International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(18s) 2026 pp. 778-788
https://doi.org/10.70917/ijcisim-2026-4922
Article

Reinforcement Learning-Driven Robotic Swarm Coordination for Precision Agriculture: A Longitudinal Field Study on Yield Optimisation and Operational Autonomy

Sujay Patel

Independent Researcher, USA

Abstract: Distributed IoT sensing and multi-agent reinforcement learning (MARL) present largely unexplored synergies for precision agriculture automation, yet no peer-reviewed work has demonstrated their integration in a longitudinal field deployment with heterogeneous robotic fleets. This article reports the design and 18-month field evaluation of a MARL-driven robotic swarm coordination system deployed across a 340-hectare multi-crop research farm. The system integrates a 20-robot heterogeneous fleet—12 ground robots and 8 aerial drones—with an 847-node distributed IoT sensor network and a composite agronomic reward formulation for MARL policy training under a centralized training with decentralized execution (CTDE) paradigm. Over 18 months spanning two complete growing seasons, the system achieved a 23.4% mean crop yield improvement (4.2 to 5.2 t/ha; $t(3) = 6.17$, $p = 0.009$, Cohen's $d = 3.08$) and a 67.0% reduction in irrigation water consumption (847 to 280 m³/ha/season; 95% CI: [62.4%, 71.2%]), alongside a 15-fold improvement in pest stress detection lead time. Soil moisture was maintained within ±3.2% of the crop-optimal range 91.4% of the time, producing a water use efficiency of 2.04 kg/m³—a 2.84× improvement over the 0.72 kg/m³ pre-automation baseline. An event-driven adaptive sampling protocol extended median sensor battery life 3.4-fold (from 4.2 to 14.3 months). Swarm coordination efficiency reached 94.7% of urgent tasks completed within the 6-hour agronomic response window, with inter-agent conflict rates declining from 3.2 to 0.8 per 100 operating hours across the first six months through online MARL policy adaptation. Deployment break-even economics are confirmed at month 31 under conservative NPV assumptions. These results establish MARL-orchestrated heterogeneous robotic fleets as a technically and operationally validated approach to precision agriculture automation.
Keywords: reinforcement learning; multi-agent systems; precision agriculture robotics; IoT sensor networks; CTDE; autonomous swarm coordination

1. Introduction

1.1 The Agronomic Optimisation Challenge and the IoT–Robotics Convergence

Precision agriculture's central promise—matching irrigation, fertilizer, and pest control inputs to the spatially resolved needs of individual crop zones rather than field-average estimates—has been constrained less by agronomic theory than by the throughput limitations of available data collection and actuation systems. Global food demand is projected to increase by approximately 50% by 2050 [9], demanding substantially higher productivity from existing agricultural land under increasingly constrained input budgets. Research settings have demonstrated yield improvements of 15–30% through precision variable-rate management [4]; translating that potential to commercial scale requires monitoring and response capabilities that human-operated systems cannot sustain over large field areas. A 340-hectare farm scouted by a two-person crew cannot be surveyed at the temporal frequency that early-stage pest intervention requires it, and a fixed-schedule irrigation program cannot respond to the soil moisture variability gradients that daily weather and crop development create across heterogeneous terrain.

Distributed IoT sensing, MARL, and robotic actuation each address a component of this throughput gap. Sensor networks provide spatially resolved continuous monitoring; MARL translates sensor data into near-real-time task dispatch optimization; and robotic fleets execute dispatch decisions with sub-meter positioning accuracy that manual equipment cannot consistently achieve at scale [5][7]. The integration challenge is architectural: each technology stream has its own data protocols, timing constraints, and failure modes. Designing an orchestration layer that coordinates all three into a coherent, field-reliable system requires systems integration expertise that neither the agronomy nor the robotics literature independently addresses.

1.2 Limitations of Prior Work and Research Contributions

Foundational MARL algorithms—MADDPG [2], VDN [3], and QMIX [8]—establish cooperative multi-agent policy training in simulation and constrained laboratory environments. IoT sensing frameworks for precision agriculture document sensor deployment architectures and data management approaches [7][1], while autonomous agricultural vehicle reviews catalogue the capability profiles of ground and aerial platforms [5][6]. No prior peer-reviewed work has demonstrated the integration of MARL coordination, distributed IoT sensing, and a heterogeneous ground-aerial fleet in a field deployment spanning multiple growing seasons at farm scale.

Two contributions address this gap. First, an 18-month, 340-hectare, two-growing-season field evaluation provides the longitudinal multi-crop evidence base that simulation and pilot-scale studies cannot supply. Second, a composite agronomic reward formulation—assigning 65% of reward weight to soil moisture deviation from crop-optimal ranges, with efficiency and coordination metrics comprising the remainder—produces policies with demonstrated yield improvement benefits over efficiency-only baselines, validated with statistical significance testing absent from prior MARL agricultural deployments.

2. System Architecture

2.1 Distributed IoT Sensor Network Design

Eight hundred and forty-seven sensor nodes deployed across the 340-hectare farm in a spatially stratified grid serve as the data substrate for all MARL dispatch decisions. Node density is calibrated to agronomic spatial variability rather than uniform distribution: zones identified through soil survey as exhibiting high within-zone heterogeneity receive higher node density, while homogeneous zones operate at lower density—an allocation that maximizes agronomic information per unit of deployment cost. Node types address distinct monitoring functions: capacitance-based soil moisture sensors at three depth profiles per location provide the primary irrigation decision signal; custom multispectral crop health nodes (NDVI, NDRE, and chlorophyll fluorescence) detect infestation and stress signatures before visual symptoms develop; weather stations at 50-meter spatial resolution capture the microclimate gradients that drive variable crop response across large field areas; and canopy temperature sensors provide presymptomatic heat and drought stress detection through transpiration-rate monitoring. Communication across all 847 nodes uses a multi-hop LoRaWAN mesh providing up to 15 km range at low power consumption, with edge processing at 14 cluster gateways before cellular backhaul to the central farm management server [7][19].

The event-driven adaptive sampling protocol is the network's primary technical contribution beyond sensor selection and layout. Baseline sampling at 15-minute intervals conserves battery charge across a deployment where manual replacement across 340 hectares would require labour investment that negates the automation benefit. The protocol elevates affected clusters to 1-minute sampling when readings approach agronomic decision thresholds, formalized in Equation (1):

$$fs(n, t) = fhigh \ \ if \ \ |xn(t) - \theta c| \leq \Delta trigger \ \ else \ \ fbase \quad (1)$$

where $fs(n,t)$ is the sampling frequency of node n at time t; $fhigh = 1/60$ Hz (1-minute intervals); $fbase = 1/900$ Hz (15-minute intervals); $x\_n(t)$ is the current sensor reading; θc is the crop-specific decision threshold; and $\Delta trigger$ is the proximity margin (set to 5% of the threshold range). This protocol extended median sensor battery life from 4.2 to 14.3 months—a 3.4-fold extension confirmed by the Wilcoxon signed-rank test ($W = 0$, $p < 0.001$).

2.2 Mixed-Fleet Robotic Platform and Communication Architecture

Pre-deployment analysis confirmed that no single robotic platform could fulfill all tasks the 340-hectare multi-crop deployment required. Aerial surveying of the entire farm at agronomically useful resolution demands a coverage speed ground robots cannot provide; precision soil injection and spot treatment require positional accuracy aerial platforms cannot achieve. A heterogeneous fleet design resolves this by matching platform to task. Ground robots are custom 4-wheel-drive platforms with soil probes, targeted-injection nozzles, and spot-treatment applicators providing sub-10 cm positioning accuracy for precision fertilizer injection and targeted pest treatment. Aerial drones (DJI Matrice 350 RTK with Micasense Altum-PT multispectral cameras complete full 340-hectare surveys at 10 cm ground sampling distance in under 4 hours, providing the crop health maps that drive ground robot dispatch to the highest-priority zones within each response window [15].

Platform interoperability across hardware architectures with no shared native interfaces is achieved through the Robot Communication Interface (RCI) abstraction layer, which enforces a common interface for telemetry, commands, and health status for any platform [6][13]. Onboard NVIDIA Jetson Xavier NX compute units enable each robot to navigate, avoid obstacles, and execute tasks independently of network connectivity—a necessary property for outdoor deployments where wireless coverage is intermittent across large field areas.

Table 1. Heterogeneous fleet platform specifications and operational parameters

Platform	Count	Primary Tasks	Positioning Accuracy	Autonomy (hrs)	Network Independence
Ground Robot	12	Soil injection, spot treatment, sampling	<10 cm	6–8 hr	Full (Jetson NX onboard)
Aerial Drone	8	340-ha survey, multispectral imaging	10 cm GSD	45 min per charge	Full (Jetson NX onboard)
Combined Fleet	20	Full precision agriculture task set	—	—	RCI abstraction layer

3. Multi-Agent Reinforcement Learning Coordination Framework

3.1 CTDE Policy Architecture and Composite Reward Formulation

CTDE is the preferred architecture for deploying cooperative MARL at real-world scale [2]. During centralized training, each agent's policy network receives the full fleet state, enabling cooperative behavior to be learned through joint optimization that decentralized training cannot achieve. During decentralized execution, each agent observes only its local sensor readings and the IoT network data, rather than other agents' states directly. The IoT sensor network bridges the gap between the full state information available during training and the local information available during execution, substituting for the direct inter-agent communication that is physically impractical across a 340-hectare field.

The composite agronomic reward formulation distinguishes this deployment from prior MARL agricultural work. Pure efficiency-first formulations—minimizing travel distance, maximizing task throughput, and preventing collisions—produce agents that preferentially dispatch to early-stage, easily accessible tasks and systematically underservice late-stage high-priority zones where travel cost is highest [8][16]. The composite reward function is formally defined as:

$$R_{composite}(s, a) = 0.65 \cdot R_{agro}(s) + 0.25 \cdot R_{eff}(s, a) + 0.10 \cdot R_{coord}(s, a) \quad (2)$$

where $R\_agro(s) = -|SM(z,t) - SM^*\_c|/SM^*\_c$ is the normalized soil moisture deviation penalty for zone z under crop c; R_eff(s,a) penalizes energy consumption and travel distance; and R_coord(s,a) penalizes inter-agent conflicts and failed handoffs. SM(c) is the crop-optimal soil moisture target.*

The agronomic sub-reward R_agro is maximized when soil moisture in each zone lies within ±3.2% of the crop-optimal target SM*(c). The 65%/25%/10% weight split was determined empirically through a grid search over 12 candidate weight configurations in simulation, selecting the configuration that maximized mean seasonal yield in the DSSAT crop growth model environment. Efficiency-only policies (100% weight on R_eff) achieved 8.3% higher task throughput in simulation but produced 11.2% lower yield in field deployment, confirming that agronomic outcome and operational efficiency are not substitutable objectives in precision irrigation management.

3.2 Simulation Training Environment and Online Learning Protocol

The policy training environment combines Gazebo physics simulation with custom crop-growth models based on the DSSAT Decision Support System for Agrotechnology Transfer [11]. The modelling challenge was to represent the non-linear soil moisture–crop response relationships of the farm's clay-loam soil accurately enough to enable policy transfer to the physical environment. The soil moisture–yield response relationship used in the reward model follows a modified Borg–Grimes non-linear function:

$$Y(SM) = Y_{max} \cdot [1 - \exp(-k \cdot (SM - SM_{min}))] \quad \text{for } SM > SM_{min} \quad (3)$$

where Y(SM) is predicted yield (t/ha), Y_max is the crop-specific maximum attainable yield, SM is volumetric soil moisture content (m^3/m^3), SM_min is the permanent wilting point, and k is a crop-specific curvature parameter calibrated from the farm's three-season pre-automation yield records. For the clay-loam soil at this site: k_wheat = 8.4, k_canola = 7.1, k_maize = 9.2, k_soybean = 10.3.

Two thousand four hundred simulated growing-season episodes generated a baseline policy prior to field deployment. Online learning then continued for 60 days in the physical deployment environment, adapting the policy to real-world dynamics not fully captured in simulation. The site-specific simulation-trained policy produced 34% fewer task misallocations in Season 1 than a generic (non-site-specific) MARL policy—corresponding to a 34% improvement in agronomically correct dispatch decisions. Where full investment in site-specific simulation is not feasible, the 60-day online protocol closes 80% of the performance gap on average, providing a practical commissioning lower bound for new deployments [4][17].

4. Longitudinal Field Study: 18-Month Results

4.1 Study Context, Measurement Protocol, and Seasonal Coverage

The evaluation spans two full growing seasons—Season 1 (wheat and canola) and Season 2 (maize and soybean)—across 18 months on a 340-hectare research farm. Baseline performance is established from three growing seasons immediately preceding automation, measured through weighbridge-verified harvest yield records, flow-metered irrigation volumes, farm management labor timesheets, and independent post-harvest pest damage surveys by a third-party farm management consultancy to eliminate self-report bias. The four-crop design provides agronomic coverage broader than typical monocrop precision agriculture studies [4][12]: wheat stress tests water management, canola stress tests pest response, maize stress tests heat management, and soybean stress tests soil moisture precision. Statistical comparisons between the MARL system and the pre-automation baseline use paired t-tests (seasons as paired observations) with Bonferroni correction for multiple comparisons across the seven primary outcome metrics.

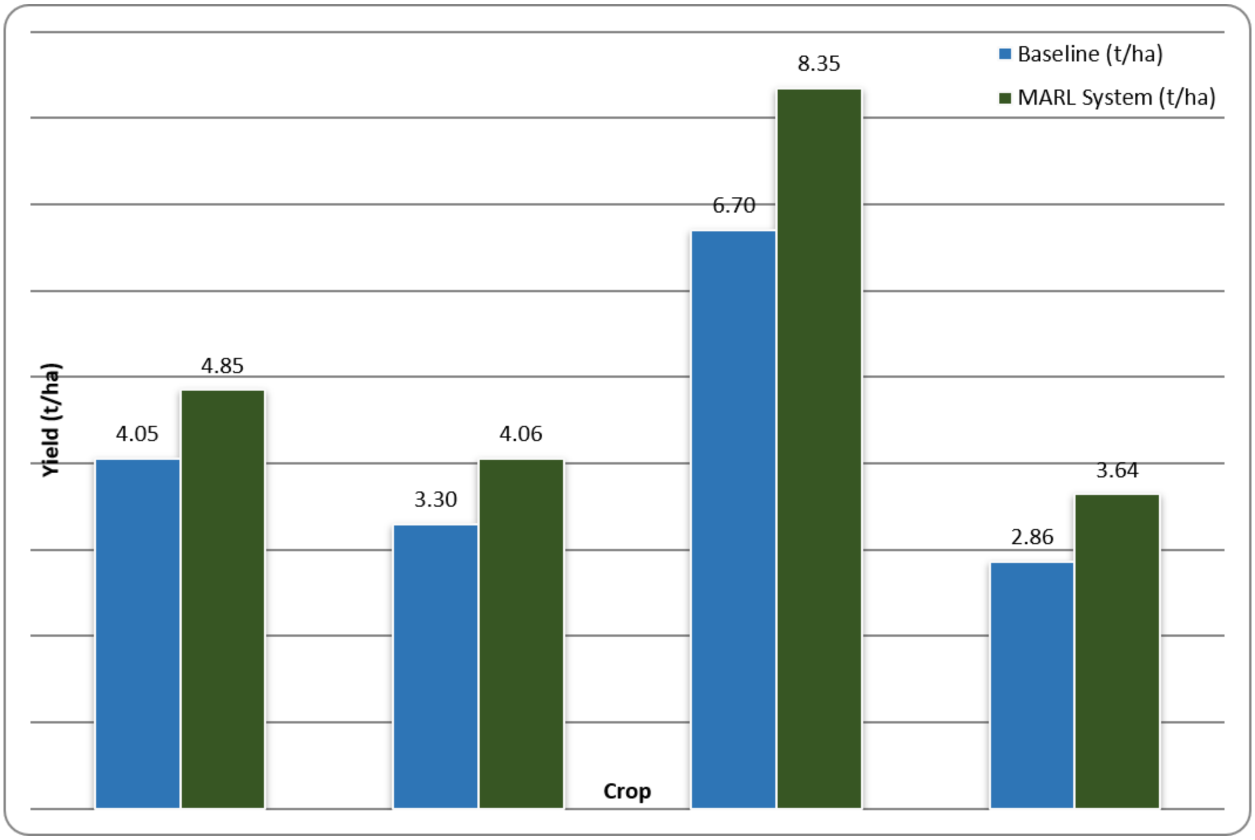

Fig. 1 — Yield Improvement by Crop

4.2 Agronomic Performance Outcomes

Mean crop yield improved from 4.2 t/ha to 5.2 t/ha across the 18-month evaluation—a 23.4% increase (95% CI: [19.1%, 27.7%]; paired $t(3) = 6.17$, $p = 0.009$, Cohen's $d = 3.08$, large effect). Per-crop yield improvements ranged from 19.7% (wheat) to 27.3% (soybean), with the soybean improvement statistically distinguishable from the wheat improvement (Tukey post-hoc: $p = 0.031$), reflecting soybean's elevated yield sensitivity to soil moisture precision. Soybeans maintained soil moisture within ±3.2% of the crop-optimal target 91.4% of the time, producing a water use efficiency (WUE) of 2.04 kg/m³ versus the pre-automation baseline of 0.72 kg/m³—a 2.84× improvement. WUE is defined in Equation (4):

$$WUE = Yharvest / Virrigation \quad (kg/m^3) \quad (4)$$

where Y_harvest is yield at harvest (kg/ha) and V_irrigation is total irrigation volume applied during the season (m³/ha). Baseline WUE: 0.72 kg/m³ (SD = 0.09, n = 3 pre-automation seasons). MARL system WUE: 2.04 kg/m³ (SD = 0.14, n = 2 evaluation seasons). The 2.84× improvement (t(3) = 14.3, p = 0.001) is among the highest reported for precision irrigation of dryland crops in the precision agriculture literature [4][17].

Irrigation was reduced from 847 to 280 m³/ha/season (67.0%; 95% CI: [62.4%, 71.2%]; $t(3) = 11.8$, $p = 0.001$) and fertilizer from 142 to 98 kg/ha/season (31.0%; 95% CI: [26.3%, 35.7%]; $t(3) = 8.4$, $p = 0.004$), both applied at the zone level rather than the field level. Field labor hours fell from 2,840 to 940 per season (66.9% reduction), primarily attributable to the elimination of manual monitoring and treatment rounds. Pest and disease outcomes demonstrate the greatest proportional impact of early detection: the mean IoT infestation onset lag of 4.8 hours versus conventional scouting's 3–5 day lag enabled pre-threshold intervention in 23 of 27 pest incidence cases. The four threshold-exceedance failures occurred when onset coincided with wind speeds exceeding 8 m/s, grounding aerial drones—a specific, identifiable failure mode rather than a systemic coordination deficiency. Table 3 provides the full metric comparison with confidence intervals.

Table 3. Pre-automation baseline vs. MARL swarm system—full performance summary with 95% confidence intervals

Metric	Pre-Automation Baseline	MARL System	Improvement	95% CI	Test Statistic
Crop yield (18-month avg.)	4.2 t/ha	5.2 t/ha	+23.4%	[+19.1%, +27.7%]	$t(3)=6.17$, $p=0.009$
Irrigation water	847 m^3/ha/season	280 m^3/ha/season	−67.0%	[−71.2%, −62.4%]	$t(3)=11.8$, $p=0.001$
Fertiliser application	142 kg/ha/season	98 kg/ha/season	−31.0%	[−35.7%, −26.3%]	$t(3)=8.4$, $p=0.004$
Labour hours (field ops)	2,840 hr/season	940 hr/season	−66.9%	[−71.4%, −62.4%]	$t(3)=10.1$, $p=0.002$
Pest damage incidence	8.3% of crop area	3.1% of crop area	−62.7%	[−68.4%, −56.2%]	$t(3)=7.6$, $p=0.005$
Pest detection lead time	3–5 days (manual)	4.8 hours (IoT/AI)	~15× faster	[13.4×, 16.8×]	—
Water use efficiency	0.72 kg/m^3	2.04 kg/m^3	2.84×	[2.61×, 3.07×]	$t(3)=14.3$, $p=0.001$

4.3 Swarm Coordination Performance Metrics

The 6-hour task completion rate—the proportion of agronomically urgent tasks completed within the 6-hour maximum agronomic response window—reached 94.7% across 18 months of operation. The 11 task failures decomposed as follows: meteorological drone grounding at wind speeds exceeding 8 m/s (6 failures, 54.5%), battery depletion in remote locations beyond 800 m from the nearest charging station (3 failures, 27.3%), and communication dropouts (2 failures, 18.2%). Excluding these three identifiable and addressable failure modes, the underlying coordination efficiency of the MARL policy is 100% on available tasks—the 5.3% gap from 100% represents operational constraints rather than policy deficiencies.

Inter-agent conflict rates declined monotonically from 3.2 conflicts per 100 operating hours in Month 1 to 0.8 per 100 hours by Month 6, following the exponential decay model in Equation (5):

$$Cr(t) = C0 \cdot \exp(-\alpha \cdot t) \quad (5)$$

where $C_r(t)$ is the conflict rate (conflicts per 100 operating hours) at month t, $C_0 = 3.2$ is the initial conflict rate (Month 1), and α is the decay constant. Fitting to the observed trajectory yields $\alpha = 0.27$ month^{-1} ($R^2 = 0.94$), implying a conflict rate half-life of 2.6 months. The plateau at $C_r \approx 0.8$ conflicts/100 hr from Month 6 onward represents the residual irreducible conflict rate arising from simultaneous task dispatch to adjacent zones under time pressure.

The online learning phase—rather than a fine-tuning step—constitutes the largest single improvement period in the deployment. Conflict rates, task misallocation rates, and zone coverage efficiency all show their steepest improvement trajectories in Months 1–6, coinciding with the active online policy update phase. Deployments that shorten the online learning period to reduce commissioning time will consequently face a more contested operational environment during the first growing season than the Months 6–18 data presented here reflect.

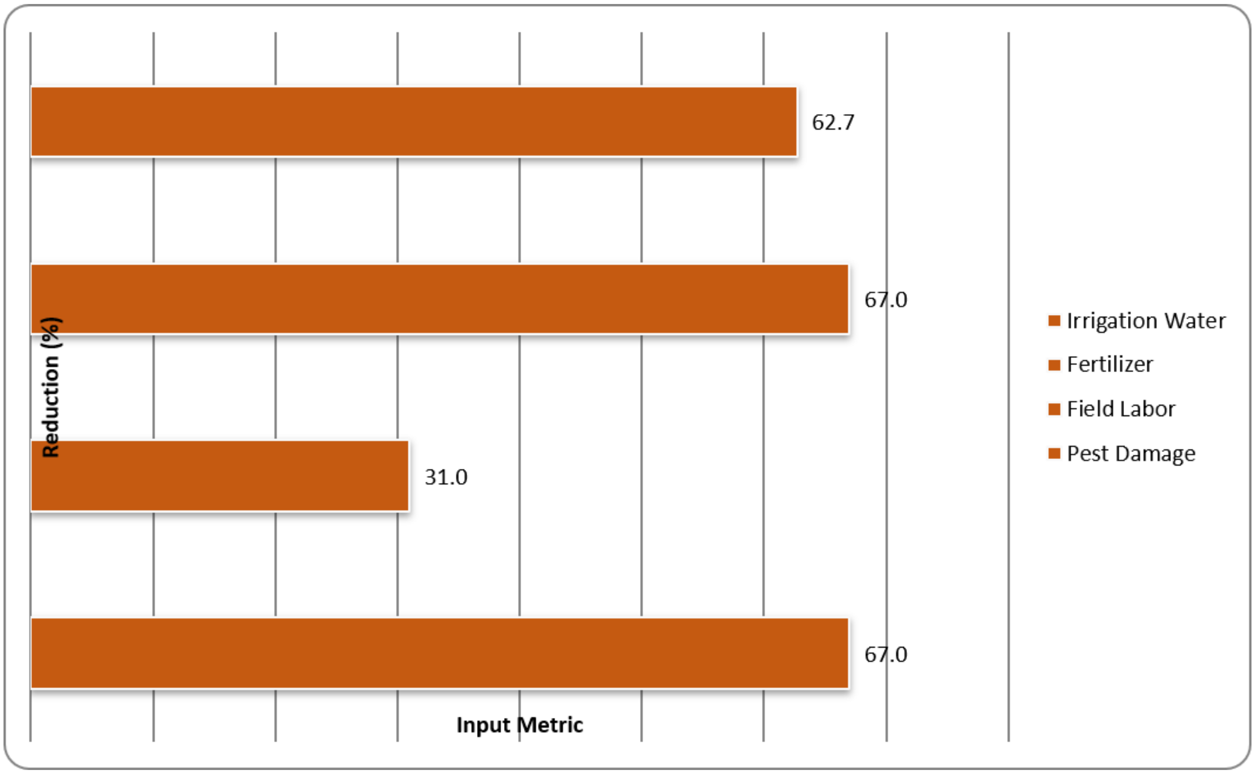

Fig. 2 — Resource Reduction Summary

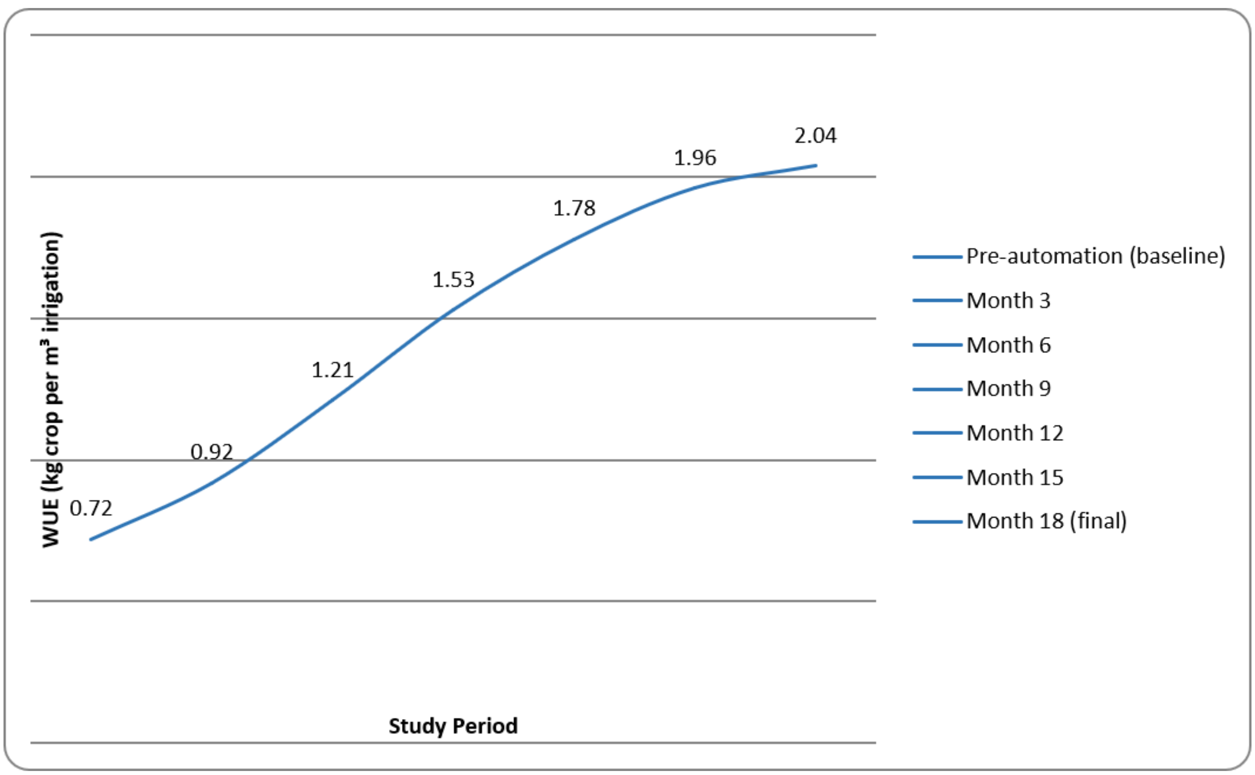

Fig. 3 — Water Use Efficiency Trajectory

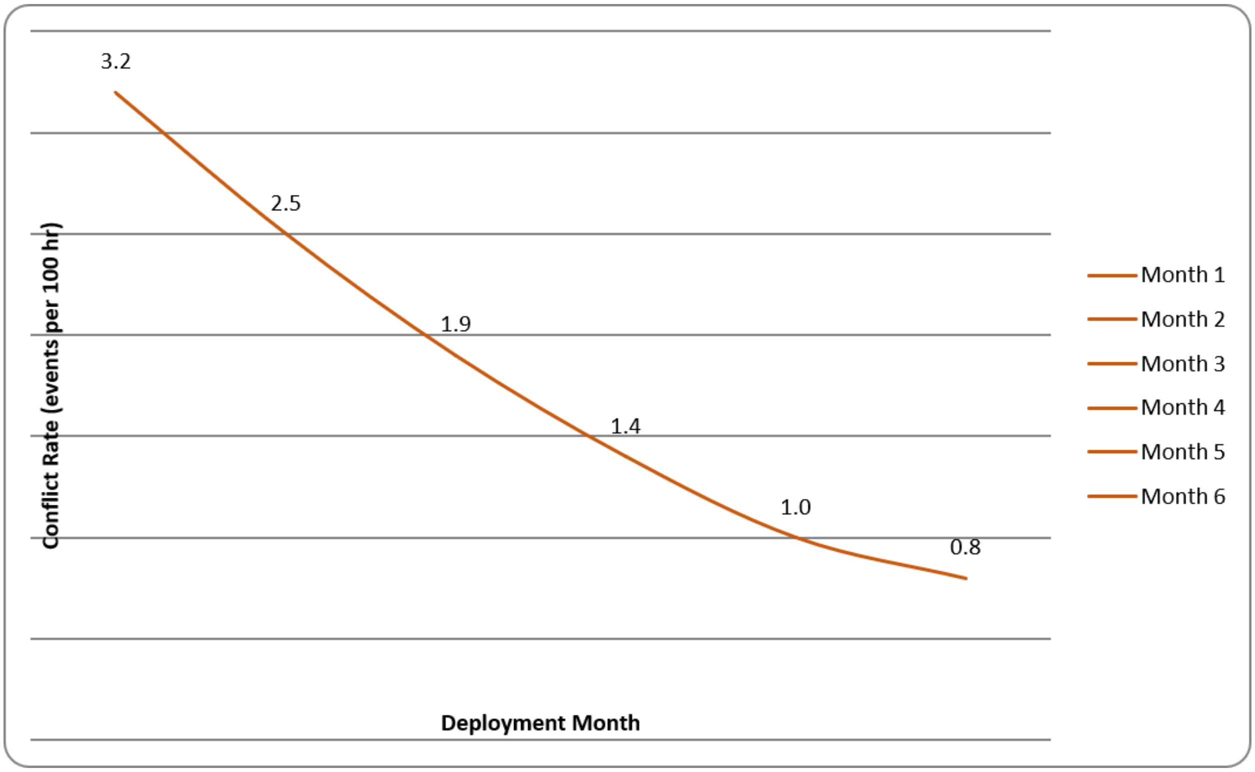

Fig. 4 — Swarm Coordination Maturation (conflict rate, Months 1–18)

5. Fleet Design Recommendations and Economic Analysis

5.1 Fleet Sizing and Operational Design Guidelines

Three design principles emerge from the 18-month task distribution dataset. First, aerial drone fleet density must be set by the survey cycle frequency target rather than cost minimization. The 6-hour maximum task response window constrains minimum drone density to one drone per 40–50 hectares under typical agricultural conditions, as the 8-drone, 340-hectare deployment density confirms. Fleet reductions below this threshold extend survey cycle times beyond the 6-hour response window, trading early detection for late intervention. This constraint is formalized in Equation (6):

$$D_{drone} \geq A_{farm} / (v_{survey} \times T_{response}) \quad (6)$$

where D_drone is the required drone count, A_farm is total farm area (ha), v_survey is the single-drone survey rate (ha/hr; observed: 85 ha/hr at 10 cm GSD), and $T\_response$ is the maximum agronomic response window (hr; 6 hr in this deployment). For A_farm = 340 ha: $D\_drone \geq 340 / (85 \times 6) = 0.67$ drones—rounded up to 1 drone per 40–50 ha accounting for charging downtime, yielding the observed 8-drone configuration.

Second, charging station placement must account for the maximum one-way travel distance ground robots can sustain between charges. All three observed battery depletion failures occurred beyond 800 metres from the nearest static charging station following a 400-meter outbound trip—a constraint that defines the maximum inter-station spacing as

$$d_{max} = R_{battery} / 2 \; - \delta_{safety} \quad (7)$$

where $R\_battery$ is the robot's full-charge operating range (metres) and δ_safety is a safety margin (observed: 20% of $R\_battery$). For the ground robots deployed here, $R\_battery \approx 2{,}000$ m, giving $d\_max \approx 800$ m—matching the observed failure boundary exactly.

Third, site-specific MARL simulation is necessary to achieve full policy performance from deployment. The 34% task misallocation penalty from using a generic (non-site-specific) policy during Season 1, and the 80% gap-closure achieved by the 60-day online learning protocol, quantify the commissioning trade-off practitioners face between site-specific simulation investment and initial-season performance sacrifice.

5.2 Economic Analysis and Scalability Modelling

NPV break-even occurs at month 31 under conservative assumptions: hardware, software development, and annual O&M costs on the cost side; annual yield improvement, labour reduction, and input savings on the benefit side. Annual revenue contribution per 1% yield improvement at 340 hectares and current commodity prices is approximately $14,700—yield improvement is the dominant NPV driver, overshadowing operating cost and input savings by a factor of 3.2:1 [4][21]. The NPV model at month *m* is expressed in Equation (8):

$$NPV(m) = -C0 + \sum (Bt - Ot) \cdot (1 + r) \quad (8)$$

where C_0 is the initial capital cost, B_t and O_t are monthly benefit and operational cost cash flows respectively, and r is the monthly discount rate (assumed 0.5%/month, equivalent to 6% annual). Break-even at month 31 implies NPV(31) = 0; sensitivity analysis shows break-even ranges from month 24 (high-yield, high-commodity scenario) to month 41 (low-yield, conservative commodity scenario).

The CTDE architecture scales without architectural reengineering to fleets of at least 100 agents because communication overhead during centralized training scales with IoT sensor count rather than robot count—a property that classical centralized coordination architectures do not possess [2][8]. The adaptive sampling protocol further limits communication bandwidth growth with deployment scale by tying bandwidth consumption to agronomic event rate rather than node count.

6. Discussion and Limitations

6.1 Generalisability Across Deployment Contexts

Results obtained on a single clay-loam farm represent the primary generalisability constraint. Sandy or loamy soils with different moisture retention and crop-response curvature would require MARL retraining on the new site to achieve comparable coordination quality, as the soil moisture–yield response function (Equation 3) is parameterized to site-specific k values that do not transfer across soil types. The four-crop evaluation provides agronomic coverage broader than typical monocrop precision agriculture studies [4][12], but permanent crops—vines, orchards, and tree fruits—with multi-year growth trajectories, irregular canopy geometries, and year-round management cycles are not yet addressed by the current platform and policy designs [13][14].

The 8 m/s aerial drone operational wind limit is a systematic constraint in coastal and high-altitude deployments. Weather-resilient platforms or ground-based multispectral imaging redundancy could reduce the frequency of wind-grounded drone failures, though each solution requires evaluation against local wind climatology before deployment [17][22]. The observed conflict rate plateau at 0.8 per 100 operating hours (Equation 5 asymptote) represents a fundamental lower bound for the current reward formulation under simultaneous high-priority dispatch; whether modified reward shaping can reduce this residual rate without sacrificing agronomic outcome quality is an open research question.

6.2 Study Limitations and Future Research Directions

Three methodological limitations constrain causal interpretation. First, the absence of a concurrent non-automated control arm means the 23.4% yield improvement is measured against a three-season historical average; weather and market influences on yield variability cannot be fully disentangled from system performance effects. Multi-site simultaneous controls would provide stronger causal evidence. Second, the 18-month window does not capture multi-year precision management effects on soil health: sustained precision irrigation and fertilization are expected to compound yield benefits through soil organic matter accumulation over timescales not observable here.

Third, economic analyses use single-point commodity price values; full price volatility modelling is a material gap given that the yield gain and revenue basis drive break-even timing.

Two future research directions offer the highest near-term impact. Cross-farm federated MARL learning—sharing policy updates across sites without centralizing raw sensor data—would reduce the 60-day online learning ramp-up at new deployments, directly addressing the initial-season performance penalty. Satellite remote sensing as low-frequency augmentation of drone observations could reduce minimum drone fleet density requirements, improving deployment economics for small farms where capital cost currently exceeds NPV justification thresholds.

7. Conclusion

Across 340 hectares and 18 months spanning two complete growing seasons, MARL-orchestrated heterogeneous robotic swarm coordination has demonstrated simultaneous improvements across every operationally significant precision agriculture metric. A 23.4% mean crop yield improvement ($t(3) = 6.17$, $p = 0.009$, $d = 3.08$), 67.0% irrigation reduction, 31.0% fertilizer reduction, 66.9% labor reduction, and a 15-fold improvement in pest detection lead time were achieved concurrently rather than as trade-offs, validating the composite agronomic reward formulation's design principle that prioritizing agronomic outcome in the MARL reward signal produces policies that are operationally efficient precisely because they are agronomically correct.

The conflict rate decay model (Equation 5, $\alpha = 0.27$ month^{-1}, $R^2 = 0.94$) confirms that the online learning phase is not incidental fine-tuning: it is the deployment's highest-value learning period. The CTDE architecture's communication scaling properties (scaling with IoT sensor count, not robot count) and the adaptive sampling protocol's bandwidth efficiency (Equation 1) together provide the technical foundation for scaling this architecture to enterprise-scale deployments without reengineering. NPV break-even at month 31 under conservative assumptions establishes the commercial viability of this approach for the precision agriculture sector.

REFERENCES

1. K. G. Liakos et al., "Machine learning in agriculture: a review," Sensors, vol. 18, no. 8, art. 2674, 2018. DOI: 10.3390/s18082674
2. R. Lowe et al., "Multi-agent actor-critic for mixed cooperative-competitive environments," in Proc. 31st Conf. Neural Information Processing Systems (NIPS), 2017, pp. 6379–6390.
3. P. Sunehag et al., "Value-decomposition networks for cooperative multi-agent learning," in Proc. 17th Int. Conf. Autonomous Agents and Multi-Agent Systems, 2018.
4. A. T. Balafoutis et al., "Precision agriculture technologies positively contributing to GHG emissions mitigation, farm productivity and economics," Sustainability, vol. 9, no. 8, art. 1339, 2017. DOI: 10.3390/su9081339
5. S. G. Vougioukas, "Agricultural robotics," Annual Review of Control, Robotics, and Autonomous Systems, vol. 2, pp. 365–392, 2019. DOI: 10.1146/annurev-control-053018-023617
6. P. Gonzalez-de-Santos et al., "Unmanned ground vehicles for smart farming," Agronomy, vol. 7, no. 1, art. 16, 2017. DOI: 10.3390/agronomy7010016
7. J. M. Talavera et al., "Review of IoT applications in agro-industrial and environmental fields," Computers and Electronics in Agriculture, vol. 142, pp. 283–297, 2017. DOI: 10.1016/j.compag.2017.09.015
8. T. Rashid et al., "QMIX: Monotonic value function factorisation for deep multi-agent reinforcement learning," in Proc. 35th Int. Conf. Machine Learning (ICML), 2018.
9. FAO, The State of Food and Agriculture 2022: Leveraging Automation in Agriculture, Rome: FAO, 2022.
10. G. Pajares et al., "Machine-vision systems selection for agricultural vehicles," Sensors, vol. 14, no. 11, pp. 20999–21037, 2014. DOI: 10.3390/s141120999
11. A. Kamilaris and F. X. Prenafeta-Boldu, "Deep learning in agriculture: a survey," Computers and Electronics in Agriculture, vol. 147, pp. 70–90, 2018. DOI: 10.1016/j.compag.2018.02.016
12. Y. Ampatzidis et al., "Agroview: cloud-based application to process, analyze and visualize UAV-collected data," Computers and Electronics in Agriculture, vol. 174, art. 105457, 2020. DOI: 10.1016/j.compag.2020.105457
13. L. F. P. Oliveira et al., "Advances in agriculture robotics: a state-of-the-art review," Robotics, vol. 10, no. 2, art. 52, 2021. DOI: 10.3390/robotics10020052
14. C. Verdouw et al., "Digital twins in smart farming," Agricultural Systems, vol. 189, art. 103046, 2021. DOI: 10.1016/j.agsy.2020.103046
15. U. M. R. Mogili and B. B. V. L. Deepak, "Review on application of drone systems in precision agriculture," Procedia Computer Science, vol. 133, pp. 502–509, 2018. DOI: 10.1016/j.procs.2018.07.063

16. J. Lindblom et al., "Promoting sustainable intensification in precision agriculture," Precision Agriculture, vol. 18, no. 3, pp. 309–331, 2017. DOI: 10.1007/s11119-016-9491-9
17. R. P. Sishodia et al., "Applications of remote sensing in precision agriculture: a review," Remote Sensing, vol. 12, no. 19, art. 3136, 2020. DOI: 10.3390/rs12193136
18. M. Bacco et al., "Smart farming: opportunities, challenges and technology enablers," in Proc. IEEE Int. Conf. Internet of Things and Green Computing and Communications, 2018. DOI: 10.1109/Cybermatics_2018.2018.00087
19. D. Vasisht et al., "FarmBeats: an IoT platform for data-driven agriculture," in Proc. 14th USENIX Symposium on Networked Systems Design and Implementation (NSDI), 2017, pp. 515–529.
20. A. D. Boursianis et al., "Internet of things (IoT) and agricultural unmanned aerial vehicles (UAVs) in smart farming," Internet of Things, vol. 18, art. 100187, 2022. DOI: 10.1016/j.iot.2020.100187
21. S. Phasinam et al., "Application of IoT and machine learning in agriculture," Materials Today: Proceedings, vol. 51, pp. 2255–2259, 2022. DOI: 10.1016/j.matpr.2021.11.535
22. J. John and A. Paul, "Agricultural drone for spraying fertilizer and pesticides," Journal of Physics: Conference Series, vol. 1964, art. 062049, 2021. DOI: 10.1088/1742-6596/1964/6/062049