Secure And Distributed Edge Intelligence For Consumer-Grade Smart Agriculture

Jiangtao Deng[1,2], Shuya Zhao[3], Jijing Cai[3], Yuchao Xia[3], Meilei Lv[2,*], Kai Fang[3,*], Hailin Feng[3], and Thippa Reddy Gadekallu[3,4]

[1] *College of Automation, Hangzhou Dianzi University, Hangzhou, China*
[2] *College of Electrical and Information Engineering, Quzhou University, Quzhou, China*
[3] *College of Mathematics and Computer Science, Zhejiang A&F University, Hangzhou, China*
[4] *Division of Research and Development, Lovely Professional University, Phagwara, India*
[*] *Corresponding author. E-mail: 37014@qzc.edu.cn, Kaifang@zafu.edu.cn*

Received: December 23, 2025; Accepted: January 26, 2026

Consumer-grade Agricultural AIoT (Agri-AIoT) systems increasingly rely on cloud-based intelligence, which introduces latency, privacy, and connectivity limitations. These limitations are particularly severe for heterogeneous consumer devices operating under strict cost, energy, and computational constraints. This review advocates a shift from cloud-centric architectures toward distributed, edge-centric intelligence. We examine lightweight model compression techniques, including pruning, quantization, and knowledge distillation, that enable real-time and on-device decision-making. We further analyze security threats such as data poisoning and adversarial attacks that arise in decentralized agricultural systems. Privacy-preserving learning mechanisms, including Federated Learning, are discussed as key enablers of collaborative intelligence without raw data sharing. By integrating lightweight Artificial Intelligence techniques with AI-native networking principles, this paper provides a unified perspective on distributed intelligence in consumer Agri-AIoT ecosystems. We conclude that the convergence of these approaches is essential for building sustainable, secure, and self-adaptive consumer-grade agricultural electronics.

Keywords: AI-Native Networking; Resource-Constrained Devices; Agri-AIoT; Lightweight AI

http://dx.doi.org/10.6180/jase.202608_31.033

1. Introduction

Smart agriculture has become a strategic pathway to improve productivity, sustainability, and resilience under growing food-security and resource constraints. National- and sector-level roadmaps increasingly highlight the need to integrate sensing, computing, and intelligent equipment to replace manual labor with machines and human cognition with computation, and to promote staged, differentiated adoption toward 2035 [1]. Technically, the past decade of agricultural IoT (Agri-AIoT) research has established a rich foundation of architectures, sensors, networking protocols, and cloud-based analytics pipelines [2, 3]. However, much of this literature implicitly assumes industrial-grade infrastructure, characterized by enterprise budgets, professional operators, reliable power supply, and stable wired or managed wireless connectivity. These conditions rarely hold in consumer-grade settings. For consumer users such as home gardeners and small-scale growers, cost sensitivity is a primary driver that fundamentally shapes system design and adoption. Unlike industrial deployments supported by enterprise investment, consumer Agri-AIoT devices must operate under strict bill-of-materials (BOM) constraints, favoring low-cost sensors, microcontrollers, and commodity wireless connectivity. These cost

limitations not only restrict available computing and energy resources, but also discourage persistent cloud dependence, thereby directly motivating the transition toward lightweight, edge-centric intelligence. In recent years, a clear shift has emerged: consumer electronics Agricultural AIoT (Agri-AIoT) is moving intelligent agriculture from industrial farms into urban balconies, home gardens, and smallholder operations [4, 5]. The market is already populated with intuitive products: applications such as PlantNet and iNaturalist provide AI-assisted identification and diagnosis for pests and plants [6]; consumer irrigation controllers (e.g., Rachio) automate watering schedules; and consumer drones are increasingly used for low-cost crop health mapping and inspection. This consumer transition changes the problem definition from centralized "farm management systems" to a heterogeneous edge ecosystem, where smartphones, microcontrollers, low-cost sensors, small single-board computers (SBCs), and drones must collectively support real-time perception and actuation under strict cost and energy constraints.

A comprehensive and up-to-date review by Dembani et al. [7] systematically surveys FL applications in agriculture, covering disease detection, yield prediction, and resource management, while explicitly identifying open challenges related to data heterogeneity, communication constraints, stakeholder trust, and ethical considerations. Complementing this survey perspective, Sharma et al. [8] investigate a blockchain-enabled FL framework for securing IoT communications in precision agriculture, demonstrating how distributed ledger technologies can enhance privacy, trust, and attack resilience without sharing raw data. In parallel, Ding et al. [9] propose a double-layer blockchain-based FL architecture tailored for Agricultural IoT, addressing geo-distribution, heterogeneous data, and device-level risks through hierarchical aggregation and adaptive noise control. Together, these recent works underscore both the growing maturity of FL-based agricultural intelligence and the persistent technical and governance challenges that motivate the need for resource-aware, secure, and consumer-friendly Agri-AIoT frameworks. In contrast, this paper is the first to strictly focus on the *consumer-grade* Agri-AIoT ecosystem. We uniquely analyze the security-efficiency trade-offs inherent to low-cost hardware, bridging the gap between lightweight AI techniques and the specific threat landscape of unmanaged, non-expert environments. Despite the rapid adoption of consumer Agri-AIoT, user experience remains fragmented and heavily cloud-dependent [10]. Consider a home gardener, Maria. She uses a smartphone app to diagnose a blighted tomato leaf, but latency delays the cloud response. Meanwhile, her irrigation controller waters heavily based on cloud weather data, unaware that her phone has just detected a fungal disease requiring dry conditions. Her drone captures early infestation signals but lacks local protocols to offload data, coordinate with other devices, or trigger timely interventions. Maria's devices operate in isolation, revealing the limitations of today's cloud-centric pipeline: coordination failures across heterogeneous endpoints, delayed response under intermittent connectivity, and privacy risks when sensitive household and garden data are processed on third-party servers [11]. Meanwhile, academic research is actively supporting this consumer-oriented shift by exploring low-cost deployments and lightweight edge intelligence. For example, Varol et al. proposed a low-cost network TAP device built on a Raspberry Pi 4 [12], illustrating how consumer-grade hardware can be repurposed for practical monitoring and edge analytics. Li et al. developed lightweight algorithms to deploy pest detection models on smartphones [13], reflecting the broader push toward on-device intelligence. In parallel, cloud–edge–device collaborative computing has been recognized as a promising architecture for smart agriculture, enabling distributed intelligence and adaptive decision-making while addressing the limitations of purely centralized designs [14]. However, consumer-grade deployments introduce a critical gap that remains under-discussed: the security, privacy, and networking challenges that arise when coordinating heterogeneous, resource-constrained devices into a unified intelligent system.

This gap becomes especially salient as collaborative intelligence mechanisms (e.g., federated learning) enter agricultural practice. Cross-silo federated learning has been explored to facilitate data sharing in agri-food supply chains without exchanging raw data, demonstrating clear potential for multi-stakeholder collaboration [15]. Recent work further shows that federated learning can be engineered to be resource-efficient for agricultural disease recognition on Internet-of-Agriculture-Things devices [16]. At the same time, comprehensive reviews emphasize that agricultural federated learning must confront practical barriers, such as data heterogeneity, unreliable communication, limited compute budgets, and stakeholder trust, before it can be widely deployed in real-world ecosystems [7]. These constraints are even more pronounced in consumer environments, where devices rely on unstable home Wi-Fi and battery-powered microcontrollers. Communication overhead thus becomes a primary bottleneck: iterative parameter exchange can dominate training time and drain energy, motivating communication-efficient federated learning strategies that combine device selection, quantization,

and resource-aware transmission [17].

Crucially, security risks do not diminish in consumer-grade systems; they often become more severe due to low-cost hardware, limited protection mechanisms, and weak operational security. Systematic reviews on smart agriculture cybersecurity highlight that attacks on agricultural ICT infrastructures can cause disproportionate societal impact, and identify persistent gaps in detection, evaluation frameworks, and practical defenses [18]. Moreover, agricultural AI introduces model-centric threats beyond traditional IoT compromise. Gao et al. systematically categorize threats to agricultural AI into adversarial example attacks, poisoning attacks, and backdoor attacks, and discuss the tangible risks these attacks pose to applications such as irrigation scheduling and plant disease detection [19]. Therefore, a consumer-grade Agri-AIoT ecosystem that is both intelligent and trustworthy requires security and privacy mechanisms that are explicitly resource-aware and network-aware, rather than retrofitted from industrial systems. To address these challenges, this review advocates a paradigm shift: moving intelligence from centralized cloud platforms to a secure, decentralized, AI-driven network of edge devices. We argue that consumer Agri-AIoT should evolve into an edge-centric ecosystem that deploys secure, lightweight AI models directly on devices and leverages Artificial Intelligence Native Networks (AINNs) [20] principles to coordinate distributed intelligence. Concretely, this paper explores how AINNs mechanisms can be adapted to consumer-grade Agri-AIoT. Key mechanisms include federated learning (FL) for collaborative model building [7, 15, 16], edge–cloud orchestration for computational management [14], and intelligent data flow for optimized communication [17]. We analyze how these approaches operate under tight BOM and energy constraints, while remaining robust against emerging cybersecurity and AI-specific threats [18, 19]. Figure 1 illustrates the conceptual transition from cloud-centric consumer devices to an AINN-coordinated edge intelligence fabric.

This edge-centric direction is dictated by consumer market realities. Unlike industrial equipment, consumer devices face strict cost and power constraints, limited compute and memory, and intermittent connectivity. The core challenge is to balance AI capability with hardware limits: models must be sufficiently accurate for meaningful tasks (e.g., disease diagnosis, pest identification, irrigation control) yet lean enough to operate for extended periods on low-cost microcontrollers or mobile devices without continuous cloud dependence [4, 5, 13]. These constraints also amplify the importance of secure operation, because resource limitations often preclude heavyweight defenses, while user trust is fragile when household data and actionable agronomic recommendations are involved [18, 19]. To underscore this paradigm shift, we explicitly distinguish Consumer Agri-AIoT from Industrial Agri-AIoT across three dimensions: cost, expertise, and resources. Unlike industrial systems driven by enterprise budgets and professional operators, consumer devices face strict Bill of Materials (BOM) constraints and must provide plug-and-play simplicity for non-experts. Furthermore, they rely on low-cost microcontrollers constrained by battery life and unstable home Wi-Fi, contrasting sharply with the high-performance and often managed infrastructure of industrial farms [2, 3]. These disparities confirm that industrial solutions cannot simply be downscaled; instead, consumer-grade ecosystems require lightweight, device-native intelligence and distributed coordination mechanisms that explicitly address communication bottlenecks and trust risks [17, 18].

The main contributions of this article are summarized as follows:

- We articulate the fundamental distinctions between consumer Agri-AIoT and industrial Agri-AIoT, explaining why industrial solutions cannot be trivially downscaled and motivating the need for lightweight, device-native intelligence.
- We provide a comprehensive review of lightweight edge intelligence mechanisms for consumer ecosystems—including model compression, edge deployment, cloud–edge–device orchestration, and federated learning—highlighting how communication-efficient and resource-aware designs are essential for scalability and real-time operation.
- We synthesize the emerging threat landscape for agricultural AI and consumer Agri-AIoT, and map these risks to feasible, resource-aware defenses, forming a coherent security and privacy blueprint for future deployments.

2. Materials and methods

Deploying artificial intelligence on resource-constrained consumer agricultural devices requires model architectures and optimization techniques that explicitly account for limitations in computation, memory, and energy. This section describes the lightweight model families and compression methods considered in this study, together with their typical deployment contexts in consumer Agri-AIoT systems. The methods emphasized here prioritize practicality, deployability, and compatibility with low-cost hardware

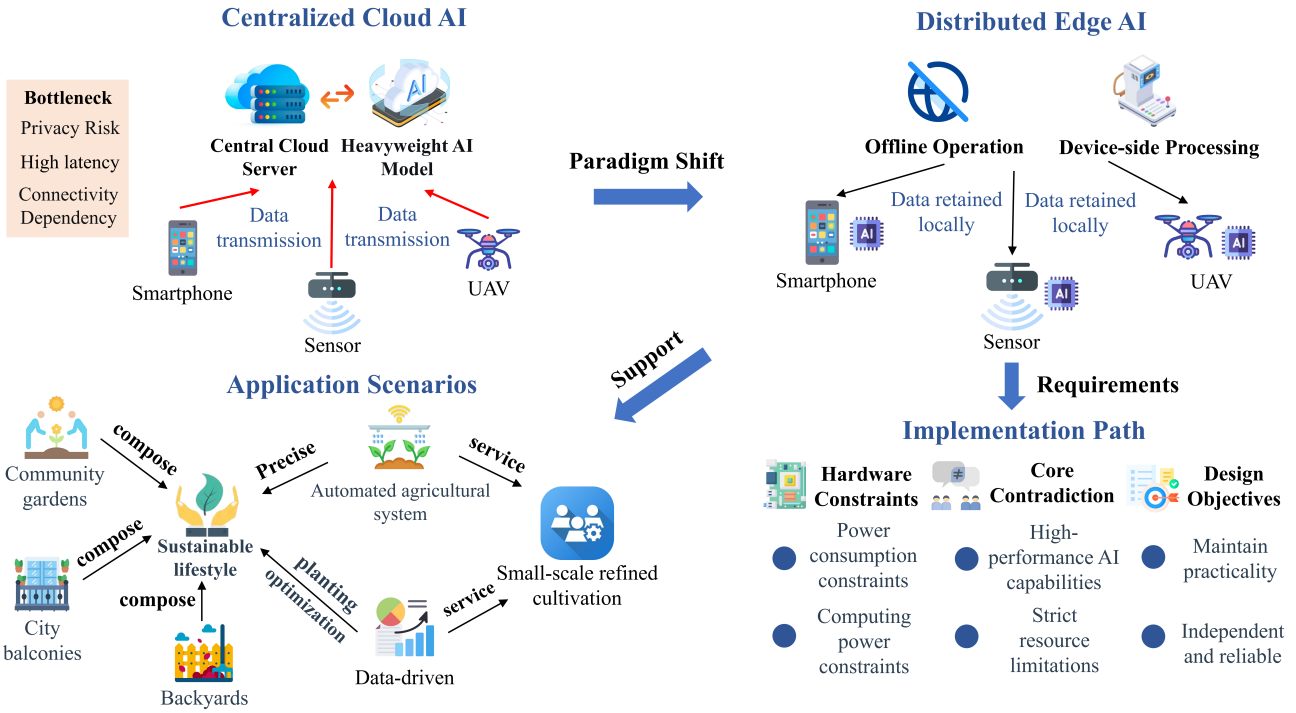

Fig. 1. From cloud-centric consumer Agri-AIoT to secure and distributed edge intelligence: applying AINNs principles (e.g., federated learning, edge–cloud orchestration, and intelligent data flow) to coordinate heterogeneous consumer devices under resource and trust constraints. Core Contradiction: The persistent tension between the high computational demand of modern AI models and the strict power, cost, and memory constraints of consumer-grade hardware.

platforms commonly used in home gardening, small-scale farming, and consumer smart agriculture products.

2.1. Methodological Scope and Device-Centric Metrics

To compare candidate models and optimization strategies in a way that reflects real consumer deployments, we focus on three device-centric dimensions that dominate feasibility on commodity hardware: model footprint, inference latency, and power/energy cost. Model footprint captures not only parameter storage but also runtime buffer requirements that can become the primary bottleneck on microcontrollers. Inference latency is interpreted as end-to-end response time from sensor acquisition and preprocessing to model output and postprocessing, since user experience in consumer applications depends on the full pipeline rather than kernel-level speed alone. Power and energy are treated as practical constraints linked to duty cycles and battery capacity; for always-on sensors, the energy per inference and the frequency of inference are often more relevant than peak power.

2.2. Lightweight Model Selection for Consumer Agri-AIoT

To represent the spectrum of AI capability achievable under heterogeneous consumer hardware constraints, we consider a set of widely adopted lightweight neural network models. These models are selected to span different trade-offs among accuracy, latency, and power, while remaining deployable under tight bill-of-materials (BOM) and energy budgets. Table 1 summarizes representative models, including approximate model size, typical inference latency, qualitative accuracy level, and indicative power consumption, together with their common deployment hardware and target Agri-AIoT tasks. Compact architectures such as MobileNetV3-Small and MobileOne-S1 align well with microcontroller-class platforms (e.g., Cortex-M series and ESP32), supporting always-on sensing and real-time triggering. More expressive models, such as TinyYOLOv4 and MobilePlantViT, are better suited to SBCs or entry-level accelerators where higher throughput enables multi-object detection and richer feature reasoning. Quantized ResNet-18 is included as a reference point for smartphone-class devices that perform periodic or on-demand analysis, providing a familiar baseline when comparing lightweight variants against a commonly used backbone. These mod-

els were explicitly selected to represent distinct hardware tiers in the consumer ecosystem: MobileOne-S1 represents the ultra-low-power MCU class capable of always-on operation, while ResNet-18 serves as a baseline for high-end consumer smartphones where periodic, high-accuracy diagnosis is feasible.

2.3. Deployment Workflow on Consumer Hardware

Translating research models into consumer-ready Agri-AIoT deployments typically requires an engineering workflow that bridges training-time objectives and device-time constraints. In this review, we treat deployment as a pipeline that starts with task specification (e.g., classification, detection, segmentation, or forecasting) and a baseline model trained on data representative of consumer scenarios, such as phone-captured images under variable illumination and cluttered backgrounds. The baseline is then optimized through compression and acceleration, converted into device-supported formats (e.g., TensorFlow Lite or ONNX Runtime Mobile), and profiled on target hardware. Profiling is essential because measured latency and energy depend not only on model topology but also on runtime kernels, memory bandwidth, preprocessing, and postprocessing. Iterative refinement is often required to satisfy footprint and energy budgets while preserving sufficient task utility for end users.

2.4. Model Compression and Optimization Techniques

Several model compression and optimization techniques are considered to enable edge deployment under consumer constraints. These techniques can be applied individually or in combination, depending on the target device class and application requirements.

Pruning: Neural network pruning reduces redundancy by removing low-importance parameters or structures. In consumer deployments, structured pruning (e.g., removing channels or blocks) is often favored because it yields direct latency benefits on commodity inference engines without requiring specialized sparse kernels. Unstructured pruning can achieve higher compression ratios but may not translate to wall-clock speedups unless sparse computation is supported. In practice, pruning is typically followed by fine-tuning to recover accuracy, and the final pruning configuration is chosen to balance memory savings against robustness and generalization under real-world noise.

Quantization: Quantization reduces numerical precision to decrease memory traffic and accelerate inference. Post-training quantization provides a low-friction path to deployment when retraining is infeasible, whereas quantization-aware training is adopted when accuracy preservation is critical under low precision. For consumer systems, integer quantization is especially relevant because it aligns well with embedded integer arithmetic and can reduce energy per inference. Mixed-precision strategies are also used when only a subset of layers is sensitive to quantization noise, allowing most of the model to benefit from low precision while keeping critical layers at higher precision.

Knowledge Distillation: Knowledge distillation enables compact models to inherit the behavior of larger, high-capacity teachers. This method is well matched to consumer Agri-AIoT settings where a teacher can be trained centrally with more diverse data and compute, while the distilled student is deployed to devices with limited resources. Distillation can be implemented by matching soft outputs or aligning intermediate representations, and it can be combined with pruning and quantization to further reduce footprint and latency. In practice, distillation is useful when consumer devices must provide stable decision quality despite limited local training data or intermittent connectivity.

Runtime-Aware Acceleration: Beyond parameter reduction, practical efficiency depends strongly on deployment runtimes and operator implementations. Operator fusion, kernel selection, memory reuse, and hardware delegates (e.g., DSP/NPU offload on mobile chipsets) can materially change latency and energy profiles. Consequently, model selection and compression are considered jointly with the target runtime stack so that theoretical reductions in FLOPs translate into measurable device-level improvements.

2.5. TinyML for Microcontroller-Based Devices

For ultra-low-power platforms, TinyML enables on-device inference directly on microcontrollers such as ARM Cortex-M and ESP32. The main constraints in these settings are limited flash storage for model parameters, very small SRAM for activation buffers, and strict energy budgets under always-on operation. As a result, TinyML deployments typically adopt compact architectures, integer quantization, and simplified preprocessing pipelines. In addition, event-driven designs are common: inference is triggered only when sensors detect meaningful changes (e.g., moisture anomalies or motion events), which reduces unnecessary computation and preserves battery life.

From an AI-native networking perspective, microcontroller-based TinyML nodes can act as intelligent filters by transmitting only salient events or anomalies rather than raw sensor streams. This reduces network load and enables long-duration operation in

Table 1. Comparative Overview of Lightweight AI Models in Agri-AIoT

Model	Size (MB)	Latency (ms)	Accuracy	Power	Typical Deployment Hardware	Agri-AIoT Tasks	Key Consumer Consideration
MobileNetV3-Small	∼2.5	∼25	Medium	Very Low	High-end MCUs (Cortex-M7), ESP32-S3	Weed detection; Pest classification; Plant segmentation	Always-on inference under strict power budgets.
EfficientNet-Lite0	∼4.5	∼45	Medium-High	Low	Entry-level SBCs (RPi Zero), Smartphones	Disease diagnosis; Flower detection; Crop monitoring	Balanced accuracy for consumer diagnostic tools.
TinyYOLOv4	∼10	∼120	High	Medium	Mid-range SBCs (RPi 4), AI Accelerators	Multi-weed detection; Fruit counting; Pest localization	Multi-object detection with moderate energy cost.
ResNet-18 (8-bit)	∼11.0	∼80	High	Medium-High	General-purpose edge CPUs, Smartphones	Fine-grained disease classification; Species identification	Reference model for periodic visual analysis.
MobileOne-S1	∼4.0	∼18	Medium-High	Very Low	DSP-enabled MCUs, Mobile CPUs	Ultra-fast weed detection; Drone navigation	Low-latency inference for real-time response.
MobilePlantViT	∼7	∼70	High	Medium	AI-capable SBCs (RPi 4, Jetson Nano)	Disease diagnosis; Multi-symptom analysis	Transformer-based reasoning for complex scenes.

Note: The qualitative power categories are indicative ranges commonly used in consumer-grade edge deployments: *Very Low* (<100 mW, suitable for always-on or battery-powered devices), *Low* (100–500 mW), *Medium* (0.5–2 W), and *Medium-High* (>2 W). Actual power consumption depends on the specific hardware platform, runtime, and workload characteristics.

consumer environments with unstable Wi-Fi. It also supports standalone consumer products, such as soil sensors and pest traps, that remain functional without mandatory cloud connectivity or a dedicated hub.

2.6. Edge–Cloud–Device Orchestration in Consumer Ecosystems

Consumer Agri-AIoT systems are inherently heterogeneous, and practical deployments often benefit from orchestration across microcontrollers, smartphones, and SBCs. In this review, orchestration is treated as a methodological component that assigns tasks based on resource availability and latency requirements. Microcontrollers are naturally suited to always-on sensing, simple detection, and actuation triggers; smartphones provide interactive, mid-scale vision inference and user-facing decision support; and SBCs or accelerators can sustain multi-object detection, local aggregation, and periodic model updates. Coordinating these roles requires careful choices about where inference runs, when data is transmitted, and how often models are updated, particularly when communication and battery budgets dominate the feasibility envelope.

To operationalize this, we propose a threshold-based decision mechanism. Tasks are processed locally by default to preserve privacy. Offloading to the cloud is only triggered if: (1) local battery reserves exceed 20%, ensuring device longevity; and (2) network latency is below 100ms, ensuring user responsiveness. Conversely, if the confidence score of the local lightweight model falls below a safety threshold (e.g., 70%), the system requests cloud verification regardless of latency.

3. Results and discussion

Our analysis of lightweight AI models in consumer-grade Agri-AIoT systems reveals a persistent tension between computational efficiency and security assurance. Model compression techniques, such as pruning, quantization, and parameter sharing, are indispensable for fitting inference into microcontrollers, smartphones, and low-cost SBCs. However, the same compression steps that reduce latency and energy consumption can also remove architectural redundancy that often functions as an implicit buffer against perturbations and distribution shifts. Consequently, compressed models may present sharper or simplified decision boundaries and narrower safety margins, increasing susceptibility to adversarial manipulation, training-time contamination, and privacy leakage when deployed in the wild.

Result 1: Security risks in AI-native Agri-AIoT systems are amplified—not mitigated—by model lightweighting and consumer-grade deployment constraints. This result contradicts the intuition that "smaller systems are safer." In practice, consumer devices operate under strict compute, memory, and battery budgets, which constrains both model capacity and the feasibility of heavyweight defenses. As a result, the protection surface shrinks while the attack surface remains broad: low-cost sensors remain easy to spoof, wireless connectivity remains intermittent and difficult to harden, and model updates and inference interfaces remain exposed. The outcome is a mismatch in which adversaries can exploit vulnerabilities faster than devices can detect, verify, or recover, turning individual weaknesses into ecosystem-level reliability risks.

3.1. Systemic Threats in AI-Native Consumer Agriculture

As consumer agricultural devices evolve from isolated tools into collaborative AI-native networks, system reliability becomes tightly coupled to the integrity of individual nodes. In such environments, compromise is rarely contained. Errors can propagate through shared models, collective inference, and distributed actuation. As a result, attacks that

target a single weak endpoint can distort group-level behavior across the network.

Data Poisoning: Federated learning (FL)-enabled agricultural applications allow decentralized devices to collaboratively train shared models. Yet the same openness that enables participation also permits malicious or faulty clients to submit corrupted updates. Lightweight aggregators often lack strong integrity verification because robust defenses (e.g., intensive anomaly scoring, multi-round auditing, or robust aggregation under strong adversaries) can be computationally and communication expensive. This creates a subtle but high-impact failure mode: poisoning does not need to crash the system to succeed; it only needs to shift the learned decision boundary in a way that degrades critical classes (e.g., confusing severe disease with benign stress), which may remain unnoticed until incorrect interventions occur at scale. In community-driven scenarios, this risk is amplified by the non-independent and identically distributed nature of consumer agricultural data. Unlike standardized industrial greenhouses, home gardens exhibit extreme variance in soil types, lighting conditions, and plant species. This heterogeneity makes it difficult for the central aggregator to distinguish between a malicious model update (poisoning) and a legitimate update derived from a unique environmental context.

Adversarial Attacks: Vision-based tasks, including plant disease and weed detection, remain vulnerable to adversarial perturbations. Consumer settings expand the feasible threat space beyond purely digital attacks. Physical perturbations (e.g., stickers, printed patterns, or texture manipulations) are realistic because sensors operate in uncontrolled environments with variable lighting, motion blur, and cluttered backgrounds. Lightweight models, particularly those optimized aggressively for latency, may be less robust to such perturbations because they sacrifice representational redundancy for speed. Once an adversarial input is injected, the impact can cascade: erroneous diagnoses can trigger unnecessary chemical treatments, incorrect irrigation adjustments, or false alerts shared across connected devices. Because consumer endpoints often implement minimal on-device validation, they can serve as low-resistance entry points that undermine the integrity of the broader inference pipeline.

Model Inversion and Membership Inference: Cloud-assisted agricultural services commonly expose inference APIs for diagnosis, recommendation, or prediction. Even when raw data are not explicitly shared, repeated queries can enable attackers to infer sensitive features of the training distribution or determine whether a specific user's data contributed to model training. In collaborative ecosystems, such leakage weakens the trust foundation required for sustained participation: users may avoid contributing images or sensor readings if they fear their household environment, crop conditions, or management patterns can be inferred indirectly. The result is a systemic degradation in both privacy and data availability, where privacy risks suppress participation and reduced participation further destabilizes model generalization.

Hardware-Level Threats: Low-cost sensors, actuators, and microcontrollers often lack secure boot, encrypted communication, and strong authentication. This makes them susceptible to firmware tampering, command injection, replay attacks, and falsified sensing. In an AI-native network, falsified hardware inputs are not merely local measurement errors; they can be integrated into shared semantic representations, shaping collective learning signals and downstream actuation decisions. A compromised soil sensor can therefore distort not only one irrigation controller but also the aggregated understanding of environmental state used by neighboring devices, especially when models or policies adapt online. This transforms hardware compromise into a mechanism for steering network-wide optimization toward incorrect states, undermining both safety and resource efficiency.

Result 2: In AI-native Agri-AIoT systems, security threats are inherently systemic, as localized compromise propagates through shared models, collective inference, and distributed control. This result emphasizes that consumer Agri-AIoT security cannot be evaluated purely at the device level. The relevant unit of risk is the ecosystem: devices exchange data, share model updates, and coordinate actions, so integrity failures become contagious. Therefore, the primary design objective is not merely preventing compromise, but ensuring that compromise does not silently propagate and that the network can maintain acceptable behavior under partial trust and partial failure.

3.2. Security as an Enabler of Trustworthy AI-Native Networks

Given the above threat mechanisms, security in consumer Agri-AIoT must be treated as an enabling layer for trustworthy AI-native networking rather than an auxiliary add-on. Figure 2 summarizes the relationship between the identified threats and feasible defenses under resource constraints. To ensure practical deployability, we explicitly map these defensive layers to the hardware constraints outlined in Table 1. Resource-constrained microcontrollers (e.g., Cortex-M7) are limited to lightweight adversarial detection and simple sanity checks, as they lack the memory for complex cryptographic operations. Conversely,

smartphone-class devices and edge gateways can support more intensive mechanisms such as partial Homomorphic Encryption and Secure Aggregation, enabling a tiered defense architecture that aligns security overhead with device capability.

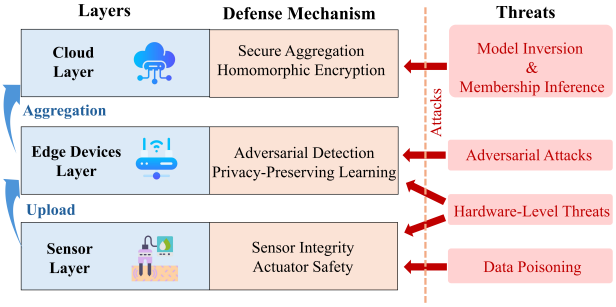

Fig. 2. A layered view of the Agri-AIoT cybersecurity landscape, mapping prominent threats to their targeted system components and corresponding defensive strategies discussed in this paper.

Lightweight Adversarial Detection: Because consumer endpoints cannot afford heavy defensive stacks, lightweight detection methods provide a pragmatic first line of protection. Input consistency checks, noise heuristics, and simple feature-level validation can screen obvious adversarial or corrupted inputs before they enter the shared inference pipeline. These methods do not guarantee security against adaptive attackers, but they can reduce the probability that low-effort attacks propagate across devices. In ecosystem terms, lightweight detection acts as a containment mechanism: it may not prevent all attacks, but it can reduce the amplification factor of a compromised node.

Privacy-Preserving Federated Learning: Federated learning supports collaborative model improvement without centralizing raw data, aligning with consumer privacy expectations. However, FL introduces practical limits in consumer environments: gradient leakage and poisoning remain plausible, and communication overhead can be substantial under unstable residential Wi-Fi. As a result, privacy-preserving FL in consumer Agri-AIoT must be paired with communication-efficient strategies (e.g., update sparsification or selective participation) and trust mechanisms that reduce the influence of untrusted updates. Under these conditions, FL becomes not merely a learning protocol but a coordination substrate whose security properties directly shape network reliability.

Secure Aggregation and Differential Privacy: Secure aggregation prevents reconstruction of individual updates, while differential privacy reduces the risk of inference-based leakage. Together, they provide a trust anchor for collaborative learning, especially when contributors are heterogeneous and partially untrusted. Importantly, these mechanisms must be calibrated against resource constraints: excessive cryptographic overhead or overly strong noise injection can degrade usability or model utility. Thus, their practical role in consumer Agri-AIoT is to provide bounded privacy guarantees while preserving enough signal for meaningful learning.

Homomorphic Encryption for Selective Inference: Homomorphic encryption enables computation on encrypted data and can protect highly sensitive queries in cloud-assisted inference. Nevertheless, its computational cost limits applicability for real-time consumer use. In near-term consumer deployments, HE is better positioned for selective, non-latency-critical workloads, such as periodic analysis of sensitive historical logs or high-value diagnostic queries where privacy risk dominates latency cost. Wider adoption will likely depend on hardware acceleration and more efficient HE schemes.

3.3. Case Study Discussion: FarmBot as a Representative Consumer Agri-AIoT System

FarmBot highlights how consumer-grade robotics can bridge cyber and physical consequences in Agri-AIoT. To rigorously analyze these risks, we categorize the attack vectors and corresponding defenses inherent to this split architecture:

- API Vulnerabilities (High-Level): The Raspberry Pi controller runs a general-purpose OS exposed to network services, susceptible to remote command injection. Defense: Implementation of strict network firewalls and containerized application logic to limit privilege escalation.

- Firmware Tampering (Low-Level): The Arduino microcontroller often lacks secure boot, allowing attackers to flash malicious firmware that ignores safety stops. Defense: Enforcing hardware-based root-of-trust and requiring signed firmware updates for all attached microcontrollers.

- Semantic Risks (Network-Level): A compromised unit acts as a Trojan Horse, injecting physically valid but contextually wrong data (e.g., reporting dry soil during rain) to corrupt shared models. Defense: Deploying neighbor-assisted consistency checks and anomaly detection algorithms at the edge aggregator.

This case demonstrates a generalizable pattern for consumer Agri-AIoT: internal device trust boundaries matter as much as external connectivity. Security must therefore

be rooted in hardware identity and authenticated internal communication, not only in perimeter defenses. Moreover, because consumer devices are increasingly embedded in collaborative AI-native ecosystems, a single compromised robot becomes both a physical hazard and a semantic hazard, capable of distorting collective perception and decision-making.

Result 3: Consumer-grade agricultural robotics highlight the necessity of hardware-rooted trust and secure internal communication as prerequisites for scalable AI-native intelligence.

The results collectively indicate that achieving secure and distributed edge intelligence in consumer Agri-AIoT requires co-design across model efficiency, learning protocols, and trust infrastructure. Lightweight AI expands deployability but narrows defensive margins; collaborative networking improves intelligence but increases systemic coupling; and consumer hardware lowers cost but widens the attack surface. Addressing these tensions demands resource-aware defenses that prioritize containment, integrity of shared learning, and hardware-grounded trust across heterogeneous device fleets.

4. Conclusions

The future of consumer-grade Agri-AIoT relies on integrating security, sustainability, and user empowerment. Trusted Execution Environments provide hardware-rooted protection against OS compromises, while Explainable Artificial Intelligence strengthens trust by visually highlighting diagnostic symptoms. Green AI ensures autonomous long-term operation. A major pain point is the lack of unified security standards across diverse device brands, where a single vulnerable hub can compromise the entire network. Future developments must prioritize cross-brand security frameworks (akin to Matter) to ensure interoperability. Establishing these unified standards is essential for allowing sensors, drones, and smartphones to seamlessly collaborate as a coherent, trusted network without introducing systemic weak points.

Additionally, future research must reconcile inherent conflicts between defense and efficiency. For instance, simultaneously deploying Model Pruning and Differential Privacy creates tension: pruning removes the parameter redundancy often required to tolerate the noise injected by privacy mechanisms. Developing "security-aware compression" frameworks is therefore critical to balancing model sparsity with privacy guarantees. Specifically, we envision a potential framework that treats compression as a multi-objective optimization problem. Unlike standard pruning, which minimizes model size subject only to accuracy constraints, security-aware compression introduces a "robustness loss" (e.g., based on adversarial margins or Lipschitz continuity) into the objective function. This mechanism ensures that parameters critical for tolerating input noise, which are often mistaken for useless redundancy in traditional compression, are preserved, thereby theoretically reconciling the conflict between lightweight design and defense requirements. These directions converge toward an AI-native networking systems paradigm fusing edge intelligence with adaptive communication. Intelligent offloading and reinforcement learning will optimize resources based on dynamic conditions. Lightweight Software-Defined Networking and Network Function Virtualization architectures enable flexible QoS to prioritize latency-critical alerts, while semantic networking minimizes overhead. Together, these innovations foster a secure, self-organizing ecosystem powered by AI-native networking systems principles. Finally, widespread adoption of consumer-grade Agri-AIoT solutions must also account for practical economic and usability barriers, as device cost, maintenance overhead, and limited technical literacy among non-expert home gardeners can significantly constrain real-world deployment and long-term adoption.

References

[1] C. Zhao, J. Li, and X. Feng, (2021) *"Development Strategy of Smart Agriculture for 2035 in China"* **Strategic Study of Chinese Academy of Engineering *23***: 1–9.

[2] M. S. Farooq, S. Riaz, A. Abid, T. Umer, and Y. B. Zikria, (2020) *"Role of IoT technology in agriculture: A systematic literature review"* **Electronics *9***(2): 319. DOI: 10.3390/electronics9020319.

[3] V. K. Quy, N. V. Hau, D. V. Anh, N. M. Quy, N. T. Ban, S. Lanza, G. Randazzo, and A. Muzirafuti, (2022) *"IoT-enabled smart agriculture: architecture, applications, and challenges"* **Applied Sciences *12***(7): 3396. DOI: 10.3390/app12073396.

[4] B. Swaminathan, S. Palani, S. Vairavasundaram, K. Kotecha, and V. Kumar, (2023) *"IoT-Driven Artificial Intelligence Technique for Fertilizer Recommendation Model"* **IEEE Consumer Electronics Magazine *12***(2): 109–117. DOI: 10.1109/MCE.2022.3151325.

[5] D. Muhammed, E. Ahvar, S. Ahvar, M. Trocan, M.-J. Montpetit, and R. Ehsani, (2024) *"Artificial Intelligence of Things (AIoT) for smart agriculture: A review of architectures, technologies and solutions"* **Journal of Network and Computer Applications *228***: 103905. DOI: 10.1016/j.jnca.2024.103905.

[6] L.-B. Chen, X.-R. Huang, G.-Z. Huang, and S.-Y. Kuo, (2025) *"An Orchid Classification Scheme Using Deep Learning for Automated Packaging in Production Lines"* **IEEE Consumer Electronics Magazine *14***(1): 65–76. DOI: 10.1109/MCE.2024.3434910.

[7] R. Dembani, I. Karvelas, N. A. Akbar, S. Rizou, D. Tegolo, and S. Fountas, (2025) *"Agricultural data privacy and federated learning: A review of challenges and opportunities"* **Computers and Electronics in Agriculture *232***: 110048. DOI: 10.1016/j.compag.2025.110048.

[8] I. Sharma and V. Khullar, (2025) *"Blockchain-enabled federated learning-based privacy preservation framework for secure IoT in precision agriculture"* **Journal of Industrial Information Integration *44***: 100765. DOI: 10.1016/j.jii.2024.100765.

[9] Q. Ding, X. Yue, Q. Zhang, Z. Xiong, J. Chang, and H. Zheng, (2024) *"Bc2FL: Double-Layer Blockchain-Driven Federated Learning Framework for Agricultural IoT"* **IEEE Internet of Things Journal**: DOI: 10.1109/JIOT.2024.3485208.

[10] D. Das, V. Udutalapally, and S. P. Mohanty, (2021) *"Consumer Technologies for Smart Agriculture"* **IEEE Consumer Electronics Magazine *10***(4): 49–50.

[11] X. Yi, Z. Zheng, P. Wu, T. R. Gadekallu, and K. Fang, (2025) *"Multi-Scale Tea Bud Grading Detection in Complex Tea Garden Scenes Based on a GenAI Training Framework"* **Computers and Electronics in Agriculture *239***: 111032. DOI: 10.1016/j.compag.2025.111032.

[12] M. Varol and M. İskefiyeli, (2025) *"A low cost compact network TAP device with Raspberry Pi 4"* **Engineering Science and Technology, an International Journal *70***: 102118. DOI: 10.1016/j.jestch.2025.102118.

[13] S. Li, Z. Yuan, R. Peng, D. Leybourne, Q. Xue, Y. Li, and P. Yang, (2024) *"An effective farmer-centred mobile intelligence solution using lightweight deep learning for integrated wheat pest management"* **Journal of Industrial Information Integration *42***: 100705. DOI: 10.1016/j.jii.2024.100705.

[14] P. Yu, F. Teng, W. Zhu, C. Shen, Z. Chen, and J. Song, (2025) *"Cloud–edge–device collaborative computing in smart agriculture: Architectures, applications, and future perspectives"* **Frontiers in Plant Science *16***: 1668545.

[15] A. Durrant, M. Markovic, D. Matthews, D. May, J. Enright, and G. Leontidis, (2022) *"The role of cross-silo federated learning in facilitating data sharing in the agri-food sector"* **Computers and Electronics in Agriculture *193***: 106648. DOI: 10.1016/j.compag.2021.106648.

[16] M. Aggarwal, V. Khullar, N. Goyal, and T. A. Prola, (2024) *"Resource-efficient federated learning over IoAT for rice leaf disease classification"* **Computers and Electronics in Agriculture *221***: 109001. DOI: 10.1016/j.compag.2024.109001.

[17] M. Chen, N. Shlezinger, H. V. Poor, Y. C. Eldar, and S. Cui, (2021) *"Communication-efficient federated learning"* **Proceedings of the National Academy of Sciences *118***(17): e2024789118. DOI: 10.1073/pnas.2024789118.

[18] M. Campoverde-Molina and S. Luján-Mora, (2025) *"Cybersecurity in smart agriculture: A systematic literature review"* **Computers & Security *150***: 104284. DOI: 10.1016/j.cose.2024.104284.

[19] Y. Gao, S. A. Camtepe, N. H. Sultan, H. T. Bui, A. Mahboubi, H. Aboutorab, M. Bewong, R. Islam, M. Z. Islam, A. Chauhan, et al., (2024) *"Security threats to agricultural artificial intelligence: Position and perspective"* **Computers and Electronics in Agriculture 227**: 109557. DOI: 10.1016/j.compag.2024.109557.

[20] Y. Chen, R. Li, Z. Zhao, C. Peng, J. Wu, E. Hossain, and H. Zhang, (2024) *"NetGPT: An AI-native network architecture for provisioning beyond personalized generative services"* **IEEE Network *38***(6): 404–413. DOI: 10.1109/MNET.2024.3376419.