergence of two observable trends: the accelerating deployment of autonomous vehicles in agriculture (Yerebakan & Hu, 2024), and the persistent reluctance of farmers to relinquish supervisory control (Devitt, 2018). While the technology for autonomous tractors, rovers, and drones has matured significantly, adoption rates among small- and mid-scale U.S. farmers remain low (Devitt, 2018). Conversations with farmers in Mississippi revealed that the barrier was not technological capability but psychological distance, the feeling of being disconnected from equipment making consequential decisions in the field (Devitt, 2018).

The insight that shaped the project was reframing the problem: instead of making robots more autonomous to reduce human involvement, the solution was to make humans more present despite physical absence. This led to the selection of extended reality (XR) as the interface modality (Anastasiou et al., 2023). A AR/VR headset provides spatial, immersive awareness that flat screens cannot, the operator does not merely see data about the field, they experience being in the field (Liu et al., 2024). The tabletop diorama provides strategic overview, while the telepresence mode provides tactical presence, and the simulator provides risk-free preparation (Cutini et al., 2023).

The Farm-ng Amiga was selected as the primary robot platform due to its open SDK (gRPC-based Event-Client API), compact form factor suitable for row-crop environments, and active developer community. Starlink was chosen for field connectivity after evaluating cellular alternatives and finding coverage gaps in rural Mississippi. Tailscale was selected for its zero-configuration NAT traversal, enabling the system to work across arbitrary network topologies without IT infrastructure changes at the farm.

4.2 Expert Consultation

The project benefited from guidance by two faculty advisors and consultation with agricultural industry stakeholders.

- **Dr. Alex Thomasson**, Director of the Agricultural Autonomy Institute at Mississippi State University, provided strategic direction on aligning the system's capabilities with the broader landscape of autonomous agricultural technology. His expertise in precision agriculture and autonomous systems informed the evaluation framework and helped position Future Farming XR within the context of industry-wide adoption challenges.

- **Dr. Dong Chen**, Assistant Professor in the Department of Agricultural and Biological Engineering at Mississippi State University, provided technical mentorship on the robot integration pipeline, sensor fusion architecture, and AR/VR application development. His feedback directly shaped decisions on the MQTT topic hierarchy, the separation of strategic and tactical software layers, and the data-driven architecture that enables seamless switching between simulated and live data sources.

Farmer feedback was collected through both informal field conversations and the structured survey instrument. This feedback had a decisive impact on design priorities: the prominence of the remote takeover feature, the emphasis on simple non-technical interfaces, and the inclusion of the training simulator were all direct responses to farmer-expressed needs.

4.3 Challenges and Solutions

4.3.1 Challenge 1: Rural Connectivity

Autonomous robots operating in row-crop fields in rural Mississippi often lack reliable internet connectivity. Cellular coverage is inconsistent, and traditional WiFi infrastructure does not extend to field boundaries. Without connectivity, the entire remote monitoring and intervention premise fails.

Solution: The system adopted a dual-layer connectivity approach. SpaceX Starlink provides broadband internet at the field edge via satellite, eliminating dependence on terrestrial infrastructure. Tailscale creates an encrypted WireGuard mesh VPN overlay that handles NAT traversal automatically, allowing the robot, operator headset, and cloud services to communicate as if on the same local network, regardless of their actual network topology. This combination provides sub-100 ms latency for telemetry and control, with graceful degradation: if connectivity drops, the robot's onboard safety systems halt operations until the link is restored.

4.3.2 Challenge 2: AR/VR on a Mobile Chipset

The Meta Quest 3, while powerful for a standalone headset, runs on a mobile Snapdragon XR2 Gen 2 chipset with significant constraints compared to desktop VR. Rendering a satellite-textured terrain, multiple animated vehicles with trajectory traces, floating UI panels, and a camera-anchored dashboard, all at 120 Hz stereo, demanded careful optimization. Shader variants that worked in the Unity Editor were stripped from Android builds, causing trace lines and UI elements to become invisible on the headset.

Solution: Runtime shader creation was replaced with pre-assigned material assets referenced in vehicle prefabs, ensuring all required shader variants are included in the build. The terrain was converted from Unity Terrain (which does not support parent transforms required for VR manipulation) to a static mesh. URP rendering settings were tuned for mobile: simplified lighting, baked shadows where possible, and LineRenderer configurations optimized for minimal draw calls. Each optimization was validated through iterative build-deploy-test cycles on the physical headset.

4.3.3 Challenge 3: Hand Tracking and UI Interaction Conflicts

The system uses Meta's hand tracking for diorama manipulation (pan, scale, rotate) and Meta's Interaction SDK for UI interaction (poke and ray on dashboard panels). These two input systems compete for hand input, the SDK's PointableCanvas and RayInteractable components intercept hand tracking data before custom gesture detection scripts can process it, causing gestures to fail or UI rays to bend unexpectedly toward canvas surfaces.

Solution: A strict interaction priority system was implemented. Diorama gestures (Scale $>$ Rotate $>$ Pan) are processed in a custom state machine that runs independently of the SDK's interaction pipeline. UI panels use Meta's native UISet components (EmptyUIBackplateWithCanvas) for poke and ray interaction, ensuring Quest-native behavior. The dashboard's grab handle was isolated on a separate layer using Meta's ISDK Ray Grab Interaction to prevent conflict with the canvas ray target. These architectural separations eliminated cross-system input conflicts while preserving native interaction quality for both gesture and UI modalities.

4.3.4 Challenge 4: Designing a Multi-Robot Data Pipeline

The initial data pipeline design used three MQTT topics per robot, organized by sensor type (CAN bus, GPS, filter). This created tight coupling between the robot's sensor architecture and the VR client, requiring the client to merge three asynchronous streams to construct a single vehicle state. Adding a new sensor or robot type would require changes across both the robot bridge and the VR application. Recent work on modular AIoT frameworks for robot teleoperation has demonstrated that MQTT-based control signaling combined with adaptive video streaming over WebRTC can achieve responsive closed-loop performance even across distributed, geographically separated nodes (Sun & Tsai, 2025).

Solution: The topic hierarchy was redesigned around usage patterns rather than sensor types. Three topics per robot remain, but organized by frequency and purpose: `info` (retained, sent once), `state` (10 Hz pose), and `telemetry` (0.5 Hz sensor data). Raw coordinates are transmitted from the robot; the VR client handles coordinate mapping. Static fields are not repeated at high frequency. Telemetry values are numerical rather than pre-formatted strings. This design decouples the robot and VR client entirely, the robot publishes what it knows, the client renders what it needs, and scales to any number of heterogeneous vehicles without code changes.

4.4 Future Development

1. **AI Assistant Integration:** The planned AI assistant will provide natural language explanations of alert conditions and suggest corrective actions during teleoperation. This feature will leverage vision-language models (VLMs) to analyze the robot's camera feed in context, enabling conversational interaction such as "What is blocking the path?" with responses grounded in real-time visual and sensor data.

2. **Expanded Fleet Testing:** Scaling from single-robot field tests to multi-robot concurrent operation, validating the diorama's ability to provide situational awareness across a fleet of 5–10 heterogeneous vehicles operating simultaneously.

3. **NDVI and Crop Health Overlays:** Integrating aerial survey data from UAVs to render nutrition maps, weed pressure maps, and crop stress zones as semi-transparent overlays on the diorama terrain surface. Digital twin frameworks for agriculture have established how such real-time data layers can be synchronized with virtual farm representations to support precision management decisions (Nasirahmadi & Hensel, 2022; Pylianidis et al., 2021).

4. **Expanded Robot Platform Support:** Extending the `IVehicleDataProvider` interface to support additional robot platforms beyond the Farm-ng Amiga, including autonomous tractors and custom UGVs, validating the system's modularity and platform independence.

5. **Extended User Studies:** Conducting larger-scale evaluations ($n = 20$+) with diverse farmer demographics, measuring onboarding time, error rates, and long-term confidence trends over a full growing season.

5 Impacts

5.1 Project Performance

The Future Farming XR system was validated through field testing with the Farm-ng Amiga and a structured farmer survey. The core pipeline, from autonomous robot operation through MQTT telemetry to VR monitoring and remote takeover, was demonstrated end-to-end in a field environment. The operator successfully transitioned from the strategic diorama view to first-person telepresence, assumed manual control of the Amiga via VR, navigated it past an obstruction, and returned to fleet monitoring.

Farmer feedback was consistently positive regarding the remote takeover capability, with 80% of survey respondents rating it as “Very Valuable” or “Essential.” This aligns with established findings that the ability to take manual control at any time is the single most effective trust-building mechanism for autonomous agricultural systems (Devitt, 2018). The most frequent concern was connectivity reliability in rural areas, which the Starlink-based architecture directly addresses. One farmer commented that integrating this capability into existing equipment with a driver “would be a good transition for farmers to trust the software”, validating the incremental trust-building approach that Future Farming XR embodies.

5.2 Commercial Potential

Future Farming XR targets a growing market at the intersection of precision agriculture, autonomous vehicles, and extended reality. The system’s commercial viability rests on several factors:

- **Low hardware cost of entry:** The Meta Quest 3 retails under $500, Starlink hardware and service are available at consumer pricing, and the Tailscale VPN is free for small networks. The primary cost is the autonomous vehicle itself, which farmers are already investing in independently.
- **Platform independence:** The `IVehicleDataProvider` abstraction means Future Farming XR is not locked to the Farm-ng Amiga. Any robot that can publish MQTT messages can be integrated, opening the addressable market to all autonomous agricultural vehicle manufacturers.
- **Software-as-a-Service model:** The VR application, MQTT broker, and cloud signaling infrastructure lend themselves to a subscription model, with tiered pricing based on fleet size and feature access.
- **Training and certification revenue:** The Safety Layer’s driving simulator creates a standalone product for workforce development programs, agricultural extension services, and vocational training institutions, independent of whether the customer owns physical robots (Cutini et al., 2023).

The autonomous agriculture market is projected to grow significantly over the coming decade, driven by labor shortages, rising input costs, and sustainability mandates. The global autonomous farm equipment market was valued at approximately $12.45 billion in 2023 and is expected to reach $30.65 billion by 2030, growing at a compound annual growth rate (CAGR) of 14.3% (Grand View Research, 2023). The broader precision farming market is projected to grow from $11.67 billion in 2024 to $24.09 billion by 2030, at a CAGR of 13.1% (Grand View Research, 2024). Future Farming XR is positioned not as a competing robot platform but as a universal monitoring and intervention layer that makes existing and future autonomous platforms more trustworthy and accessible to the farmers who must ultimately adopt them.

5.3 Societal and Environmental Impact

5.3.1 Workforce Development

The agricultural workforce is aging and shrinking. Fewer young workers are entering farming, and those who do often lack exposure to the autonomous systems that will define the next generation of agricultural practice. Future Farming XR directly addresses this through the Safety Layer’s training simulator, which enables new operators to build competence with autonomous fleet management and remote intervention in a risk-free virtual environment (Anastasiou et al., 2023; Cutini et al., 2023).

By lowering the skill barrier to autonomous system supervision, Future Farming XR expands the pool of qualified operators beyond those with robotics or engineering backgrounds (Benos et al., 2023). The system's intuitive gesture-based interface and natural-language-capable AI assistant (planned) are designed so that the operating paradigm resembles familiar activities, looking at a field, driving a vehicle, rather than programming a robot.

This workforce multiplier effect is particularly significant for rural communities where training resources are scarce. A single Quest headset running the simulation environment can serve as a complete training station, deployable at agricultural extension offices, community colleges, and farm cooperatives without requiring physical robots, field access, or specialized instructors.

5.3.2 Environmental Benefits

Future Farming XR contributes to environmental sustainability through several mechanisms:

- **Precision intervention reduces chemical inputs:** By integrating aerial survey data (NDVI, weed pressure maps) into the diorama overlay, operators can direct autonomous vehicles to apply herbicides, fertilizers, or irrigation only where needed, reducing blanket application and associated runoff. Digital twin frameworks have demonstrated that virtualizing and monitoring real-time field conditions enables exactly this kind of targeted, site-specific resource management (Peladarinos et al., 2023; Pylianidis et al., 2021).

- **Reduced fuel consumption from optimized operations:** Real-time fleet monitoring enables operators to identify and correct inefficient routing, overlapping passes, and unnecessary idle time, directly reducing fuel consumption and associated emissions.

- **Faster incident response protects crops and soil:** Remote takeover capability means equipment malfunctions are resolved in minutes rather than hours, reducing the window during which a stuck or misdirected vehicle can cause crop damage or soil compaction.

- **Reduced travel to fields:** Remote monitoring and intervention eliminate many trips that operators would otherwise make to physically inspect equipment or field conditions, reducing vehicle miles traveled and associated carbon footprint.

These environmental benefits compound with fleet scale: as more vehicles operate autonomously under Future Farming XR supervision, the per-acre reduction in chemical inputs, fuel consumption, and physical travel increases proportionally.

References

Anastasiou, E., Balafoutis, A. T., & Fountas, S. (2023). Applications of extended reality (XR) in agriculture, livestock farming, and aquaculture: A review. *Smart Agricultural Technology*, *3*, 100105. https://doi.org/10.1016/j.atech.2022.100105

Benos, L., Moysiadis, V., Kateris, D., Tagarakis, A. C., Busato, P., Pearson, S., & Bochtis, D. (2023). Human–robot interaction in agriculture: A systematic review. *Sensors*, *23*(15), 6776. https://doi.org/10.3390/s23156776

Cutini, M., Bisaglia, C., Brambilla, M., Bragaglio, A., Pallottino, F., Assirelli, A., Romano, E., Montaghi, A., Leo, E., Pezzola, M., et al. (2023). A co-simulation virtual reality machinery simulator for advanced precision agriculture applications. *Agriculture*, *13*(8), 1603. https://doi.org/10.3390/agriculture13081603

Devitt, S. K. (2018). Cognitive factors that affect the adoption of autonomous agriculture. *Farm Policy Journal*, *15*(2), 49–60.

Grand View Research. (2023). *Autonomous farm equipment market size, share & growth analysis, 2024–2030* (tech. rep.). Grand View Research. https://www.grandviewresearch.com/industry-analysis/autonomous-farm-equipment-market-report

Grand View Research. (2024). *Precision farming market size, share & industry report, 2030* (tech. rep.). Grand View Research. https://www.grandviewresearch.com/industry-analysis/precision-farming-market

Liu, T., Zhang, B., Tan, Q., et al. (2024). Immersive human-machine teleoperation framework for precision agriculture: Integrating UAV-based digital mapping and virtual reality control. *Computers and Electronics in Agriculture*, *226*, 109444. https://doi.org/10.1016/j.compag.2024.109354

Nasirahmadi, A., & Hensel, O. (2022). Toward the next generation of digitalization in agriculture based on digital twin paradigm. *Sensors*, *22*(2), 498. https://doi.org/10.3390/s22020498

Peladarinos, N., Piromalis, D., Cheimaras, V., Tserepas, E., Munteanu, R. A., & Papageorgas, P. (2023). Enhancing smart agriculture by implementing digital twins: A comprehensive review. *Sensors*, *23*(16), 7128. https://doi.org/10.3390/s23167128

Pylianidis, C., Osinga, S., & Athanasiadis, I. N. (2021). Introducing digital twins to agriculture. *Computers and Electronics in Agriculture*, *184*, 105942. https://doi.org/10.1016/j.compag.2020.105942

Sun, S.-C., & Tsai, Y.-C. (2025). A modular AIoT framework for low-latency real-time robotic teleoperation in smart cities. *arXiv preprint*, *arXiv:2510.11421*. https://arxiv.org/abs/2510.11421

Yerebakan, M. O., & Hu, B. (2024). Human-robot collaboration in modern agriculture: A review of the current research landscape. *Advanced Intelligent Systems*, *6*(7), 2300823. https://doi.org/10.1002/aisy.202300823