Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X Volume 4, Issue 5, 2026

CROSS-LAYER RESILIENCE IN UAV-ENABLED DIGITAL AGRICULTURE: A SYSTEMATIC MAPPING REVIEW

Khakoo Mal[*1], Basit Raza[2]

[*1]*Department of Computer Science, Sukkur IBA University, Sukkur, Sindh, Pakistan.*
[2]*Institute of Computer Science, Shah Abdul Latif University, Khairpur Mirs, Sindh, Pakistan*

[*1]khakoo.mal@iba-suk.edu.pk, [2]basitraza.computerscience@gmail.com

DOI: https://doi.org/10.5281/zenodo.20354998

Keywords
UAV; drone; digital agricul-ture; precision agriculture; sys-tematic mapping; edge AI; IoT; 5G; 6G; resilient agriculture; cross-layer systems; deploy-ment evaluation

Article History
Received: 25 March 2026
Accepted: 04 May 2026
Published: 23 May 2026

Corresponding Author: *
Khakoo Mal

Abstract
Unmanned aerial vehicles (UAVs) have become a common part of digital agriculture, yet much of the literature still treats them as aerial cameras at-tached to crop-monitoring workflows. That view is useful, but it misses several conditions that decide whether a field system can be trusted in prac-tice. A farm-deployed UAV workflow has to keep producing useful decisions when illumination changes, wind and dust affect image quality, rural links become unstable, onboard devices have limited compute, and battery capacity constrains both sensing and communication. This review asks: how can UAV-enabled agricultural systems remain reliable when sensing quality, edge intelligence, communication, energy, and decision requirements interact in the field? We address this question through a systematic mapping review of UAV-enabled digital agriculture and closely related edge, communication, and deployment literature. The screening process identified 1,248 records, screened 897 records by title and abstract, assessed 238 full texts, and retained 45 sources that met the final criteria for direct cross-layer synthesis. Each full text was coded using a common scheme covering the primary system layer, computation location, communication dependency, energy treatment, deploy-ment evidence, reported metrics, and decision-support role. In this review, resilience means the ability of a UAV-agriculture workflow to maintain useful sensing, inference, communication, mission execution, and decision delivery under field and system disturbances. The synthesis shows that the literature is strong in sensing and task-level perception, but much less consistent in re-porting latency, energy, connectivity, uncertainty, and farmer-facing decision value together. The paper contributes a seven-layer deployment taxonomy, a reporting-quality appraisal, a failure-mode and mitigation landscape, a com-pact quantitative reporting summary, and a minimal benchmark specification for future cross-layer UAV-agriculture evaluation.

INTRODUCTION

Digital agriculture is no longer limited to occa-sional field observation. Farms increasingly rely on sensors, UAVs, Internet of Things (IoT) devices, machine learning models, and communication net-works to monitor crop stress, detect weeds and dis-eases, plan irrigation, guide spraying, and support management decisions. UAVs are attractive in this setting because they can capture high-resolution data over large areas without requiring permanent field infrastructure [1–3, 5]. Their value is espe-cially clear in fields where satellite imagery is too coarse, fixed cameras are impractical, or manual scouting is slow and inconsistent.

The UAV-agriculture literature has expanded across crop monitoring, plant-stress assessment, weed mapping, pest surveillance, canopy analysis, smart spraying, and yield estimation [4, 6, 7, 12]. Sensor choices have also widened from RGB cameras to multispectral, hyperspectral, thermal, and LiDAR systems [9–11]. At the same time, edge computing, cloud–edge collaboration, UAV networking, and 5G/6G smart-agriculture studies have started to frame UAVs as connected cyber-physical nodes rather than passive image collectors [15, 18–21]. What remains less settled is how these strands should be evaluated together when sensing quality, onboard computation, wireless delay, battery use, and farm decisions constrain one another.

Despite this progress, many papers still evalu-ate only one part of the deployment pipeline. A disease-detection model may report accuracy but not the device on which it runs. A UAV communication paper may report throughput or delay without connecting those numbers to crop decisions. A path-planning study may optimize coverage while assuming stable sensing and communication. A cloud-based analytics system may work well in a controlled farm but become slow or unusable in rural areas with intermittent connectivity. These examples point to the same problem: UAV-enabled agriculture is often studied as a set of separate technical tasks, while real deployment behaves as a coupled system.

This paper therefore reviews UAV-enabled digi-tal agriculture from a cross-layer deployment perspective. It treats the UAV as a mobile sensing, edge-intelligence, and communication node that operates inside a physical farm environment and produces decision-support outputs. The central research question is:

How can UAV-enabled agricultural sys-tems remain reliable when farms face weak connectivity, changing environ-mental conditions, limited onboard com-putation, battery constraints, and time-sensitive decision requirements?

The paper does not propose a new UAV algorithm, drone platform, or communication protocol. Its contribution is a systematic mapping and syn-thesis framework for judging deployment readi-ness across layers. Four contributions structure the review. First, the paper consolidates UAV-enabled digital agriculture into a seven-layer frame-work linking farm conditions, sensing, onboard AI, edge intelligence, communication, energy-aware mobility, and decision support. Second, it maps the existing literature through shared coding cri-teria and a quality checklist, showing which lay-ers are well reported and which are commonly assumed away. Third, it connects failure modes to mitigation strategies reported in the literature, in-cluding field-shift handling, edge/offline inference, energy-aware routing, communication-aware task allocation, and uncertainty-aware decision thresholds. Fourth, it proposes a concrete roadmap with testable challenge problems and a minimal cross-layer benchmark specification that can guide future datasets and evaluation protocols.

Operational Meaning of Resilience

In this review, resilience is used in an operational sense. A resilient UAV-enabled agricultural work-flow is one that continues to produce a useful farm decision when field conditions and system resources change. The term therefore covers more than model accuracy. It includes sensing under variable light and weather, inference under device limits, communication under weak rural links, mission completion under battery and payload constraints, and decision delivery in a form that a farmer or agronomist can act on. Table 1 summarizes the five resilience dimensions used through-out the synthesis.

REVIEW METHODOLOGY

This review was conducted as a systematic mapping study rather than a narrative or semi-systematic review. A systematic mapping design is suitable here because the topic crosses several communities: UAV remote sensing, crop analytics, edge AI, IoT agriculture, wireless networks, energy-aware mobility, and adoption studies. The purpose was not to estimate a single pooled effect size. The purpose was to classify what has been studied, what evidence is reported, and which deployment layers remain weakly connected.

Search Strategy

The search covered work published between 2015 and April 2026, with stronger emphasis on the 2019–2026 period where UAV agriculture, edge intelligence, and 5G/6G agriculture became more visible. The databases and digital libraries con-sidered were IEEE Xplore, Scopus, Web of Sci-ence, ScienceDirect, SpringerLink, ACM Digital Library, MDPI, Frontiers, and selected preprint sources. The search combined UAV and drone terms with agriculture, sensing, AI, edge com-puting, communication, energy, and deployment terms. The main search strings were:

- (“UAV” OR “drone”) AND (“precision agricul-ture” OR “smart agriculture” OR “digital agri-culture”);
- (“UAV” OR “drone”) AND (“crop monitoring” OR “plant stress” OR “disease detection” OR “weed detection”);
- (“UAV” OR “drone”) AND (“edge AI” OR “mo-bile edge computing” OR “onboard inference” OR “cloud edge”);
- (“UAV” OR “drone”) AND (“IoT” OR “5G” OR “6G” OR “connectivity”) AND “agriculture”;
- (“UAV” OR “drone”) AND (“battery” OR “energy-aware” OR “coverage path planning”) AND “agriculture”;
- (“digital agriculture” OR “resilient agriculture”) AND (“UAV” OR “drone”).

Table 1. Operational Resilience Dimensions Used in This Review

Dimension	Practical meaning	Suggested indicators
Sensing resilience	The UAV observes crops reliably despite lighting, weather, crop-stage, soil-background, and sensor changes.	Cross-weather F1, cross-season drop, calibration error, failure-case report.
Computational re-silience	The perception or analytics work-flow runs within device, memory, and timing limits.	Inference latency, FPS, model size, memory use, fallback mode.
Communication re-silience	Alerts and data remain usable when links are weak, intermittent, or de-layed.	Upload delay, packet loss, bandwidth, p95/p99 alert delay.
Mission resilience	The UAV completes useful cover-age within battery, route, payload, and hovering constraints.	Flight time, area covered, battery drain, route length, Wh/ha.
Decision resilience	The output supports a timely and trusted agricultural action.	False-alert cost, response time, dashboard usability, farmer feedback.

The initial search was deliberately broad. Rel-evant work appears under different labels, includ-ing UAV remote sensing, precision farming, smart farming, crop robotics, agricultural IoT, edge agri-culture, aerial networks, and cloud–edge–device collaboration.

Inclusion and Exclusion Criteria

A paper was eligible when it met at least one of the following conditions: it reviewed UAVs in agricul-ture; evaluated UAV-based agricultural sensing; proposed AI-enabled UAV perception for crop, weed, disease, stress, or yield tasks; discussed UAV, IoT, edge, or cloud integration for agricul-ture; addressed communication and connectivity relevant to smart agriculture; or studied energy, mobility, regulation, adoption, or deployment con-straints that affect UAV agriculture.

Records were excluded when they were purely military UAV studies, drone-hardware papers with-out agricultural relevance, image-classification papers without UAV or deployment context, non-English papers, short notes without sufficient technical detail, inaccessible full texts, or papers whose contribution could not be mapped to any cross-layer dimension. General edge, communication, calibration, and domain-shift references were retained only when they supported the evaluation lens; they were not treated as direct evidence of UAV-agriculture performance.

Screening and Coding Procedure

The selection process followed PRISMA 2020 reporting logic [38]. After duplicate, non-English, and clearly irrelevant records were removed, 897 records were screened by title and abstract. Reports that appeared relevant to UAV agriculture or to one of the cross-layer system dimensions were kept for full-text assessment. Of the 256 reports sought for retrieval, 238 full texts were assessed.

All 238 full texts were coded at the screening level using the criteria in Table 3. The screening-level code recorded whether the paper contained usable evidence for at least one of seven layers: physical farm environment, UAV sensing, onboard perception and AI, edge intelligence, communication and connectivity, energy and mobility, and decision support. For papers that passed eligibility, a fuller code was recorded for computation location, communication dependency, energy treatment, deployment evidence, decision role, and quantitative reporting.

After full-text eligibility assessment, 45 sources met the final criteria for direct cross-layer synthesis. The remaining full texts were excluded because they lacked sufficient cross-layer evidence, deployment relevance, agricultural context, or technical detail required for the mapping dimensions. The 45-source corpus was therefore not treated as a convenience sample; it was the focused direct-synthesis set used to compare evidence across the seven layers. When several papers covered the same narrow point, priority was given to recent reviews, higher-quality primary studies, and works that reported more than one deployment dimension. The coded source table in Appendix 10 lists all 45 directly synthesized sources and their cross-layer codes.

Quality and Relevance Assessment

The quality appraisal was designed for deployment reporting rather than clinical-style risk of bias. Each directly synthesized primary or review source was checked against seven items: crop and dataset context, sensor and flight details, environmental reporting, model and computation details, connectivity assumptions, energy or mission re-porting, and decision-support relevance. A study received one point for a clearly reported item, half a point for partial reporting, and zero when the item was absent. Scores of 5.0–7.0 were treated as high reporting quality, 3.0–4.5 as medium qual-ity, and below 3.0 as low quality. The checklist is summarized in Table 4. The purpose was not to re-ject low-scoring papers automatically. Instead, the appraisal shows where the literature is strong and where reporting habits make deployment claims hard to judge.

Across the mapped corpus, reporting quality was uneven. Sensing and remote-sensing papers often described sensors, imagery, and field context, but they rarely reported inference latency, memory use, or communication delay. Edge and communication papers were stronger on architectures and delay assumptions, yet many of them used generic evaluation settings that were not tied to crop decisions. Energy and path-planning papers usually treated flight endurance more carefully, but they often assumed perception outputs rather than mea-suring how model uncertainty affects mission de-cisions. Adoption-oriented studies brought useful evidence about cost, training, and user acceptance, although they typically did not measure sensing or network performance. This imbalance is one of the main reasons a cross-layer evaluation view is needed.

EXISTING REVIEW LANDSCAPE AND REMAINING CROSS-LAYER GAPS

The existing review landscape is broad but fragmented. UAV-agriculture reviews provide strong coverage of crop monitoring, spraying, sensing platforms, and application areas [2–4, 6, 32]. Remote-sensing reviews explain the value of hyperspectral, multispectral, thermal, and high-resolution imagery for crop and stress analysis [9, 10, 13, 34]. Smart-agriculture IoT reviews dis-cuss sensor networks, cloud platforms, and farm data integration [15–17]. Edge-computing reviews focus on latency, distributed execution, resource allocation, and cloud–edge–device architectures [18–20]. UAV communication surveys provide useful background on aerial links, beyond-5G networking, line-of-sight communication, and UAV-assisted wireless systems [22–24]. Adoption and digital-agriculture studies explain cost, skills, governance, data ownership, and farmer acceptance [27–30].

These streams are valuable, but they tend to organize evidence by discipline. The result is that deployment bottlenecks are distributed across papers rather than treated as one system problem. Table 5 summarizes the main review directions and the cross-layer gaps that remain.

Evidence Matrix and Reporting Patterns

Table 6 condenses the coded evidence from the 45 directly synthesized sources. The table does not treat all sources as equivalent primary experiments. Review papers were coded for the dimensions they synthesize, while primary and technical studies were coded for the metrics they report. Numerical values were extracted only when a source explicitly reported them; otherwise the entry is marked as not reported. This conservative choice avoids giving false precision to a literature in which many deployment variables are mentioned but not measured.

Because the included studies differ widely in crop type, sensor payload, validation split, UAV platform, hardware setting, and network assumption, this review reports metric availability and PRISMA Flow

Diagram

Study selection process for the systematic mapping of UAV-enabled digital agriculture

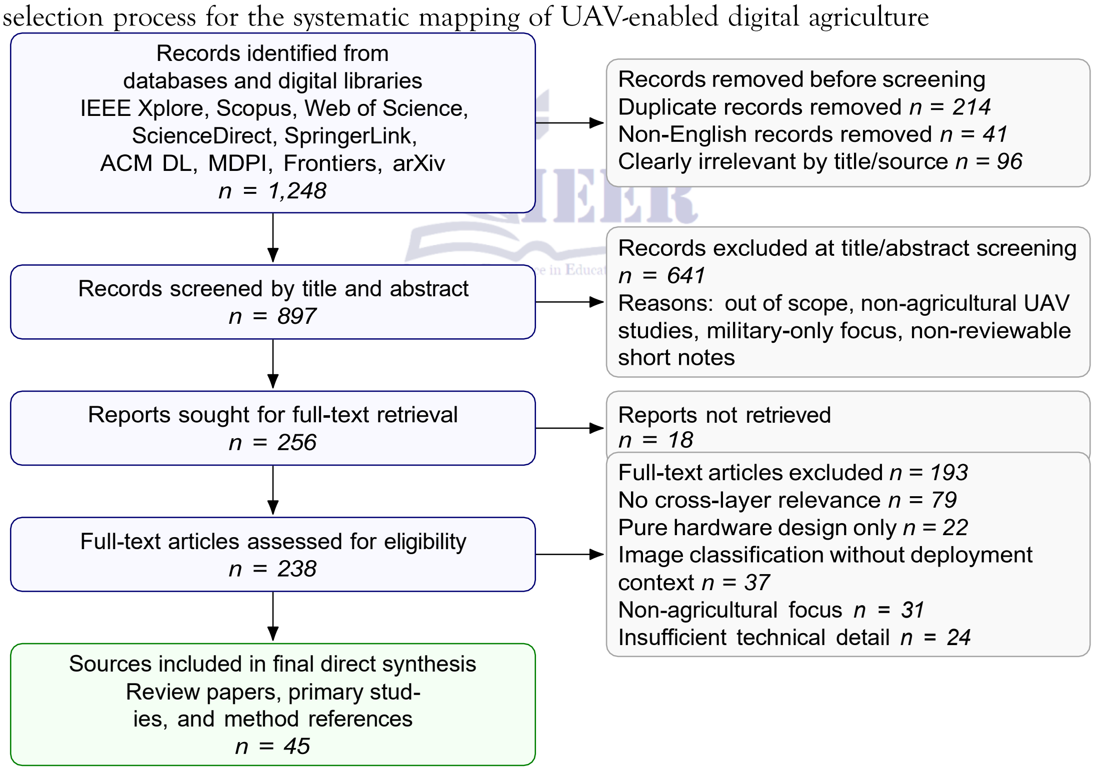

Fig. 1. PRISMA flow diagram for the systematic mapping review. The figure reports the screening path from 1,248 initial records to the 45 sources retained for final direct cross-layer synthesis.

Table 2. Reason-Wise Full-Text Exclusion After Eligibility Assessment representative reporting forms rather than pooled numerical effects.

Full-text exclusion reason	Count
No usable cross-layer evidence for the mapping dimensions	79
Pure hardware or platform design without agricultural deployment evidence	22
Image classification without UAV or deployment context	37
Non-agricultural focus after full-text inspection	31
Insufficient technical detail for reliable coding	24
Total full texts excluded	**193**

The evidence matrix points to a practical mismatch. UAV sensing studies commonly report task-level performance, but they do not always say whether the model can run on the UAV, whether the upload link can support the workflow, or how quickly the result reaches the farmer. Edge and communication papers often report delay and resource metrics, yet they may not connect those metrics to crop stress, spraying, irrigation, or disease-warning decisions. Adoption studies discuss cost and usability, but they rarely connect these issues to sensing quality or network behavior. A Q1-level research agenda for this area should therefore move from isolated metrics to cross-layer reporting.

Descriptive Mapping of Quality and Metrics

Applying the checklist in Table 4 gave the distribution shown in Table 8. The values should be read as a reporting-quality audit, not as a judgment of scientific worth. Some excellent algorithmic papers scored medium because they did not report field conditions or energy; some adoption papers scored medium because they did not include sensing or

Table 3. Study Classification Criteria Used for the Cross-Layer Mapping

Dimension	Codes used	How the code was applied
Primary layer	Environment; sensing; AI; edge; communication; energy; decision	Assigned from the dominant contribution. A disease-detection paper was coded as sensing/AI; an offloading paper was coded as edge/communication.
Computation location	Onboard; edge; cloud; hybrid; not specified	Based on where inference, analytics, model training, or decision logic was executed.
Communication dependency	Offline; intermittent; cloud-dependent; real-time; not specified	Based on whether the workflow required data upload, low-latency links, cloud processing, or local fallback.
Energy treatment	Not reported; battery mentioned; energy measured; energy optimized	Based on whether flight time, payload, battery drain, route cost, or compute-energy trade-off was discussed or evaluated.
Deployment evidence	Simulation; controlled field; real field; multi-site; adoption study; review	Based on the strongest setting reported. Single-farm experiments were coded as real field but not multi-site.
Decision role	Monitoring; alert; recommendation; autonomous action; dashboard; adoption	Based on how the technical output was connected to farm action or user interaction.
Quantitative reporting	Accuracy/perception; latency; bandwidth; energy/flight time; usability; none	Recorded only when the paper reported concrete metrics rather than general statements.

Table 4. Quality-Assessment Checklist Applied to the 45 Directly Synthesized Sources

Assessment item	What was checked	Reason for inclusion
Dataset and crop context	Crop type, field location, season, sample size, and data source	Makes it possible to judge whether findings are field-specific or reusable.
Sensor and flight details	Sensor type, altitude, resolution, platform, payload, and acquisition protocol	Connects perception results to UAV sensing constraints.
Environmental reporting	Lighting, weather, wind, crop stage, soil background, and field variability	Reveals whether field-shift risks can be interpreted.
Model and computation details	Model type, device, inference location, model size, latency, memory, or training setting	Shows whether the method can plausibly run on UAV or edge hardware.
Connectivity assumption	Offline operation, cloud upload, edge link, real-time link, or intermittent network	Indicates whether a system can work under weak rural connectivity.
Energy and mission reporting	Battery use, flight time, area coverage, path planning, payload, or compute-energy trade-off	Links sensing and AI choices to mission endurance.
Decision-support relevance	Alert, recommendation, dashboard, autonomous action, farmer feedback, or adoption evidence	Shows whether technical output leads to an agricultural decision.

Table 5. Thematic Comparison of Existing UAV-Agriculture and Adjacent Review Directions computation details.

Review direction	Main emphasis	Typical strength	Cross-layer limitation
UAV precision agriculture	Crop monitoring, mapping, spraying, vegetation indices, field applications	Strong application coverage and sensor overview	Often reports sensing benefits without latency, energy, or connectivity analysis.
Remote sensing and plant stress	Multispectral, hyperspectral, thermal, RGB, LiDAR, vegetation stress	Strong discussion of sensor capabilities and field imaging	Computation location and farmer-facing decision timing are often secondary.
AI and machine learning in agriculture	Detection, classification, segmentation, regression, yield and disease models	Clear model/task framing and accuracy reporting	Field shift, uncertainty, and hardware feasibility are not always evaluated.
IoT and smart agriculture	Field sensors, data platforms, cloud systems, automation	Good coverage of connected farm infrastructure	UAV-specific mobility, battery, and onboard inference constraints may be underdeveloped.
Edge/cloud agriculture	Edge nodes, cloud–edge collaboration, local analytics	Stronger treatment of latency and distributed execution	Agricultural decision value and UAV flight constraints are often indirect.
UAV communication and 5G/6G	UAV links, aerial networks, coverage, throughput, latency	Strong communication theory and network evaluation	Crop tasks and farm decisions are usually not the evaluation target.
Energy and path planning	Coverage, routing, battery, payload, mission scheduling	Stronger modeling of endurance and coverage	Perception uncertainty, edge AI, and communication quality are often assumed.
Adoption and sustainability	Cost, skills, policy, farmer acceptance, social effects	Strong socio-technical perspective	Technical sensing, AI, and network metrics are rarely measured together.

Two conclusions follow. First, the UAV-agriculture literature is not weak; it is uneven. It has strong sensor and task evidence, but weaker cross-layer measurement. Second, the most useful future papers will not necessarily be those with the largest models. They will be the studies that report enough context for another researcher to un-derstand when, where, and under what operating constraints the UAV system works.

CROSS-LAYER TAXONOMY FOR RESILIENT UAV-ENABLED DIGI-TAL AGRICULTURE

A deployment-ready UAV agriculture system can be described as seven connected layers. The layers are not a strict protocol stack. They are an evaluation lens that helps authors state which parts of the system are measured and which parts are assumed. Figure 2 shows the framework used in this review.

Physical Farm Environment Layer

The farm environment is the first source of uncertainty. Sun angle, cloud cover, wind, dust, crop stage, soil background, irrigation status, and terrain shape can all affect sensing and flight stability. Remote-sensing reviews have long emphasized the importance of acquisition conditions, but many AI studies still treat image collections as if they were stable benchmark datasets [10, 13, 34]. For deployment evaluation, environmental metadata should not be optional. It explains when model performance is likely to transfer and when it may fail.

UAV Sensing Layer

The sensing layer includes the UAV platform, camera, payload, flight altitude, ground sampling distance, overlap, and acquisition protocol. RGB cameras are inexpensive and useful for many scouting tasks, while multispectral and hyperspectral sensors improve vegetation and stress analysis at higher cost and payload weight [9–11]. Thermal cameras and LiDAR add further information but also increase system complexity. The deployment question is not simply which sensor is most accurate. It is which sensor provides enough information within the cost, payload, endurance, and maintenance limits of the target farm.

Table 6. Evidence Matrix Across the 45 Directly Synthesized Sources

Evidence group	Sources	Common metrics	Computation re-porting	Connectivity re-porting	Energy report-ing	Main synthesis finding
UAV-agriculture reviews	8	Application coverage; sensor and task categories	Mostly descriptive	Limited or indirect	Mentioned more than measured	Establish breadth, but rarely evaluate the full deployment chain.
Remote sensing and plant stress	7	Vegetation indices; mapping accuracy; spectral separability	Usually post-processing or unspecified	Rarely measured	Flight details partly reported	Field and sensor context are stronger than compute and network reporting.
AI and agricultural ML	5	Accuracy; F1; mAP; IoU; model comparison	Device and latency often incomplete	Usually assumed offline or cloud	Rarely tied to battery	Model performance is visible; deployment feasibility is harder to judge.
IoT and smart-agriculture integra-tion	6	Architecture; sensing coverage; data flow; platform functions	Cloud/edge split sometimes dis-cussed	Stronger than in vision papers	Usually a system concern, not a measured outcome	Integration is well motivated, but UAV-specific constraints are not always tested.
Edge and cloud–edge computing	4	Latency; offloading cost; resource allocation; device placement	Central focus	Often explicit	Sometimes dis-cussed	Edge execution addresses ru-ral latency, but crop-level de-cision value is rarely mea-sured.
UAV communication and 5G/6G	5	Throughput; delay; coverage; link relia-bility	Mostly network-level	Central focus	Sometimes linked to endurance	Useful communication evi-dence often sits outside agri-cultural decision workflows.
Energy, mobility, and coverage	2	Route length; coverage; battery/endurance; route cost	Usually assumed or simplified	Often assumed	Central focus	Mission planning is key, but it is rarely coupled with AI confidence or connectivity.
Adoption and digital-agriculture governance	4	Adoption intention; cost; skills; policy; usability	Usually not tech-nical	Contextual	Cost and main-tenance empha-sized	Adoption work explains why technically sound systems may fail in

						smallholder set-tings.
Evaluation-method references	4	Tail latency; calibra-tion; uncertainty; distribution shift	Evaluation lens	Supports p95/p99 and fallback re-porting	Not UAV-specific	These works justify measur-ing reliability beyond average accuracy.

Onboard Perception and AI Layer

The AI layer converts sensed data into crop-level information. Common tasks include disease detection, weed mapping, canopy segmentation, water-stress assessment, yield estimation, and pest surveillance [14, 35, 36]. The literature often reports accuracy, F1, mAP, or segmentation scores. These metrics are useful but incomplete. A model that performs well under one season or sensor may degrade under different illumination, crop stage, or farm management. Calibration and uncertainty reporting are therefore important when outputs are used for spraying, irrigation, or disease alerts [42, 43, 45].

Edge Intelligence Layer

Edge intelligence determines whether decisions can be made locally, near the field, or only after cloud upload. UAVs can run lightweight models onboard, offload data to a farm-edge node, or send data to a cloud platform [18–20]. Each choice changes latency, privacy, energy use, and dependence on connectivity. For rural farms, edge-first and offline-first designs are often more realistic than cloud-only pipelines. The key reporting variables are device type, model size, memory use, inference latency, frame rate, and fallback behavior when the network is unavailable.

Communication and Connectivity Layer

Connectivity affects whether UAV outputs reach the decision point in time. UAV-to-ground, UAV-to-edge, UAV-to-cloud, and UAV-to-UAV links can support real-time monitoring, swarm coordination, and cloud-assisted analytics [22–24]. In agriculture, however, connectivity may be intermittent or expensive. Mean latency alone is not enough; tail latency is often the value that determines whether an alert is late [44]. Future UAV-agriculture papers should report upload delay, packet loss, bandwidth, p95 or p99 alert delay, and the local fallback mode used when connectivity fails.

Table 7. Quantitative Reporting Summary Extracted from the Direct-Synthesis Corpus

Metric family	Studies reporting	Reported values or forms	Main comparability issue
Perception performance	17/45	Accuracy, F1, mAP, IoU, vegetation-index accuracy, or stress-separability measures; task-level values were typically high in controlled settings but varied by crop, sensor, and split.	Different crops, sensors, labels, and validation protocols prevent a defensible pooled effect size.
Latency and computation	8/45	Inference latency, offloading delay, model size, memory use, frame rate, or edge resource cost.	Hardware, batch size, network state, and model compression settings were often missing or not comparable.
Connectivity and network quality	9/45	Throughput, upload delay, packet loss, coverage, link reliability, or p95/p99-style delay recommendations.	Many values came from network-level or adjacent studies rather than agricultural decision workflows.
Energy, flight time, and mission cost	13/45	Flight time, battery/endurance, route length, covered area, payload notes, or energy-aware path cost.	Energy was usually separated from perception accuracy and communication load.
Uncertainty and shift behavior	5/45	Calibration, uncertainty, cross-domain evaluation, or distribution-shift protocols.	Most UAV-agriculture studies did not report uncertainty under held-out farms, seasons, or sensors.
Usability and adoption	7/45	Cost, farmer acceptance, training needs, dashboard requirements, or adoption barriers.	Usability evidence was rarely linked to telemetry, alert delay, or model confidence.

Table 8. Quality-Appraisal Summary by Dominant Evidence Group

Dominant evidence group	High	Medium	Low	Typical reporting pattern
UAV sensing and remote sensing	5	7	3	Field and sensor context are often good; compute and connectivity are weaker.
AI and perception	2	4	3	Accuracy is usually reported; latency, memory, calibration, and energy are inconsistent.
IoT, edge, and communication	3	6	2	Architecture and delay are clearer; crop decision relevance is often indirect.
Energy, mobility, and coverage	1	2	1	Battery and route constraints are visible; perception uncertainty is usually simplified.
Adoption and decision support	1	4	1	Cost and usability are discussed; technical telemetry is rarely measured.
Method/evaluation references	0	0	0	Used as evaluation support rather than appraised as UAV-agriculture studies.

Total appraised UAV/agriculture sources	12	23	10	Most papers cover one or two layers well; few report the whole deployment path.

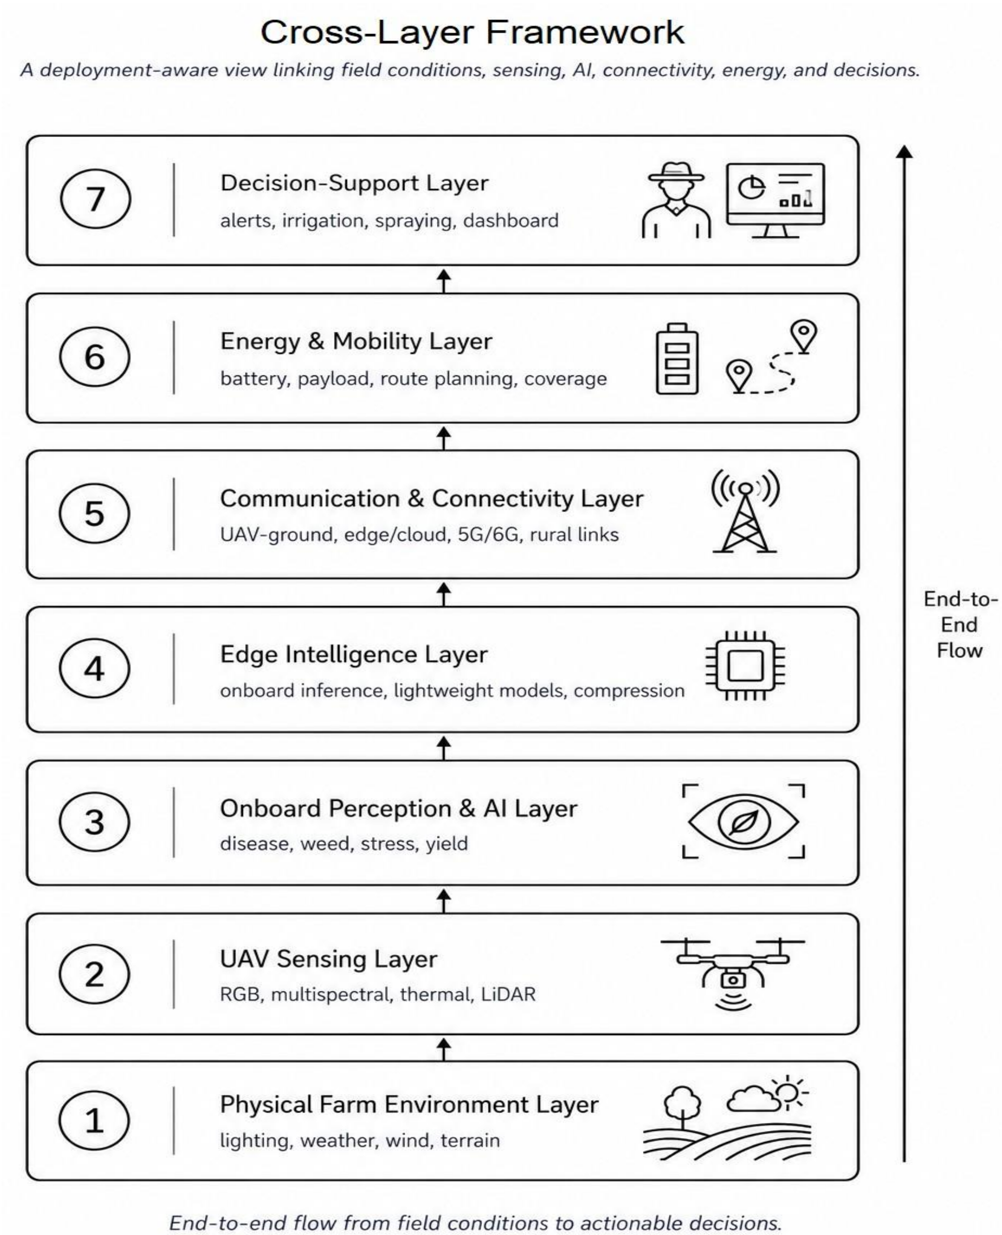

Fig. 2. Cross-layer framework for UAV-enabled digital agriculture, linking field conditions, sensing, onboard AI, edge intelligence, connectivity, energy-aware mobility, and decision support.

Energy and Mobility Layer

Battery capacity, payload weight, route planning, hovering time, sensing frequency, computation, and communication all affect UAV endurance. Coverage path planning studies show that route design is central to area coverage and mission time [26]. Yet energy is often reported separately from perception and communication. A heavier sensor may improve stress detection but shorten flight time. A larger model may improve accuracy but increase onboard compute load. A communication-heavy workflow may drain energy through repeated upload or hovering. Deployment-aware evaluation should therefore report flight time, covered area, battery drain, payload, route length, and compute-energy trade-off.

Decision-Support Layer

The final layer is the farm decision. A UAV output becomes useful when it supports an alert, recommendation, dashboard, spraying action, irrigation plan, scouting route, or management decision. Adoption studies remind us that technical performance is not enough; farmers also consider cost, training, trust, interface language, maintenance, regulation, and service availability [27–29]. A deployment-ready UAV system should report not only what the model predicts but also how the result is delivered, how quickly it arrives, and whether the user can act on it.

Table 9. Cross-Layer Taxonomy of UAV Agriculture Systems

Layer	Typical components	Deployment question	Suggested indicators
Physical farm environment	Lighting, weather, wind, dust, terrain, crop stage	Will the system remain reliable when field conditions change?	Illumination, weather, wind, crop stage, season, soil background.
UAV sensing	RGB, multispectral, hyperspectral, thermal, LiDAR, altitude, GSD	Which sensor is sufficient within cost and payload limits?	Sensor type, GSD, altitude, overlap, resolution, payload weight.
Onboard perception and AI	Detection, segmentation, classification, regression, uncertainty	Does the model generalize across farms, seasons, sensors, and weather?	F1, mAP, calibration, cross-domain drop, confidence quality.
Edge intelligence	Onboard inference, edge offloading, pruning, quantization, split computing	Can the model run within compute, memory, latency, and battery limits?	Model size, latency, FPS, RAM, device type, fallback mode.
Communication and connectivity	UAV-ground, UAV-edge, UAV-cloud, 5G/6G, rural links	Can decisions be delivered when connectivity is weak or delayed?	Bandwidth, loss, upload delay, p95/p99 alert delay, offline behavior.
Energy and mobility	Battery, payload, path planning, coverage, hovering	How should the UAV balance sensing quality, compute, communication, and coverage?	Flight time, battery drain, route length, area covered, Wh/ha.
Decision support	Alerts, dashboards, recommendations, spraying, irrigation	Does the system produce timely and trusted farm decisions?	Alert delay, decision accuracy, usability, cost, farmer feedback.

DEPLOYMENT CHALLENGES AND CROSS-LAYER FAILURE MODES

Deployment failures often occur between layers rather than inside one layer. A sensing method can fail because light changed. A model can fail because a crop stage was not represented in training data. An edge workflow can fail because the device cannot process images before the UAV leaves the area. A cloud workflow can fail because the link is slow. A decision system can fail because the alert is too late or not trusted. Table 10 summarizes the main failure modes.

Solutions Landscape

The reviewed literature does not provide one universal solution. Instead, it offers partial remedies that work under specific assumptions. Table 11 maps common failure modes to mitigation strategies and to the type of improvement usually reported. The limitations column is important: a mitigation that improves one layer may move the burden to another layer.

Table 11 also shows why mitigation cannot be judged from one layer alone. Edge inference can reduce dependence on a weak rural link, but it may increase onboard compute and battery demand. Energy-aware routing can extend coverage, yet it may change image overlap, viewing angle, and the time available for inference. Calibration is not only a machine-learning metric; it affects whether a farmer should trust an alert when the scene differs from the training data. These trade-offs are the practical reason for evaluating UAV agriculture as a coupled system.

Table 10. Deployment Challenges, Failure Modes, and Reporting Frequency

Failure mode	Affected layers	Practical impact	Recommended metric	Mapped sources
Illumination and weather shift	Environment, sensing, AI	Incorrect crop-stress or disease predictions under new light, weather, or crop-stage conditions.	Cross-weather F1, calibration error, failure-case report.	14/45
Weak rural connectivity	Communication, edge, decision	Delayed upload, late alert, or failure of cloud-dependent workflows.	p95/p99 alert delay, upload delay, fallback success.	13/45
Battery and payload limitation	Sensing, edge, energy	Reduced coverage, shorter mission time, or forced reduction in sensing quality.	Flight time, area covered, battery drain, Wh/ha.	12/45
Cloud dependence	Edge, communication, decision	System becomes unusable or slow when internet access is intermittent.	Cloud-only versus edge-only latency and availability.	9/45
Model generalization failure	Environment, sensing, AI	Poor transfer across farms, seasons, sensors, or crop varieties.	Cross-farm performance drop, ECE, Brier score.	11/45
Uncertain or overconfident alerts	AI, decision	False spraying, missed disease warning, or reduced farmer trust.	Confidence reliability, false-alert cost, review rate.	5/45
Cost and adoption barrier	Decision, deployment	Smallholders cannot afford, maintain, or trust the system.	Cost per hectare, training time, usability score.	7/45

Table 11. Solutions Landscape: Reported Mitigations, Improvement Patterns, and Remaining Limits

Failure mode	Mitigation reported in the literature	Observed improvement pattern	Remaining limitation	Representative sources
Field and illumina-tion shift	Multi-season acquisition, spectral sensing, aug-mentation, domain-shift evaluation	Better stress separability or smaller cross-condition performance loss when variation is represented	Few datasets cover many farms, seasons, sensors, and weather states to-gether	[9, 10, 13, 45]
High model cost on-board	Lightweight CNNs, prun-ing, quantization, edge of-floading, cloud–edge collab-oration	Lower inference delay or smaller memory footprint when compact models are used	Energy and accuracy trade-offs are not always reported together	[18–20, 35]
Weak connectivity	Edge-first inference, store-and-forward upload, UAV-edge links, local dashboards	Reduced dependence on cloud upload and shorter lo-cal decision delay	Requires careful synchro-nization and may reduce access to large cloud models	[15, 18, 19, 22]
Battery and cover-age limits	Energy-aware route plan-ning, payload-aware sens-ing, mission scheduling	Longer coverage or lower route cost under explicit mission constraints	Often assumes fixed per-ception quality and stable communication	[23, 24, 26]
Overconfident pre-dictions	Calibration, uncertainty esti-mation, abstention, human-in-the-loop review	Better confidence reliabil-ity and safer thresholding under dataset shift	Calibration can degrade under new farms or sen-sors without external val-idation	[42, 43, 45]
Low-resource adop-tion	Shared UAV services, sim-pler dashboards, regional-language interfaces, training support	Improved practical accessi-bility and user fit	Adoption evidence is of-ten survey-based rather than linked to technical telemetry	[5, 27–29]

A DEPLOYMENT-AWARE EVALUA-TION VIEW

A deployment-aware UAV-agriculture paper should make its assumptions visible. It should state where inference runs, which link is required, how much delay is tolerable, how the UAV energy budget is affected, and how the output supports a farm decision. The minimum evaluation set de-pends on the contribution, but a balanced report should include four groups of metrics.

First, perception metrics should still be reported: accuracy, F1, mAP, IoU, calibration, and cross-domain performance where relevant. Second, system metrics should report device type, inference latency, memory footprint, model size, bandwidth, upload delay, packet loss, and tail latency. Third, mission metrics should include flight time, battery use, payload, route length, area covered, and mission completion. Fourth, decision metrics should include alert delay, false alert cost, usability, farmer feedback, and action relevance.

Average values are not enough for time-sensitive systems. A farm alert that is fast on average but occasionally delayed by several minutes may not be suitable for spraying, irrigation, or disease-warning workflows. Tail-latency thinking from edge systems is therefore relevant to UAV agriculture, even when the original work is not agriculture-specific [44]. Similarly, uncertainty evaluation under dataset shift is relevant because a UAV model trained in one farm may be used under different light, crop stage, sensor, or season [43, 45].

RESEARCH GAPS AND FUTURE ROADMAP

Figure 3 summarizes the roadmap derived from the mapping study and groups future work into five deployment-oriented research directions.

Connectivity-Aware UAV Intelligence

Future UAV systems should decide when to process data onboard, when to use a nearby edge device, and when to upload to the cloud. This choice should depend on link quality, urgency, battery state, and the value of the decision. Two testable problems follow. First, for disease-alert missions under intermittent rural links, can an edge-first UAV workflow reduce p95 alert delay without reducing decision quality compared with a cloud-only workflow? Second, can link-aware task allocation keep alert delivery within a predefined time budget when bandwidth changes during flight?

Energy-Aware Edge AI

Energy-aware AI should report not only model accuracy but also the battery cost of sensing, inference, communication, and hovering. Two challenge questions are useful. Can model compression improve covered area per battery cycle while keeping crop-stress detection within an acceptable accuracy loss? Can mission planners jointly select route, sensor mode, and inference mode to minimize energy per actionable hectare rather than only route length?

Field-Shift Robust Perception

Field-shift robustness should become a normal evaluation requirement. A model should be tested across farms, seasons, weather, crop stages, and sensors where possible. Two hypotheses can guide future work. Models trained with environmental metadata and uncertainty-aware thresholds will produce fewer unsafe high-confidence errors under new farm conditions. External validation across farms and seasons will reveal larger performance drops than random train-test splits, especially for disease and stress tasks.

Cross-Layer Benchmarks

Future benchmark datasets should pair UAV imagery with telemetry, environmental context, network state, energy use, and decision labels. A benchmark should not only contain images. It should record the operating context that determines whether a UAV system can make a useful decision. The minimum specification proposed here is given in Box 13.

Two benchmark challenges follow from this specification. The first is a cross-layer alert-delivery challenge: given UAV imagery, environmental metadata, link state, and battery state, deliver a crop alert within a latency and energy budget. The second is a field-shift challenge: train on one set of farms, seasons, and sensors, then evaluate on held-out farms or seasons with calibration and uncertainty metrics.

Low-Resource and Farmer-Centered Deployment

Many UAV systems are developed for well-resourced experimental settings. Smallholders may need different deployment models: shared UAV services, lower-cost sensors, offline dashboards, regional-language interfaces, and simple

Table 12. Minimum Reporting Checklist for Deployment-Aware UAV-Agriculture Studies

Reporting group	Minimum information to include
Perception and uncertainty	Accuracy or F1/mAP/IoU, held-out farm or season split where possible, calibration or confidence reliability, and representative failure cases.
Compute and edge execution	Inference location, processor or edge device, model size, memory use, frame rate, latency distribution, and fallback behavior.
Communication and alert delivery	Link type, bandwidth, packet loss, upload delay, p95/p99 alert delay, cloud/edge availability, and offline mode.
Energy and mission operation	Battery capacity or level, payload, route length, flight time, covered area, hovering time, and estimated energy per hectare.
Decision and user context	Alert type, recommended action, time sensitivity, interface language, cost per hectare, training need, and farmer or agronomist feedback.

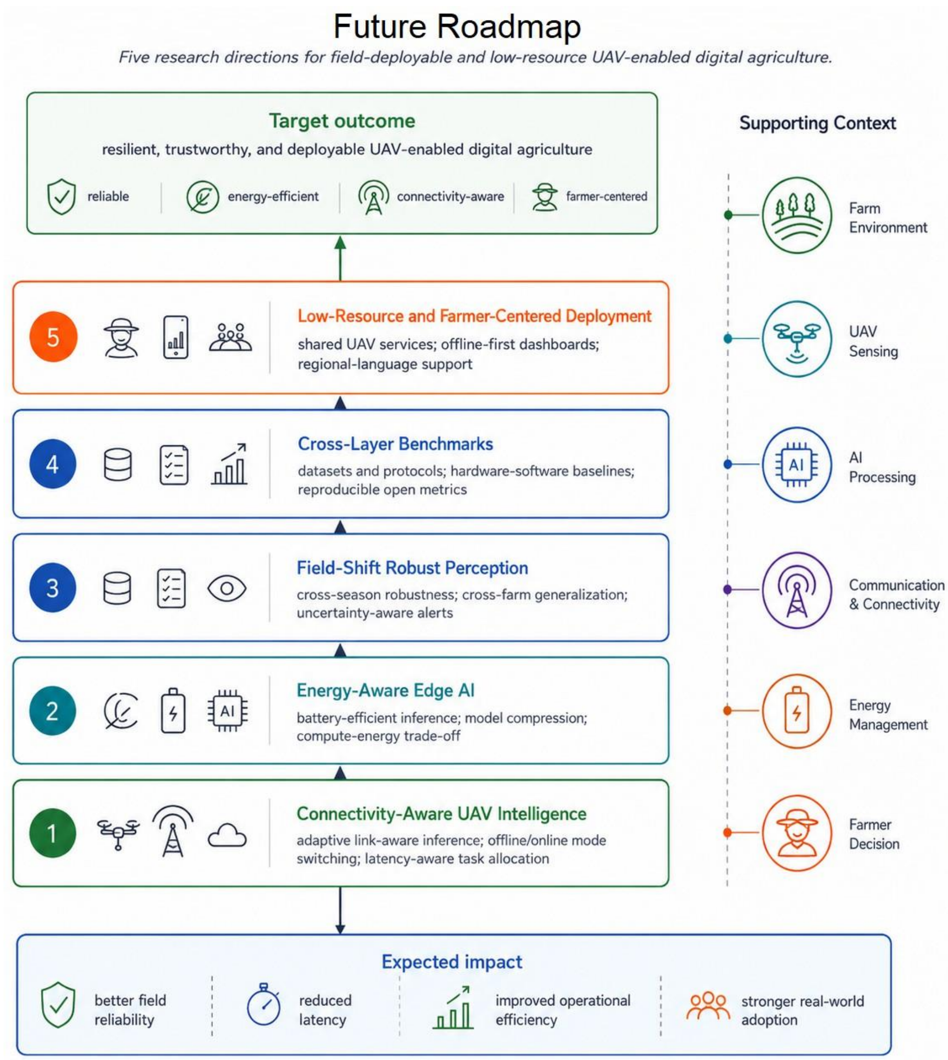

Fig. 3. Future roadmap for resilient UAV agriculture. The roadmap groups the next research directions into connectivity-aware intelligence, energy-aware edge AI, field-shift robust perception, cross-layer

benchmarks, and low-resource farmer-centered deployment. maintenance workflows. Two research questions are practical here. Can shared UAV services re-duce per-hectare cost while preserving decision timeliness? Can regional-language dashboards and uncertainty-aware alerts improve farmer trust com-pared with technical dashboards designed only for experts?

Table 13. Box 1. Minimal Cross-Layer Benchmark Specification for UAV-Enabled Digital Agriculture

Required component	Minimum fields to report
UAV image data	Raw or processed images; sensor type; altitude; ground sampling distance; camera settings; geolocation; timestamp; crop and field identifier.
Environmental logs	Weather, illumination condition, wind, temperature, humidity, crop stage, soil/background condition, and season.
Computation logs	Inference location, device type, model size, memory use, inference latency, frame rate, and local fallback mode.
Connectivity logs	Link type, bandwidth, packet loss, upload delay, p95/p99 alert delay, disconnection events, and cloud/edge availability.
Energy and mission logs	Battery level, flight time, route length, payload, hovering time, area covered, and estimated energy per hectare.
Ground-truth decisions	Crop condition labels, agronomist or expert annotation, disease/stress severity, recommended action, and time sensitivity of the decision.
Farmer usability indicators	Dashboard language, alert interpretability, training need, cost per hectare, response time, trust feedback, and maintenance burden.
Evaluation splits	Cross-farm, cross-season, cross-weather, cross-sensor, and leave-one-field-out splits, with random splits reported only as a baseline.

Table 14. Research Gaps, Testable Hypotheses, and Reporting Practices

Gap	Why it matters	Testable challenge problem	Suggested reporting practice
Vision-heavy evaluation	Accuracy alone does not show whether a UAV system is deployable.	Compare random-split performance with cross-farm, cross-season, and cross-weather performance.	Report accuracy with field metadata, calibration, and external validation.
Weak connectivity handling	Rural links may delay or block cloud workflows.	Test cloud-only, edge-only, and hybrid workflows under controlled link degradation.	Report bandwidth, loss, upload delay, p95/p99 alert delay, and fallback behavior.
Battery and payload constraints	Heavy sensors and models reduce flight time and coverage.	Jointly evaluate model choice, sensor mode, route length, and energy per hectare.	Report battery drain, flight duration, payload, route length, and area coverage.
Cloud dependence	Cloud-only systems can fail when connectivity is intermittent.	Measure whether edge-first inference maintains usable decision quality during disconnection.	Report online and offline modes separately.
Limited uncertainty awareness	Overconfident wrong alerts can cause unnecessary spraying or missed action.	Test calibration and abstention under held-out farms and seasons.	Report ECE, Brier score, confidence curves, and false-alert cost.
Low-resource adoption	Technical systems may not fit smallholder conditions.	Compare expert dashboard, simplified dashboard, and regional-language alert interface.	Report cost, training time, usability, maintenance needs, and farmer feedback.

DISCUSSION

The mapping study leads to a simple but important conclusion: UAV agriculture cannot be evaluated from one layer alone. A better sensor may be too expensive or too heavy for routine missions. A stronger model may be too slow for onboard in-ference. A cloud platform may become unusable under weak rural connectivity. A route that max-imizes coverage may still collect poor images if wind, illumination, or crop-stage variation is ig-nored. A high-accuracy model may cause harm if it is overconfident under field shift.

For researchers, the implication is that papers should state both the measured layer and the as-sumed layers. A crop-disease paper should say whether it tested new farms, seasons, or sensors. An edge AI paper should report device type, model size, memory use, and tail latency. A communi-cation paper should connect network performance to agricultural decision timing. A path-planning paper should state what happens to sensing quality and decision confidence when route cost is opti-mized. A deployment paper should describe cost,

Table 15. Illustrative Application of the Seven-Layer Lens to Representative Literature Streams usability, training, and farmer trust.

Representative stream	What the literature reports well	What the seven-layer lens makes visible
UAV sensing and re-mote sensing reviews [10, 13]	Sensor choices, field imaging conditions, vegetation indices, and crop-stress applications.	Computation location, alert delay, battery cost, and farmer-facing action are often outside the evaluation boundary.
UAV-IoT and smart-agriculture integration [5, 15]	Data flow, IoT integration, cloud or edge architecture, and farm-management use cases.	UAV payload, flight endurance, model uncertainty, and time-critical decision value need more explicit measurement.
UAV communication and edge-computing studies [20, 22, 24]	Throughput, offloading, cov-erage, latency, and resource-management models.	Crop-level decision timing and agricultural consequences of network delay are usually implied rather than evaluated.

For system designers, the cross-layer view en-courages joint design. Sensor selection, model ar-chitecture, offloading policy, wireless link, battery budget, and decision interface should be planned together. The strongest practical system is not nec-essarily the one with the highest isolated accuracy. It is the one that gives a useful decision under the constraints of the target farm.

The same view is relevant beyond agricultural imaging. UAV-enabled agriculture increasingly sits inside space–air–ground and edge-intelligent systems, where sensing, communication, local au-tonomy, and decision delivery are designed to-gether. A crop-alert workflow, for example, de-pends on the sensor payload, the onboard model, the link to an edge or cloud node, the remaining mission energy, and the decision interface used by the operator. This makes the review directly rele-vant to integrated sensing, edge intelligence, UAV networking, and autonomous decision-support re-search.

For journal reviewers and dataset builders, the results suggest a clear reporting standard. UAV-agriculture datasets should include operating con-text. Experiments should include at least one deployment-aware split. Evaluation should report failure cases rather than only best-case accuracy. These practices would make the field more compa-rable and would reduce the gap between impres-sive demonstrations and reliable farm use.

LIMITATIONS OF THIS REVIEW

This review has four limitations. First, the fi-nal 45-source direct-synthesis corpus was se-lected through explicit cross-layer eligibility crite-ria rather than by counting every UAV-agriculture publication that appeared in the initial search. The approach is suitable for a systematic mapping re-view, but it still depends on the transparency of the coding scheme and exclusion reasons. Sec-ond, reporting quality varied across the literature, which limits the extraction of comparable numerical values for latency, energy, communication, and usability. The quantitative summary therefore reports counts and reporting patterns rather than a pooled meta-analysis. Third, some edge, communication, calibration, and domain-shift references are adjacent methodological evidence rather than direct UAV-agriculture experiments. They are used to strengthen the evaluation lens, not to claim that their results directly validate agricultural UAV de-ployments. Fourth, the proposed benchmark spec-ification is a research recommendation. It should be refined through community datasets, field cam-paigns, and consultation with agronomists, UAV operators, and farmers.

CONCLUSION

This review asked how UAV-enabled agricultural systems can remain reliable when field conditions, sensing, AI, connectivity, energy, and decision requirements interact. The systematic mapping shows that the literature is rich in UAV sensing and crop-perception studies, but less consistent in reporting the system factors that determine deployment readiness. Latency, energy, connectivity, uncertainty, and farmer-facing decision value are often mentioned but not measured together.

The paper therefore argues for a cross-layer view of UAV-enabled digital agriculture. UAVs should be evaluated not only as aerial cameras, but as mobile sensing, edge-intelligence, and communication nodes operating inside uncertain farm environ-ments. The proposed taxonomy, quality appraisal, failure-mode landscape, roadmap, and benchmark specification provide a practical way to organize future work. A stronger next generation of UAV-agriculture studies will report not only whether a model works, but where it works, how fast it runs, how much energy it costs, how it behaves un-der weak connectivity, and whether the resulting decision is useful to the farmer.

REFERENCES

C. Zhang and J. M. Kovacs, "The application of small unmanned aerial systems for precision agriculture: A review," *Precision Agriculture*, vol. 13, no. 6, pp. 693–712, Dec. 2012, doi: 10.1007/s11119-012-9274-5.

D. C. Tsouros, S. Bibi, and P. G. Sarigiannidis, "A review on UAV-based applications for precision agriculture," *Information*, vol. 10, no. 11, Art. no. 349, Nov. 2019, doi: 10.3390/info10110349.

R. Guebsi, R. Mami, and A. Ben Ali, "Drones in precision agriculture: A comprehensive review," *Drones*, vol. 8, no. 5, Art. no. 198, May 2024, doi: 10.3390/drones8050198.

J. Agrawal and M. Y. Arafat, "Transforming farming: A review of AI-powered UAV technologies in precision agriculture," *Drones*, vol. 8, no. 11, Art. no. 664, Nov. 2024, doi: 10.3390/drones8110664.

S. Khan, T. Mazhar, T. Shahzad, M. A. Khan, S. Guizani, and H. Hamam, "Future of sustainable farming: Exploring opportunities and overcoming barriers in drone-IoT integration," *Discover Sustainability*, vol. 5, Art. no. 470, 2024, doi: 10.1007/s43621-024-00736-y.

A. Hafeez, M. A. Husain, S. P. Singh, A. Chauhan, M. T. Khan, N. Kumar, A. Chauhan, and A. Soni, "Implementation of drone technology for farm monitoring and pesticide spraying: A review," *Information Processing in Agriculture*, vol. 10, no. 2, pp. 192–203, Jun. 2023, doi: 10.1016/j.inpa.2022.02.002.

D. Olson and J. Anderson, "Review on unmanned aerial vehicles, remote sensors, imagery processing, and their applications in agriculture," *Agronomy Journal*, vol. 113, no. 2, pp. 971–992, 2021, doi: 10.1002/agj2.20595.

G. R. Radoglou-Grammatikis, P. Sarigiannidis, T. Lagkas, and I. Moscholios, "A compilation of UAV applications for precision agriculture," *Computer Networks*, vol. 172, Art. no. 107148, May 2020, doi: 10.1016/j.comnet.2020.107148.

T. Adão, J. Hruška, L. Pádua, J. Bessa, E. Peres, R. Morais, and J. J. Sousa, "Hyperspectral imaging: A review on UAV-based sensors, data processing and applications for agriculture and forestry," *Remote Sensing*, vol. 9, no. 11, Art. no. 1110, Oct. 2017, doi: 10.3390/rs9111110.

W. H. Maes and K. Steppe, "Perspectives for remote sensing with unmanned aerial vehicles in precision agri-culture," *Trends in Plant Science*, vol. 24, no. 2, pp. 152–164, Feb. 2019, doi: 10.1016/j.tplants.2018.11.007.

A. Matese and S. F. Di Gennaro, "Technology in precision viticulture: A state of the art review," *International Journal of Wine Research*, vol. 7, pp. 69–81, 2015, doi: 10.2147/IJWR.S69405.

Y. Xing, X. Liu, and X. Wang, "Integrating UAVs, satellite remote sensing, and machine learning in precision agriculture: Pathways to sustainable food production, resource efficiency, and scalable innovation," *Frontiers in Agronomy*, vol. 7, Art. no. 1670380, Jan. 2026, doi: 10.3389/fagro.2025.1670380.

J. G. A. Barbedo, "A review on the use of unmanned aerial vehicles and imaging sensors for monitoring and assessing plant stresses," *Drones*, vol. 3, no. 2, Art. no. 40, May 2019, doi: 10.3390/drones3020040.

M. F. Aslan, A. Durdu, K. Sabanci, and E. Ropelewska, "A comprehensive survey of the recent studies with UAV for precision agriculture in open fields and greenhouses," *Applied Sciences*, vol. 12, no. 3, Art. no. 1047, Jan. 2022, doi: 10.3390/app12031047.

A. D. Boursianis, M. S. Papadopoulou, P. Diamantoulakis, A. Liopa-Tsakalidi, P. Barouchas, G. Salahas, G. Karagiannidis, S. Wan, and S. K. Goudos, "Internet of Things (IoT) and agricultural unmanned aerial vehicles (UAVs) in smart farming: A comprehensive review," *Internet of Things*, vol. 18, Art. no. 100187, May 2022, doi: 10.1016/j.iot.2020.100187.

N. N. Thilakarathne, S. Bakar, P. E. Abas, and H. Yassin, "Internet of things enabled smart agriculture: Current status, latest advancements, challenges and future directions," *Heliyon*, vol. 11, no. 4, Art. no. e42136, 2025, doi: 10.1016/j.heliyon.2025.e42136.

N. R. Gatkal, S. M. Nalawade, R. K. Sahni, G. B. Bhanage, A. A. Walunj, P. B. Kadam, and M. Ali, "Review of IoT and electronics enabled smart agriculture," *International Journal of Agricultural and Biological Engineering*, vol. 17, no. 5, pp. 1–14, 2024, doi: 10.25165/j.ijabe.20241705.8496.

R. Gong, H. Zhang, G. Li, and J. He, "Edge computing-enabled smart agriculture: Technical architectures, practical evolution, and bottleneck breakthroughs," *Sensors*, vol. 25, no. 17, Art. no. 5302, Aug. 2025, doi: 10.3390/s25175302.

P. Yu, F. Teng, W. Zhu, C. Shen, Z. Chen, and J. Song, "Cloud–edge–device collaborative computing in smart agriculture: Architectures, applications, and future perspectives," *Frontiers in Plant Science*, vol. 16, Art. no. 1668545, Oct. 2025, doi: 10.3389/fpls.2025.1668545.

X. Xia, S. M. M. Fattah, and M. A. Babar, "A survey on UAV-enabled edge computing: Resource management perspective," *ACM Computing Surveys*, vol. 56, no. 3, Art. no. 78, pp. 1–36, Oct. 2023, doi: 10.1145/3626566.

F. Zhang, Y. Zhang, W. Lu, Y. Gao, G. Gong, and J. Cao, "6G-enabled smart agriculture: A review and prospect," *Electronics*, vol. 11, no. 18, Art. no. 2845, Sep. 2022, doi: 10.3390/electronics11182845.

B. Li, Z. Fei, and Y. Zhang, "UAV communications for 5G and beyond: Recent advances and future trends," *IEEE Internet of Things Journal*, vol. 6, no. 2, pp. 2241–2263, Apr. 2019, doi: 10.1109/JIOT.2018.2887086.

M. Mozaffari, W. Saad, M. Bennis, Y.-H. Nam, and M. Debbah, "A tutorial on UAVs for wireless networks: Applications, challenges, and open problems," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 3, pp. 2334–2360, 3rd Quart. 2019, doi: 10.1109/COMST.2019.2902862.

Y. Zeng, J. Lyu, and R. Zhang, "Accessing from the sky: A tutorial on UAV communications for 5G and beyond," *Proceedings of the IEEE*, vol. 107, no. 12, pp. 2327–2375, Dec. 2019, doi: 10.1109/JPROC.2019.2952892.

T. Lagkas, V. Argyriou, and P. Sarigiannidis, "UAV IoT framework views and challenges: Towards protecting drones as things," *Sensors*, vol. 18, no. 11, Art. no. 4015, Nov. 2018, doi: 10.3390/s18114015.

T. M. Cabreira, L. B. Brisolara, and P. R. Ferreira Jr., "Survey on coverage path planning with unmanned aerial vehicles," *Drones*, vol. 3, no. 1, Art. no. 4, Jan. 2019, doi: 10.3390/drones3010004.

N. Nafar, M. Fatemi, and K. Rezaei-Moghaddam, "Smart drone in sustainable agriculture: Evaluating four predictive models for technology adoption in Iran," *Smart Agricultural Technology*, vol. 12, Art. no. 101363, Dec. 2025, doi: 10.1016/j.atech.2025.101363.

M. Saiz-Rubio and F. Rovira-Más, "From smart farming towards agriculture 5.0: A review on crop data management," *Agronomy*, vol. 10, no. 2, Art. no. 207, Feb. 2020, doi: 10.3390/agronomy1000207.

L. Klerkx, E. Jakku, and P. Labarthe, "A review of social science on digital agriculture, smart farming and agriculture 4.0: New contributions and a future research agenda," *NJAS: Wageningen Journal of Life Sciences*, vol. 90–91, Art. no. 100315, Dec. 2019, doi: 10.1016/j.njas.2019.100315.

S. Wolfert, L. Ge, C. Verdouw, and M.-J. Bogaardt, "Big data in smart farming: A review," *Agricultural Systems*, vol. 153, pp. 69–80, May 2017, doi: 10.1016/j.agsy.2017.01.023.

T. Miller, M. Khan, and A. Ahmed, "The IoT and AI in agriculture: The time is now—A systematic review of smart sensing for arable crops and grasslands," *Sensors*, vol. 25, no. 12, Art. no. 3583, Jun. 2025, doi: 10.3390/s25123583.

B. K. Komatineni, R. K. Prasad, and S. Kumar, "A comprehensive review of the functionality and applications of UAVs in precision and smart agriculture," *Discover Applied Sciences*, vol. 6, Art. no. 184, 2024, doi: 10.1186/s43067-024-00184-8.

M. V. Ferro, S. Catania, and G. Laudani, "Technologies and innovative methods for precision viticulture: A comprehensive review," *Horticulturae*, vol. 9, no. 3, Art. no. 399, Mar. 2023, doi: 10.3390/horticulturae9030399.

S. Manfreda *et al.*, "On the use of unmanned aerial systems for environmental monitoring," *Remote Sensing*, vol. 10, no. 4, Art. no. 641, Apr. 2018, doi: 10.3390/rs10040641.

A. Kamilaris and F. X. Prenafeta-Boldú, "Deep learning in agriculture: A survey," *Computers and Electronics in Agriculture*, vol. 147, pp. 70–90, Apr. 2018, doi: 10.1016/j.compag.2018.02.016.

K. G. Liakos, P. Busato, D. Moshou, S. Pearson, and D. Bochtis, "Machine learning in agriculture: A review," *Sensors*, vol. 18, no. 8, Art. no. 2674, Aug. 2018, doi: 10.3390/s18082674.

G. Idoje, T. Dagiuklas, and M. Iqbal, "Survey for smart farming technologies: Challenges and issues," *Comput-ers and Electrical Engineering*, vol. 92, Art. no. 107104, Jun. 2021, doi: 10.1016/j.compeleceng.2021.107104.

M. J. Page, J. E. McKenzie, P. M. Bossuyt, I. Boutron, T. C. Hoffmann, C. D. Mulrow, *et al.*, "The PRISMA 2020 statement: An updated guideline for reporting systematic reviews," *BMJ*, vol. 372, Art. no. n71, Mar. 2021, doi: 10.1136/bmj.n71.

P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, *et al.*, "Advances and open problems in federated learning," *Foundations and Trends in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021, doi: 10.1561/2200000083.

T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," in *Proceedings of Machine Learning and Systems*, vol. 2, pp. 429–450, 2020.

H. Zhu, J. Xu, S. Liu, and Y. Jin, "Federated learning on non-IID data: A survey," *Neurocomputing*, vol. 465, pp. 371–390, Nov. 2021, doi: 10.1016/j.neucom.2021.07.098.

C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *Proceedings of the 34th International Conference on Machine Learning*, vol. 70, pp. 1321–1330, 2017.

Y. Ovadia, E. Fertig, J. Ren, Z. Nado, D. Sculley, S. Nowozin, *et al.*, "Can you trust your model's uncertainty? Evaluating predictive uncertainty under dataset shift," in *Advances in Neural Information Processing Systems*, vol. 32, 2019.

J. Dean and L. A. Barroso, "The tail at scale," *Communications of the ACM*, vol. 56, no. 2, pp. 74–80, Feb. 2013, doi: 10.1145/2408776.2408794.

P. W. Koh, S. Sagawa, H. Marklund, S. M. Xie, M. Zhang, A. Balsubramani, *et al.*, "WILDS: A benchmark of in-the-wild distribution shifts," in *Proc. 38th Int. Conf. Machine Learning*, vol. 139, pp. 5637–5664, 2021.