Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X Volume 4, Issue 6, 2026

RECONCILING PRIVACY, EXPLAINABILITY, AND FEDERATED LEARNING IN DECENTRALIZED PRECISION AGRICULTURE

Muhammad Owais[*1], Karishma Lohana[2], Dr. Mughair Aslam Bhatti[3]

[*1,2]*M.S. Data Science at Shaheed Zulfikar Ali Bhutto Institute of Science and Technology (SZABIST), Karachi, Pakistan*
[3]*Department of Robotics and Artificial Intelligence at Shaheed Zulfikar Ali Bhutto, Institute of Science and Technology (SZABIST)*

[*1]msds24101126@szabist.pk

DOI:https://doi.org/10.5281/zenodo.20729331

Keywords
Decentralized precision agriculture, explainable artificial intelligence, federated learning, privacy-preserving machine learning.

Article History
Received: 19 April 2026
Accepted: 01 June 2026
Published: 17 June 2026

Corresponding Author: *
Muhammad Owais

Abstract
Precision agriculture is decentralized, relying on heterogeneous datasets provided by farms, cooperatives, and sensing platforms; nevertheless, these datasets cannot be centrally aggregated owing to privacy, regulatory, and economic limitations. Although deep learning has achieved significant performance increases in crop monitoring and disease detection, it not only uses centralized training, which conflicts with the distributed nature of agricultural data, but also leads to issues of confidentiality, accountability, and trust. Previous studies have mostly studied federated learning, privacy protection, and explainable artificial intelligence individually, with no evaluations of their complexity in situations that are defined by non-IID data, sparse connectivity, and edge-computational limitations. We provide an analytical synthesis in this study, where we consider federated learning, privacy, and explainability as design requirements that are mutually dependent. We propose a common taxonomy of architectures, federated learning frameworks, privacy preservation methods, explainability methods, data formats, and deployment platforms. Through comparative analysis, we reveal trade-offs between predictive accuracy, interpretability fidelity, communication overhead, and privacy robustness, and new challenges, especially the instability of explanations, lack of auditable and decentralized benchmarks, and trade-off between privacy and utility.

I. INTRODUCTION

A. Societal and Technological Motivation

Agriculture is increasingly being modeled as a data-driven model of operation with crop well-being, input efficiency, and risk being inferred instead of being observed by traditional scouting. What appears to be a simple scaling issue at first face value is, in reality, a governance problem: agronomic data is created by various actors, are not evenly distributed, and have different values across the supply chain. Recent agricultural data privacy research indicates that the desire to share data by farmers is limited by confidentiality, economic risk, institutional trust, and only on model performance [14]. These limitations make the central data lake assumption structurally weak in the agricultural context because the data do not simply exist in different locations; they have different obligations, liabilities, and bargaining power. This fragmentation can be seen especially in the case of aerial and remote sensing pipelines because, despite the ostensibly benevolent motive

of gathering data, the issue of privacy remains applicable even in the context of research activities [78]. Consequently, the prevailing deep-learning community paradigm of centralized training and subsequent optimization does not clean up agricultural realities.

B. What Current AI Solutions Attempt

Technical response has also been much-needed to enhance model capacity and reduce manual feature design, with deep learning being the horsepower in disease recognition, field monitoring, and similar phenotyping. As shown by existing and modern research, this pattern of investigation has been appealing: deep neural models, when trained on sufficiently representative images, not only attain high diagnostic accuracy and surpass previous pipelines utilizing handcrafted features, but also do so more than earlier pipelines [34]. The architectures that are more recent continue to build on these benefits with attention mechanisms and multi-scale reasoning, which are more effective when symptoms are visually subtle or co-occurring [1]. Parallel work maintains that what remains practically difficult is no longer the presence of high-performing models but the possibility of implementing them in different farms and seasons without re-collected or re-labeled data having to be collected afresh [73].

Federated learning has been considered in response to the limitations of data sharing, and can be used to train shared models while preserving farm data on local devices, effectively delegating the training task to a distributed set of players [7]. One such complementary reaction has been the inclusion of post-hoc explanations with the model outputs, which has been driven by the anticipation that agronomic stakeholders would require interpretable explanations, rather than predictions, when making decisions that influence yield and cost.

C. Why These Solutions Fail in Agricultural Reality

The failure modes manifest in the divergence of agricultural data of the neat assumptions on which most benchmark-driven AI is based. First, the farm data are non-IID in a practical rather than purely statistical sense: fields vary by cultivar, soil, microclimate, management, sensors and labeling practice; thus, it is not necessarily true that having more clients would lead to better generalization. Although federated learning may be successfully applied to plant disease recognition, performance stability can drop once client distributions become sharply fragmented, and convergence behavior may be sensitive to patterns of participation that would be viewed as small nuisances in well-provided areas [62]. Second, rural deployment comes with constraints that most ML papers consider as after-thoughts, such as intermittent connectivity, limited computing hardware, and client capability that is asymmetric, which makes training cadence and model-update schedules difficult [87]. Third, privacy protection is not a checkbox; formal defenses may also have new vulnerabilities or trade-offs with operational consequences, such as gradients or updates leaking information or protection mechanisms that decrease the utility of models trained on already low-resource signals [75]. Lastly, it is impossible to make agriculture explainable as a glossy heatmap, since it is a precise explanation that is used to adjudicate accountability between agronomists, farm managers, and automated decision-support systems. Explainable deep learning in the context of explainable plant phenotyping highlights a longstanding inconvenience: explanations can seem plausible and yet remain fragile to a distribution shift, which is exactly what agriculture is subjected to [49]. This heterogeneity, poor infrastructure, and high-stakes accountability form a combination that makes the design of accuracy-first complete.

D. Fragmentation of Existing Literature

A more serious problem is that all three pillars deep learning, federated learning, and explainability are treated as individual and independent modules that can be stacked together without the need to fundamentally rethink the assumptions underlying the system. Available surveys on deep learning and task-specific studies have a tendency to predetermine architectural design and quantitative performance improvements, and often give little emphasis to the role of composition and heterogeneity of training data in generalization and slowing down

the discovery of deployment risks until expensive, post-hoc validation phases [63]. In most cases, even with rigorous technical analysis, surveys on federated learning are often described in a domain-neutral register, thus obscuring agronomic limitations, such as seasonal drift, client imbalance, and accountability requirements that go beyond standard conceptions of device heterogeneity [56]. In a similar vein, the explainability literature in the agricultural domain is also overwhelmingly concerned with interpretation methods but offers little criticism of the question of whether the explanations can remain reliable in the context of decentralized optimization or with the introduction of privacy-conserving perturbations, which are the conditions where explanations become most susceptible to challenge [12]. By intertwining these strands, they tend to be produced into a narrative structure instead of being challenged as a combination of mutually binding factors, and fail to meet the demands of field deployment. The fact that the latest studies on federated, explainable AI in smart agriculture explicitly describe integration is worthwhile; however, the wider area is still lacking in the systematic explanation of integration failures and the causes of failure persistence [4]. Research on cross-silo federated learning in the agri-food sector shows that the key bottleneck is often the institutional structure and incentive alignment as opposed to the learning algorithm itself [92]. Such observations give reason to a review not simply to list methodologies but also to critically survey the tensions that define whether such methodologies can coexist.

E. Research Gaps and Contributions

The existing body of literature in surveys that focuses solely on deep learning gives precedence to architectural novelties and peak accuracy metrics and puts concerns on data governance, deployment restrictions, and operational viability at a backburner or in the future. This is particularly troublesome in agriculture, where normal fluctuations in cultivar, seasonality, sensing modality, and field conditions essentially determine whether models can be trained, transferred, or maintained. As a result, the goals of the optimization of the DL-only surveys are in marked contrast to the realistic failures observed in the field, thus continuing to maintain a wide gulf between the performance of the algorithm and field reliability. Surveys concentrating on federated learning alone introduce an alternative, but no less important limitation: concentrating on farming actors as generic clients, FL-only syntheses obscure the agronomic importance of heterogeneity, including variations in crop varieties, management, climatic conditions, and the quality of data. The abstraction of this nature can be beneficial to theoretical analysis, but underscores the need for accountability, validation and detection of failure in the field. Convergence instability, imbalance of fairness, and failure of robustness are all revealed in agriculture earlier and more conspicuously than in many other areas, which makes such simplifications very expensive to implement. Similarly, explainable AI surveys confine explanation as a post-hoc interpretive layer that may be added to a model when it is trained. They often do not deal with the stability of explanations in the case of decentralized training executions, heterogeneous clients, or privacy-respectful perturbations. Explanation instability is not a cosmetic issue in agronomic decision making, where inconsistent attribution can switch up advice, misplace blame, undermine agronomic trust, and ultimately hinder the uptake of interpretability in making high-stake decisions by agronomy practitioners. These failures are distinctly enhanced by agriculture because the field is highly cohesive to the variability of biology, infrastructural constraints, and financial vulnerability. Privacy restrictions, explainability requirements and dynamics of learning clash in a less obvious manner than in areas where data streams are stable and governance is centralized. Thus, agriculture can be viewed as a stress test that reveals hidden weaknesses in the DL, FL, and XAI views only when taken separately.

This study provides four contributions to the analysis. First, it establishes a coherent taxonomy that locates deep learning structures, federated learning models, privacy solutions, explainability methods, and deployment environments in an agricultural system space. Second, it offers a cross-dimensional synthesis, which clearly exposes the trade-offs between robustness, scalability,

deployment realism, interpretability reliability, and privacy vulnerability, instead of making individual performance claims. Third, it enumerates and formalizes the most common forms of failure, especially those caused by non-IID federation, explanation drift, and privacy-utility collapse, in such a way that negative outcomes become subject to analysis instead of being swept under the carpet. Lastly, it defines the implications of governance and deployment that bridge the gap between technical design decisions and data sovereignty, accountability, and adoption limitations in decentralized agricultural systems, re-integrating algorithmic decision-making with real world existence and institution setting.

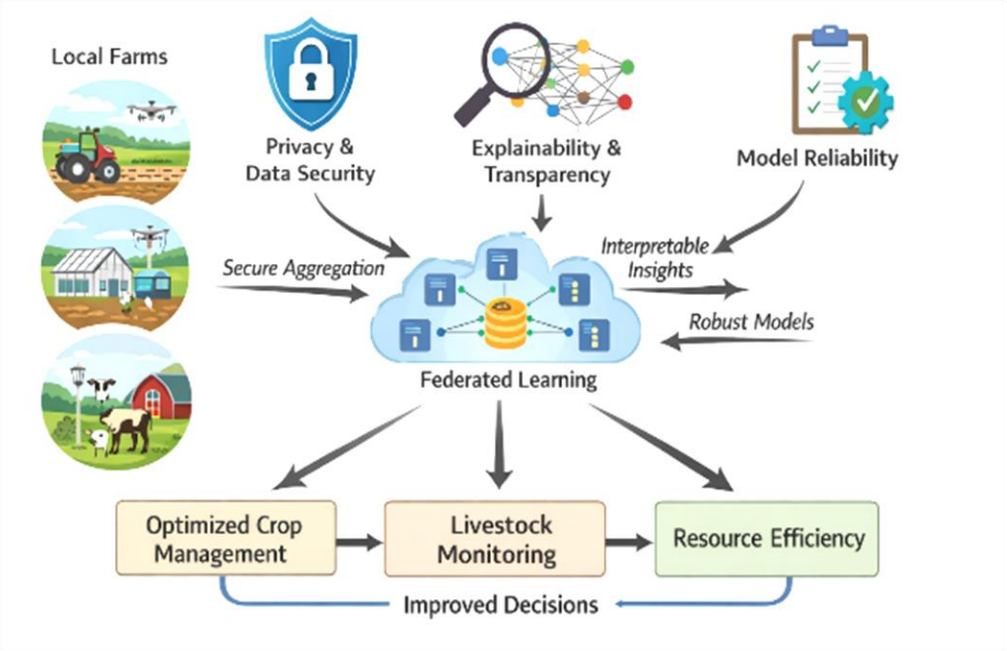

FIGURE 1. Interaction of privacy constraints, explainability reliability, and federated optimization under decentralized precision agriculture.

Figure 1 has been added to eliminate a common error of interpretation in this literature where privacy, explainability, and federated optimization are additive concepts that can be added separately. Instead, the figure 1 represents them, as two tied-up constraints, in which the reinforcement of one dimension may cause instability in another, in a way that is not immediately apparent in the case of accuracy curves alone, such as an explanation drift in a non-IID federation or a utility degradation in privacy perturbation. Making these dependencies explicit initially allows further discussion to judge of what previous work has done and what limitations it has had in a sensible view of systems, but not as a series of independent methods.

II. Background and Conceptual Foundations

This section provides the conceptual scaffolding required to read the rest of the review though not to cover familiar ground again, but to retain focus on the points that are technically problematic when learning is decentralized and privacy is non-negotiable, as well as explanations are operational specifications and not optional visualizations. The overloaded terminology used within communities, especially privacy, robustness and explainability, is explained here so that it can be consistent with the synthesis and trade-off analysis therein.

A. Deep Learning in Precision Agriculture

Deep learning in agriculture has also developed in two convergent directions, namely, diversification of task classes and increasing model capacity,

where the discipline swings between pragmatism in engineering and architectural enthusiasm. The most salient family is disease and pest recognition, which can be naturally translated into supervised image classification and provides immediate operational value in case labels are present [28]. Recent systems with strong performance are increasingly based on attention-based representations, such as transformer-based pipelines to reduce both symptom ambiguity and background clutter, which easily overwhelm shallow convolutional priors in field settings [2]. In addition to disease diagnosis, remote sensing of land cover and land-use classification has been incorporated into agricultural monitoring, where the models have to grapple with the issue of spatial heterogeneity, seasonal drift, and multi-resolution sensing; this does not tackle the problem of pattern recognition, but that of scene understanding [5]. Segmentation tasks dominate the center of this shift as they force models to identify agronomically meaningful objects as opposed to simply labeling them, and encoder-decoder networks are still considered standard hubs of that family [95]. Other segmentation families focus on different inductive biases and deployment costs; encoder-decoder architectures that are optimized to support real-time inference remain applicable in systems where agricultural robotics or on-device constraints dominate system design [90]. Modern segmentation systems, designed with consideration of semantic faithfulness with tricky scenes, are a valuable conceptual reference when agricultural tracking requires sensitivity to boundaries and context conservation instead of a crudely defined region [39]. The practical implication is that success in agricultural deep learning is often contingent on a combination of benefits that a single architecture can have but not on a single best architecture.

B. Federal Learning Fundamentals

Federated learning can best be thought of as a training regime that moves the optimization to a distributed coalition (as opposed to a centralized curator), but the tough questions of coordination, verification and the risk of failure remain open. In agricultural contexts, the cross-silo federation of cooperatives, farms, agribusiness units, or research stations is the most common example because they are relatively stable and the data have identity stewardship even when incentives are not well aligned [92]. Federation across devices is conceptually attractive to sensor-intensive agriculture but with more severe computation, connection, and participation-churn constraints, making it less reflective of most production systems, even as a narrative of massive learning in the Internet of Things [66]. The fundamental technical complication lies in the fact that non-IID data is not a nuisance parameter, but a first-order characteristic of agriculture, because farms are structurally different in cultivar, management, microclimate, sensing hardware, and labeling conventions, and gradient directions may represent different agronomic regimes [57]. The personalization strategies are usually manifested as the apparent solution, but they also create a slight conflict: a personalized model may be useful locally and at the same time cease to be readable as a common artifact because stakeholders demand consistent reasoning across locations [13]. The risks of security and privacy add to the situation even more because even non-raw-data-based updates of the models have the potential to leak sensitive data or can be used in manipulation; therefore, federated learning is more related to security engineering than distributed SGD with communication constraints [10]. As a result, federation in agriculture cannot be merely a workaround in data sharing because it recreates the meaning of generalization in the face of no single population distribution about which all stakeholders concur that it should be represented.

C. Privacy-Preserving Computation Paradigms

Privacy in decentralized learning has been described as a combination of techniques, but operationally, it is a combination of trade-offs between confidentiality, utility, latency, and verifiability, where various mechanisms fail in different ways. Differential privacy (DP) is appealing because of its capacity to provide a formal language of privacy loss, and can be implemented without adding heavy cryptography infrastructure. However, the noise necessary to provide meaningful guarantees can obstruct agronomically delicate signals (especially when

data are already sparse or weakly labeled) [3]. Secure aggregation and associated cryptographic schemes ensure that updates cannot be inspected by aggregators or attackers but also add coordination overhead and computation cost which may not be disproportionate in rural or edge limited settings [8]. Multi-party variants of secure aggregation also strive to minimize this overhead, but the additional complexity of the protocol introduces new points of vulnerability that are seldom stress-tested under realistic agricultural network conditions [18]. Multi-key aggregation protocols enhance the flexibility and coverage of threats but the additional key-management overhead and the sensitivity of protocol implementation can become obstacles in deployments where participants differ significantly in terms of IT maturity [88]. Fully homomorphic encryption permits computation on encrypted updates with strong confidentiality properties, but is currently limited by the ability to scale its latency and resource usage to the level of training a small coalition [89]. Verifiable aggregation proposes an added design goal, aggregate integrity, where clients can verify that the aggregation is correct without learning about private updates, the goals of which are significant when federation across institutions is limited by mutual mistrust. The general message is that formal privacy does not necessarily relate to deployable privacy, because the most stringent test of mechanisms can be the least practical in situations where privacy is most needed.

D. Explainability as a Socio-Technical Requirement

E. Explainability is often discussed in an agricultural setting as a means of easy visualization, but its essential agenda is governance, since explainability is used to support the justification of interventions, attribute blame, and justify decisions when stakeholders contest their results. In decision support, such as crop recommendation, explainability is connected to farmer agency and institutional responsibility, since the adoption of the recommendation is determined by the correspondence between the justification and agronomic intuition and risk aversion [43]. Therefore, visual salience cannot be brought down to the quality of explanations; plausible-looking explanations need not be stable, faithful or consistent over time and place, just when farmers and agronomists require them. Recent studies of explainable deep learning in smart agriculture highlight the practical need to associate interpretable explanations with model outputs, but also reveal how exposure to changes in data distribution and training dynamics, which are natural in decentralized environments undermine the practicality of explanations [50]. The introduction of governance-related solutions, that is, the use of explainability, integrity, and traceability solutions, is part of an unspoken understanding that agriculture is not just another application area, but a high-stakes socio-technical system where trust has to be designed, as opposed to assumed [65]. In this regard, this review considers explainability as a reliability property that has to survive non-IID federation and privacy limitations, instead of a post-hoc auxiliary that is added to model training.

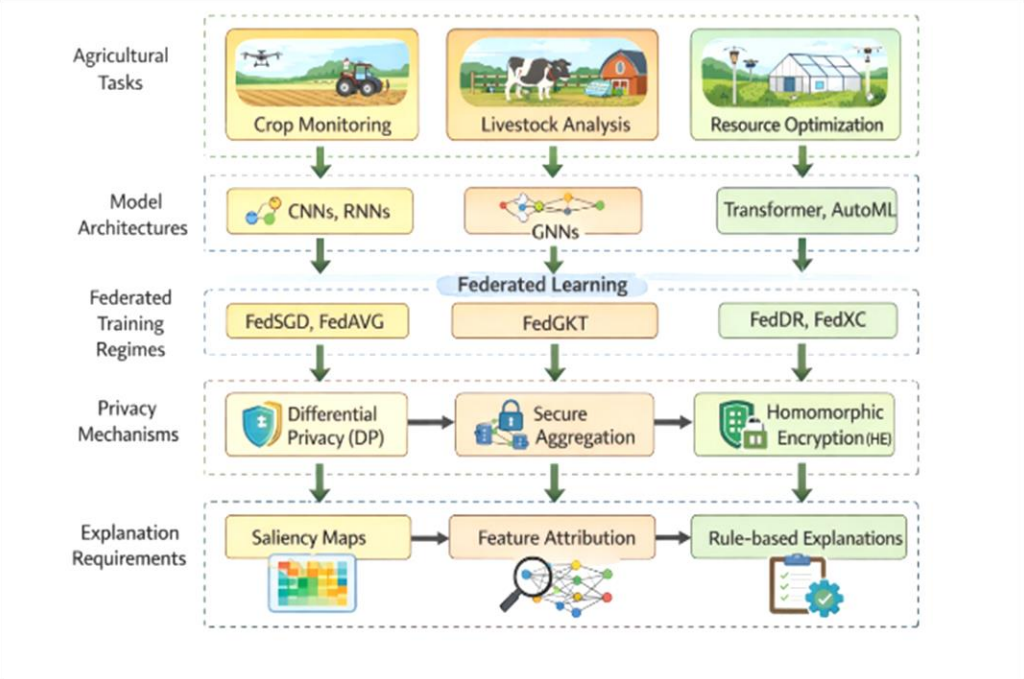

FIGURE 2. Conceptual stack linking agricultural task classes, model architectures, federated training regimes, privacy mechanisms, and explanation requirements across deployment constraints.

Included in Figure 2 is a dependency that is often implicit in the literature: architectural decisions cannot be made without federated training regimes, and frameworks cannot be made without privacy and explanation objectives once deployment constraints are taken into account. The figure 2 gives a brief map of how the decisions made during upstream, i.e. task formulation, sensing modality, and architectural inductive bias, condition the downstream feasibility, i.e. federated optimization stability, admissible privacy mechanisms, and explanation reliability. This framing avoids one of the pitfalls of reviews, namely assessing privacy, federated learning, and explainability as modular add-ons, and encourages the following synthesis, in which each of the three dimensions is reviewed as constraining the other.

III. Review Methodology

The central limitation that informed the methodology design of this research was that current reviews in the field of agricultural artificial intelligence tend to effectively enumerate techniques but not to describe why so allegedly sound methods fail to perform when implemented in the context of decentralization, privacy limitations, and accountability. Based on this, the methodology was designed to favor analytical traceability and comparative arguments instead of exhaustiveness so that every study included would add to an understanding of systemic tensions and not just bloat the coverage.

A. Review Design Rationale

The systematic-analytical hybrid design was chosen because very systematic protocols are more likely to accept formal logic of inclusion to disfavor interpretive richness, whereas very narrative reviews are more likely to experience untraceable biases of the author in synthesizing heterogeneous technical areas. Such a combination enables controlled corpus building and simultaneously, allows the flexibility to probe contradictions between learning performance, privacy assurances, and interpretability reliability that are not manifested by checklist-based synthesis alone [9]. Symbolic PRISMA-style reporting was deliberately avoided because its flow-diagram abstraction is capable of expressing procedural rigor but not the operationalization of how qualitative judgments (e.g., deployment realism or explanation validity), were operationalized in technical assessment [73].

Methodological transparency would be more effectively promoted in the agricultural AI community, where research studies vary radically in sensing modality, scale, and governance assumptions [14].

B. Data Sources and Search Strategy

Five major scholarly databases, namely IEEE Xplore, Elsevier ScienceDirect, SpringerLink, MDPI, and Wiley Online Library, were chosen because they contain high-impact publications in the field of federated learning, explainable AI, and agricultural informatics [56] and were used to conduct literature retrieval. Search queries were formulated based on Boolean combinations of terms that cut across learning paradigms, governance mechanisms, and agricultural settings (e.g., federated learning, explainable AI, privacy preservation, precision agriculture, crop disease, remote sensing, and decentralized systems) [7]. To prevent contamination of foundational algorithmic work with up-to-date deployment-oriented literature, a time frame was put between 2015 and 2025 to represent both the early adoption of deep learning and recent federated and privacy-conscious extensions [34]. The actual execution of the query and consolidation of results were performed repeatedly to reduce bias based on the database being considered as the end result, and duplicate records were sorted out before the screening process commenced to maintain the integrity of the corpus [82].

C. Inclusion Criteria

Only studies that showed particular relevance to agricultural or agri-food systems, and not agriculture as a transitory application scenario, were considered [73]. The depth had to be technical, that is, the works included needed to explain learning architecture, training regime or privacy or explainability mechanism and had to be specific enough to aid comparative reasoning and not conceptual assertion [20]. Deployment realism was considered a first-order criterion, where studies that recognized data heterogeneity, infrastructure constraints or stakeholder considerations were preferred over those that reported only results based on a sanitized benchmark dataset [92]. Peer review was compulsory, limiting the selection of journal articles and IEEE-indexed conference proceedings to maintain methodological responsibility [11].

D. Exclusion Criteria

Data sets based solely on toys or highly edited datasets were also not included because this sort of environment suppresses the heterogeneity that characterizes the risk of agricultural deployment [63]. Articles that address federated learning or explainable AI without a fundamental agricultural basis were filtered out so that they do not bring assumptions into the domain where data governance and accountability pressures are fundamentally different [83]. When they were not placed empirically or evaluatively worded that could be placed in the comparative synthesis later in the review, concept-only and position papers were excluded [30]. Peer-reviewed preprints were also not included in the search to avoid overpowering the effect of speculative assertions that had not been subjected to technical analysis [56].

E. Quality Assessment Rubric

In the context of this research, it is important to note that each of the incorporated studies was assessed using a four-dimensional rubric that tried to highlight not only the strengths of a particular study but also to reflect its weaknesses. The former is the first dimension, data realism, which estimate how far the data sources reflect the real field conditions, seasonal variability, and sensor noise typically experienced in working agriculture [78]. The second dimension, evaluation rig, questioned the baseline comparisons, experimental design and reporting practices, with a special focus on whether the claims of the study were applicable beyond a single crop, season or site [69]. The third dimension was privacy claim validity, which looked at whether the stated privacy guarantees were formalized, empirically validated, or simply assumed, and whether possible leakage routes were considered or tested empirically [75]. The fourth dimension of explainability coherence measures the relationship between the interpretive outputs and the underlying model behaviors and agronomic relevance, instead of being put forward as visually plausible but analytically inert artifacts

[49]. The rubric purposely avoided the use of numbers to allow spurious precision; instead, it enabled a qualitative comparison between different methodological traditions [12].

F. Biases and Methodological Limitations

Various possible biases were also predicted and minimized by design choices instead of post-hoc actions. The biases towards accuracy-biased outcomes in many cases remain rare in agricultural AI, thus remaining biased towards positive results in the literature, not reporting failure modes that are essential in realistic deployment [73]. In various studies that depended on a small number of publicly available datasets of plant disease, dataset-reuse bias was found; this tendency enhances the perception of progress and masks constraints in generalization [63]. The danger of marginalizing views of smallholders or agricultural systems with infrastructure constraints that cannot generate data is in the regional dominance within the literature, especially that of data-rich research institutions, where the pressures of decentralization are most intense [79]. Although that the hybrid review design will reduce some types of bias, interpretive synthesis is bound to be perceptually subjective, which cannot be fully objective without losing the level of analysis [92].

G. Replicability and Auditability

To enable auditability, search queries, database selections, inclusion criteria, exclusion criteria, and quality dimensions were consistently established and implemented before synthesis across the corpus [11]. Metadata extraction followed a standardized schema that had task type, learning paradigm, privacy mechanism, explainability approach, dataset characteristics and decomposition of deployment assumptions, and thus allowed traceable comparisons [53]. Coherent computational reproducibility is an outstanding difficulty in this area owing to unregulated code, data, and hyperparameter release, especially in agricultural research involving privacy issues that sharing could violate governance policies [14]. However, the corpus building to analytical synthesis methodology is well documented, and the review procedure can be replicated independently, even when experimental artifacts are not available [92].

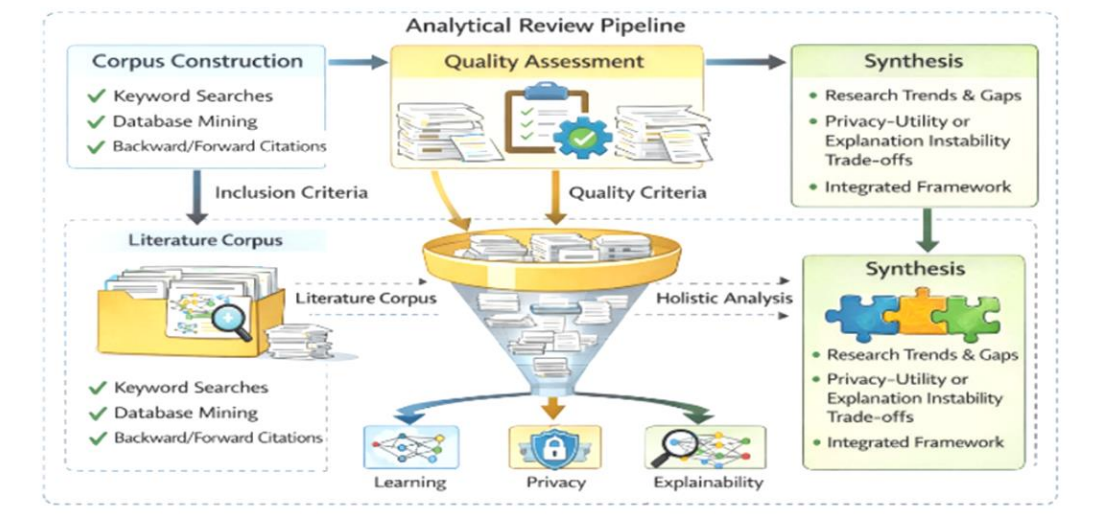

FIGURE 3. Analytical review pipeline illustrating corpus construction, quality assessment, and synthesis across learning, privacy, and explainability dimensions.

Figure 3 clearly explains how methodological choices determine the conclusions made in the literature and, as such, disprove the idea that synthesis is produced automatically through the size of the search. As the figure 3 illustrates, inclusion, exclusion, and quality assessment is a

filter that not only foregrounds structural tension, including privacy-utility trade-offs, or the instability of the explanation, but also entirely blocks incremental performance comparison. The figure 3, by presenting the review as a review pipeline as opposed to a review funnel, supports the interpretative rigor and auditability focus of the study, in contrast to symbolic adherence to procedures.

IV. Taxonomy of Models and techniques

The methodological chapter determined the standards of evidence selection and screening; however, it did not, institute a technical structure to the discipline. Thus, this chapter fills this gap by building a taxonomy that is functionally designed to fill a purpose and is not just popular. These capabilities, underlying assumptions, and limitations of a category can be defined when it is decentralized, faces privacy constraints, and meets the interpretability requirements. In contrast to traditional survey listing models and datasets, the aim here is to slice the literature at its critical points: architecture inductive bias, federated optimization structure, privacy-mechanism class, explanation modality, data modality, and deployment substrate; thus, allowing other subsequent chapters to compare trade-offs without reducing the discussion to histories of better accuracy [6].

A. Deep Learning Architectures

Families of convolutional neural networks (CNNs) are still predominant in agricultural vision pipelines because of their comparative data efficiency in cases with limited annotation budgets and their ability to retain local texture features that, in many cases directly relate to disease lesion and pest damage indications [19]. Transfer learning based on standard CNN backbones in agronomic practice is particularly appealing because it provides a viable realistic route to deployable performance in situations where it is economically infeasible to have large, balanced, and seasonally diverse field datasets [22]. In the case of architectural customization, the emphasis is usually moved to channel attention, pruning, and other efficiency-conscious adjustments that are not oblivious to deployment constraints; nevertheless, they maintain convolutional inductive priors [29]. However, one continued weakness of CNN-based research is that it is contingent on dataset regimes; the improvement in results reported in studies that use small or handpicked training distributions is often due to improved fitting and not true agronomic generalization [31].

The use of families of vision transformers has become popular in plant disease detection because of their ability to capture global conditions and long-range interactions that can help deconflict similar manifestations of disease that differ in location or comorbidity [35]. Nevertheless, the adoption of transformers in the agricultural sector cannot be boiled down to a performance narrative, it is inextricably combined with data and computing considerations, and it is far less likely to provide rich labels and homogenous imaging conditions. Accordingly, models built on a transformer are frequently highly regularized, with an augmentation regime thoroughly crafted, or trained with domain constraints, to prevent a high-performance that breaks down when curated imagery is replaced with in-field cameras [37]. The contrasting class of engineering trade-offs, dubbed hybrid architectures, involving convolutional stems and attention-based or multi-scale reasoning modules, is more characteristic of agriculture, where the representational capacity of an explicit stack of transformers is sufficiently large to represent complex symptoms at the cost of either the training or inference costs of such a stack [38]. Taxonomically, these hybrids are less important because they illustrate the way (deliberately) to tune inductive bias and maintain local texture sensitivity, but inject the global context only where the task structure requires it [40]. Finally, even architectural labels do not yield enough information to infer edge suitability; not only do counts of parameters determine feasibility, but also memory access pattern, latency variance, and resilience to noisy sensor pipelines [60]. This strain is particularly clear in segmentation and dense prediction models, on which some central agricultural systems adjacent to robotics rely, where localization by space is no less important than classification but is fast to become edge-unfriendly unless made edge-friendly. [61]

B. Federated Learning Paradigms

In agriculture, federated learning is commonly conceptualized based on FedAvg-type aggregation. However, the fundamental premise of agronomic comparison of client updates breaks down when client distributions represent real agronomic variation and not sampling noise [16]. Practically, FedAvg agricultural deployments only work well with reduced heterogeneity through either selective client choice or limiting the work to a very restricted set of acquisition conditions [17]. As a result, FedAvg is gradually being recognized in the literature not as a solution but as a diagnostic tool: the point when FedAvg fails to be viewed as an indicator of how large a divergence of clients must be overcome by any practical agricultural system [23].

The solution to non-IID behavior is often suggested as personalized federated learning, although its extension to the agricultural sphere is indirect [25]. On the one hand, personalization can enhance local performance, but on the other hand, it can also diminish cross-site comparability, which is important when stakeholders view the results of model outputs as collective agronomic knowledge as opposed to local predictions [41]. One major taxonomic difference is that between light-touch personalization, such as fine-tuning task-specific heads, and structural personalization, such as clustered or meta-learning approaches, since the second radically alters the expectation of global governance on what is collective and what is not [42]. A decentralized and peer-to-peer federated variant is commonly driven by the wish for single points of failure and reliance on trusted coordinators, which is attractive to agricultural settings with instantiated institutional trust inequality [48]. Nonetheless, there are trade-offs with decentralization, such as the complexity of coordination, increased difficulty of validation, and increased client drift, which become operationally binding enough to become more operationally constraining than the centralized coordinator it aims to substitute [55].

C. Privacy-Preserving Mechanisms

Noise-based privacy schemes, especially differential privacy, are conceptually attractive, as they offer tunable privacy loss, but are not free. Injected noise is in direct conflict with agronomic signal fidelity, particularly in situations with rare diseases or with symptoms in the early stages of progression, where subtle visual signals can be of disproportionate significance. One important difference in the taxonomy is the enforcement status of privacy, locally versus globally, or a combination of both strategies, as each option redefines the threat model and changes the position where utility collapses when the client conditions are heterogeneous [52]. Privacy encoding techniques seek to maintain utility, in which representations are altered before they are shared, thereby shifting the privacy-utility trade-off between optimization dynamics and feature construction.

Cryptography, such as aggregation assumes that the aggregator is potentially curious and secures the updates appropriately; however, it has the drawback of coordinating and communicating poorly with the number of clients and dropouts, which is more common than unique case in rural deployments. Cluster aggregation is a secure aggregation method that proposes an explicit scaling approach that restricts the extent of aggregation prior to global combination; however, it does not change the attack surface and adds a new trust assumption in clusters. Homomorphic encryption and multi-key forms provide a higher level of security of confidentiality, but systems taxonomy should be considered as an operational mechanism rather than a purely cryptographic concept, as they are limited by latency constraints and client-side computation limits. Multi-party and multi-homomorphic designs also increase flexibility but the complexity of protocols itself becomes a barrier to deployment, especially when there is an uneven distribution of technical capacity among stakeholders. A hybrid privacy architecture that merges secure aggregation with verifiability is more than confidentiality, but also integrity, in recognition of the fact that malice aggregation is a real threat in economically sensitive agricultural partnerships. The taxonomy, therefore, differentiates privacy versus inspection and privacy versus being right, which indicates that trust failures consist not only of data leakage [58].

D. Explainability Techniques

Agronomic uses of gradient-based visualizations are also common because of their intuitively map-like behavior; however, their validity must be treated in the same way as a hypothesis and not a fact, especially under conditions of distributional shift [24]. Explanation-based training can be applied to reduce the mismatch between the attention of a model and the desired visual cues of an agronomist, when the models are specifically designed with interpretive constraints; however, this correspondence often creates other sensitivities during training [44]. Attribution-based explanations should provide better-structured explanations than raw saliency maps but are susceptible to gradient perturbations in the input as well as to the stochasticity of the model a weakness exacerbated by the presence of privacy mechanisms that corrupt gradients or by the fact that federated updates vary across rounds. Issues of continuity of empirical assessment have proven that the quality of explanation should be measured qualitatively and quantitatively because, by being visually plausible, there may be a lack of consistency in explanations among samples and clients [46].

Attention explanations are also supposed to be naturally interpretable, and this is not always the case: attention can be faithful without being stable and faithful attention need not be helpful in making decisions in a human context [59]. When using an agriculture-oriented orientation of transformer pipelines, attention maps can be used to make disease cues contextual on a larger spatial scale, and the attention maps are expected to be justified by agronomic significance and not to be taken as self-explanatory artifacts [71]. Current research on the diagnosis of plant diseases emphasizes that the reliability of the explanation cannot be discussed outside of the design of the evaluation process because when the explanations are not stress-tested in conditions of realistic variability of imaging and field noise, the claims of interpretability fail.

E. Data Modalities and Fusion Strategies

The dominance of RGB imagery in the literature means that it is cheap and available to many, but the dominance also focuses on research emphasis on leaf-level diagnostics, thus limiting the overall knowledge on system-level agricultural intelligence [64]. In mobile or farmer-facing applications, RGB pipelines are subject to further limitations of on-device inference, variation of capture, and user-controlled noise, which fundamentally change the concept of robustness in assessment. The multispectral UAV benchmarks allow shifting the emphasis of diagnosis to monitoring, allowing the stress to be detected, segmented, and field mapped on a large scale, much closer to precision agriculture as an operating field [70]. The context of federated learning is specifically relevant in these remote sensing tasks since the challenges of data governance and sovereignty are not in individual form, but typically as an institution, making cross-silo federation a logical organizational choice [91].

Hyperspectral and spectroscopic modalities present a unique interpretability difficulty because the explanation should be articulated in feature spaces that are not visually natural. This drives explainable AI beyond saliency overlays into domain-consistent explanations based on spectral and physiological explanations. Therefore, the taxonomy needs to distinguish between visual interpretability and scientific interpretability where the latter needs to be resistant to sensor noise and domain-specific feature correlation [93]. IoT and sensor-fusion pipelines also take even more agricultural intelligence than imaging and pose greater privacy risks because fine-grained temporal traces can provide insight into management choices and economic operations. Hierarchical and distributed federated architectures of sensor-abundant agricultural systems underscore that fusion is equally a modeling problem as it is an orchestration problem, with local smart preprocessing, aggregation rate, and client trustworthiness driving what the global model can ultimately come to know.

F. Deployment Contexts

Cloud-based deployments make orchestration and model updates easier; however, they tend to centralize control even with local raw data, establishing a dependency on central coordinators that might be inconsistent with agricultural trust

structures. The idea of cloud-federated pipelines for disease detection shows the attractiveness of the managed infrastructure and at the same time demonstrates how connection and synchronization quickly become first-order constraints. Alternative deployment, edge-centric, is often motivated by latency alone, yet in farming, it also serves as a form of governance: local inference keeps exposure minimized, and on-farm autonomy is encouraged. However, this change reduces the viable model space and increases the significance of efficiency-conscious architectures and steady explanation behavior in the face of degraded sensing. Hybrid edge-cloud federated pipelines strive to strike a balance between responsiveness and coordination, but cannot be considered algorithms, as their failure modes tend to be infrastructural. The presence of nearby fields such as precision aquaculture indicates that federated integration of edge-cloud platforms can be successful, although only in cases where communication policies, dynamics of participation, and schedule of updates have been structured with the same care and attention as the models themselves.

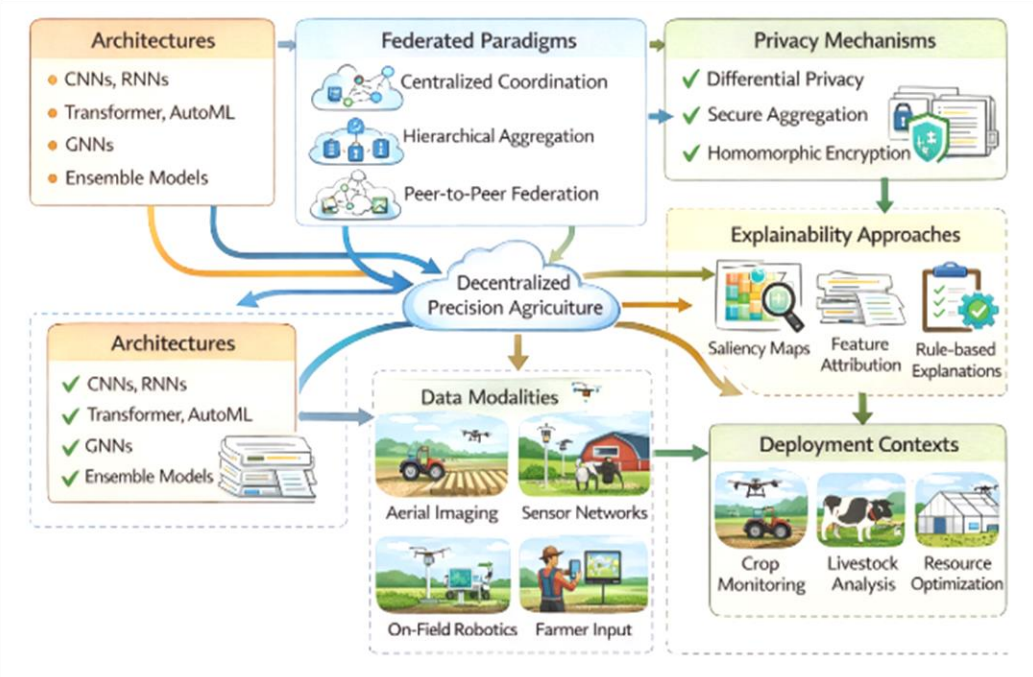

FIGURE 4. Integrated taxonomy linking architectures, federated paradigms, privacy mechanisms, explainability approaches, data modalities, and deployment contexts for decentralized precision agriculture.

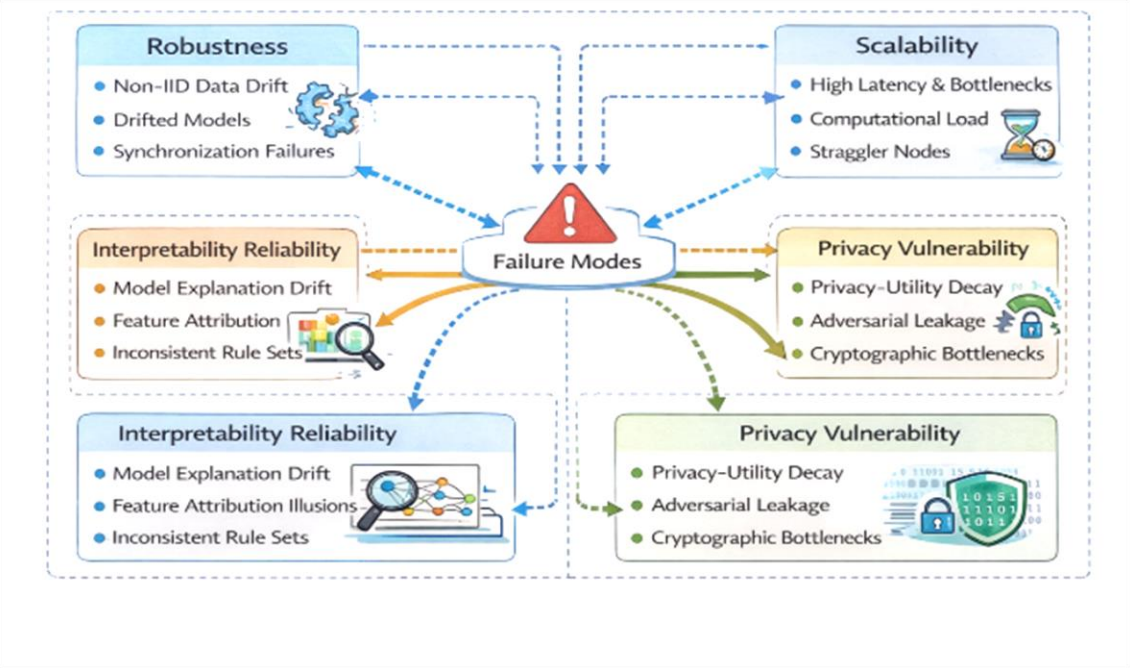

FIGURE 5. Failure-mode map illustrating dominant breakdown points across robustness, scalability, interpretability reliability, and privacy vulnerability.

Figure 4 is applied as the organizing backbone of the review, but not as an illustrative diagram. This demonstrates how decisions that seem to be made independently, such as the choice of a transformer backbone, the use of differential privacy or the inclusion of Grad-Cam are interrelated when the context of deployment is considered. The figure 4 also reveals the aspects on which the literature has under specified its assumptions, especially the coordinator of federation, assumed threat model, and reliability of explanation under non-IID clients. This allows the subsequent comparative synthesis to be performed without returning to a method-by-method catalogue.

Negative space is revealed in Figure 5: it sums up the failures of systems where headline accuracy appears to be good. The number points to common failure points aggregate non-IID instability, privacy-utility degenerates under noise, cryptographic latency, and explanation drift across clients, therefore the review can talk of failure as evidence, as opposed to failure as embarrassment. This value also prepares Chapter 5 in the sense that it establishes a neat jump over taxonomy to comparative trade-off synthesis.

V. Comparative Synthesis and Failure Analysis

This chapter is the analytical heart of the review, moving beyond the classification to a comparative analysis of the effectiveness of modern methods. It evaluates the points of success and graceful failure of these approaches and the failure points that undermine trust and deployability. Rather than ranking methods according to this accuracy or novelty, the synthesis measures the effect of design decisions on robustness, scalability, interpretability, reliability, and privacy vulnerability under conditions typical of decentralized agriculture.

A. Robustness Under Non-IID Conditions

The lack of robustness in agricultural federated learning is generally manifested at the beginning of the process, not as a full loss of accuracy but as instability in optimization. This is because heterogeneous clients introduce conflicting gradients, which in turn destabilize the initial stages of training [57]. Decentralized crop classification empirically indicates that models often stabilize to locally consistent but globally inconsistent representations, especially when the distributions of the clients are indicative of actual

agronomic heterogeneity, and not due to sampling noise [32]. Conversely, there are architectures that do not undergo catastrophic failure; rather, they have a gradual degeneration in performance, in which there is greater variability in the accuracy with clients, although the average accuracy is agreeable [67]. The situation becomes even more problematic when the disease prevalence is different across sites, and minority-pathology signals are either diluted during aggregation or overfitted at the local level without any meaningful transfer [69]. These findings indicate that robustness in agricultural contexts is more stable in learning dynamics in the presence of unending heterogeneity than in worst-case accuracy. In addition, the majority of the assessments of federated learning in agriculture are based on simulations, with only a small number of works being held in multi institutional or field settings, leaving a gap that will have a significant effect on how the claims on the strengths and scalability covered are interpreted.

B. Scalability and Communication Limits

Scalability claims tend to work in theory but fail to work in practice because of practical communication limitations with large-scale participation by clients, particularly when the size of the participating group is larger than that of a small cooperative cohort [83]. Aggregation approaches that are communication-efficient can reduce the amount of bandwidth used, but often make the assumption of trustworthy participation and synchronized updates, which is not a realistic view of rural connectivity conditions [27]. Moreover, cryptographic aggregation further worsens the issue of scalability by adding latency and coordination costs that grow super-linearly with the number of clients and with the dropout rate [26]. Empirical research suggests that algorithmic scalability may be unrelated to institutional scalability because coordination costs, fault tolerance, and negotiation of trust become dominant long before computational constraints are encountered [92]. The conceptualization of scalability in agriculture, therefore, should be that of a socio-technical, as opposed to an algorithmic property.

C. Deployment Realism

A different type of failure is demonstrated by deployment realism: operationally weak systems that are technically functional, but seldom seen by benchmark testing. The key limitation of edge-based agricultural systems is energy usage, as devices need to run day and night with small power constraints [81]. Feasible update frequency is also constrained by bandwidth, which means that model freshness and communication cost must be traded off, which is hardly ever considered in centralized evaluations [66]. Hardware heterogeneity further adds to these problems, because federated clients vary in data, as well as in computational power, and have straggler effects that bias aggregation timing and model evolution [87]. Field deployments suggest that small architectural inefficiencies will make otherwise good approaches impractical as the scale goes beyond pilot studies [21].

D. Interpretability Reliability

Failures of interpretability are more dangerous than failures of accuracy due to their subtlety, and it seems that one can offer an explanation that seems plausible but does not predict consistently across training runs or client sets. Explanation drift is a problem in that it is assumed that patterns of saliency or attribution of explanations can be used to analyze semantically similar inputs across federated regimes, resulting in inconsistent explanations [47]. The problem of cross-client inconsistencies occurs particularly in the agricultural field, where interested parties desire to have similar accounts on farms to allow them to understand agronomics together and make informed decisions regarding their policies [50]. Quantitative and qualitative evaluation studies both indicate that, with a shift in distribution, the explicatory reliability, such as the predictive performance, is worse off and it is possible that the apparent accuracy of a model can be maintained, but the reliability becomes worse and worse [51]. This imbalance demonstrates the importance of considering interpretability as a reliability constraint and not as an auxiliary visualization option.

E. Privacy Vulnerabilities

Potential privacy vulnerabilities occur in federated agricultural systems because they are accessed through a channel which is often underestimated when formal guarantees are considered adequate. Gradient-leakage attacks show that information can be leaked as model updates, even when the raw data are kept locally, especially when the agricultural data have a high spatial or temporal correlation [75]. Noise defenses reduce this risk but cause a privacy-utility crash because perturbations cause weak agronomic signals to be hidden, particularly in the case of rare diseases or early stress signals [76]. Explanation leakage is an under-researched but equally important weakness because visual or feature-based explanations may unintentionally transform site-specific properties, thus breaking confidentiality [85]. These technical risks are further increased by regulatory misalignment because different compliance requirements of transparency and data protection can push system design in competing directions when the requirements of privacy and explainability are optimized separately [79]. Combining these weaknesses, privacy in agriculture cannot be achieved by only selecting a mechanism, but must be jointly developed with learning and explanatory approaches [86].

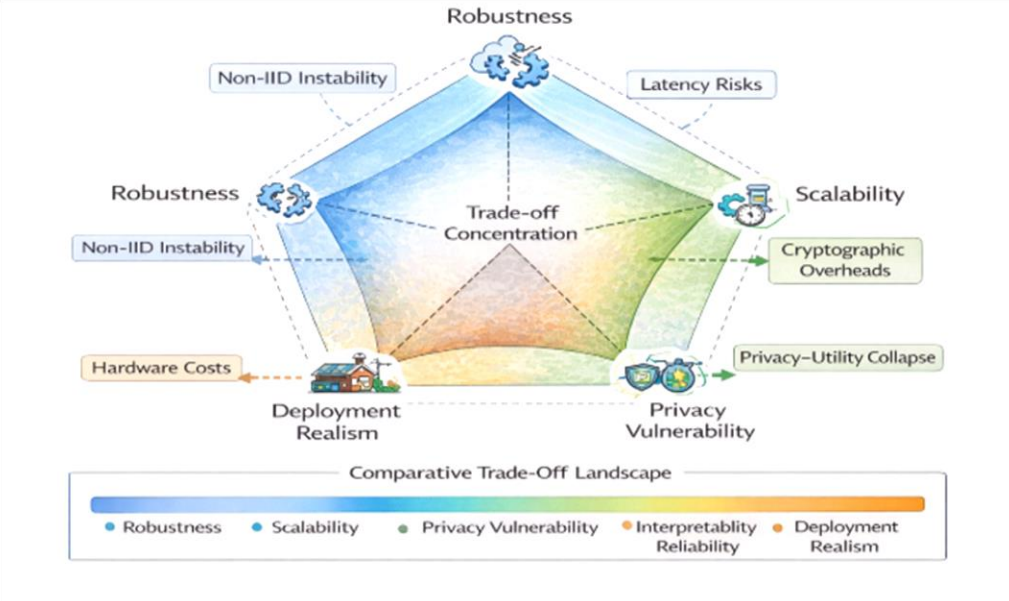

FIGURE 6. Comparative trade-off landscape illustrating interactions among robustness, scalability, deployment realism, interpretability reliability, and privacy vulnerability in decentralized agricultural AI systems.

Figure 6 integrates the multi-dimensional trade-offs revealed throughout this chapter into one analysis perspective. instead of ranking approaches, the figure 6 shows the locations of the various design options where risk is likely to be concentrated, which allows us to see why the advances along one axis are likely to push the failure towards another axis. The main point that is supported by this visualization is that there is no single technique that has been shown to be dominant in all of the dimensions, and that the AI design approach in agriculture needs to be based on compromise with constraints and not on maximal optimization.

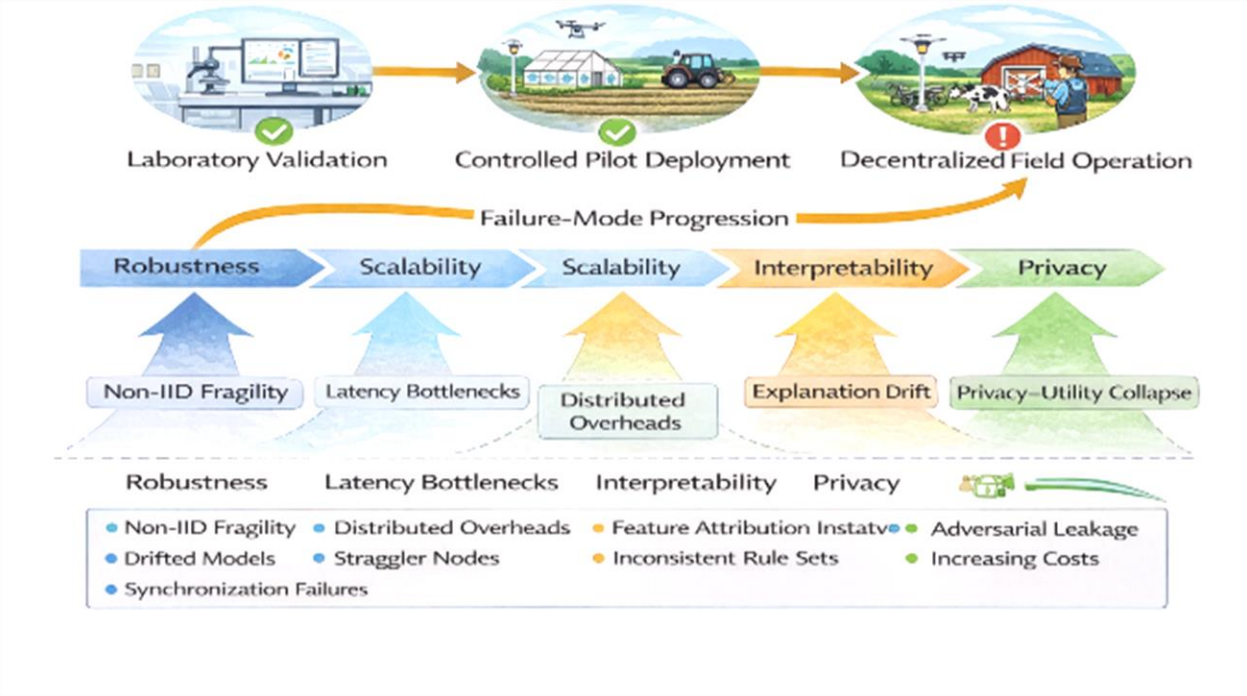

FIGURE 7. Failure-mode progression showing how robustness, scalability, interpretability, and privacy issues emerge as systems transition from controlled evaluation to real-world agricultural deployment.

Figure 7 points out that most failures are not simultaneous but manifest as systems move between the laboratory and decentralised field operation. Initial tests usually obscure deployment-critical problems like drift in explanations or communication cost, and only become apparent when it is operated over an extended period. The characterization of failure as a process as opposed to a final decision helps in legitimizing the importance of evaluation procedures that put stress systems in realistic agricultural scenarios before it is claimed that they are ready to be used.

VI. Biometric and Trend Analysis

This part is not meant to count the papers. Its aims are more direct and useful, namely, to clarify how the research agenda has been shaped, why some of the ideas have become popular and others stagnated, and what this tendency implies the validity of the existing statements. Because the corpus represents a combination of disciplines, including computer vision, privacy engineering, distributed optimization, and agricultural informatics, bibliometric surrogacy is used in this case as an indicator of maturity and coordination and not as an alternative to technical evaluation.

A. Temporal Evolution: Pre-2020 vs. Post-2020 Inflection

The history of agricultural deep learning before 2020 can perhaps be encapsulated as an exploration of competence within the framework of heavily unchallenged centralized paradigms. The research work in this period was mainly focused on showing how neural architectures could outperform hand-written pipelines when faced with visually challenging plant symptomatology, and the accuracy of such networks on existing datasets was the major measure of improvement [15]. Centralization was implicitly seen as the default engineering configuration and not a hypothesis to be justified and the issue of data governance, ownership or coordination were effectively externalized. Reflectively, the key aspect in these early pipelines is not the simplicity, but the rarity, with which they were designed to withstand the conditions that make agricultural AI intrinsically arduous, namely

seasonality, field heterogeneity, changing cultivars, and lumpy labeling economics [33].

After 2020, the research agenda is characterized by a strong extension as the concept of governance becomes part of the technical discussion. The concept of decentralization turns out to be a first-class design variable, and not a peripheral issue that can be deployed [54]. Privacy is no longer a rhetorical issue, but an explicit system constraint, with explicit mechanisms and threat models, and explainability becomes post-hoc visualization, operational accountability in a decision-support environment, where predictions are being taken. This inflection is not just the increase in the number of publications, but the change in the willingness of the community that high accuracy on curated benchmarks is insufficient to close the gap between the new environment and the real farming [45]. In addition, it understands the fact that meaningful deployment requires the treatment of learning systems as socio-technical artefacts and not solitary algorithms. A more muted though no less momentous secondary inflection comes along, with a growing pressure in regard to reliability. Modern literature goes beyond adding ingredients to federated learning such as privacy and explainability, and instead questions the possibility of these components existing together without compromising each other. The ideas of the stability of the explanation, verifiability, and the possibility of its deployment are not presented as optional enhancements; they are preconditions that reveal deep contradictions in designing the system [68].

B. Thematic Clusters

In this temporal development, the corpus naturally drift towards the formation of four thematic clusters, which have also been formed in parallel and have only recently started converging. What is important is not only the presence of these clusters but the reality that the assumptions behind them more often than not rarely go hand in hand, thus explaining why so many combined frameworks are convincing on paper and fractile in reality. The first group, dedicated to deep learning as a tool in the field of agricultural perception, has the largest presence in the initial timeline and is still significant since perception problems are relatively easy to define and assess. It's typical attributes are the model-centric improvement, dataset-centric assessment, and negligence of governance. Although the strongest contributions make representational capacity or image variability robustness, weaker contributions implicitly represent that access to data and labeling scales linearly- which seldom holds outside pilot deployments [74].

The second focuses on federated learning of distributed agricultural intelligence and enlarges the time when decentralization becomes inevitable. It focuses on training protocols, aggregation strategies and non-IID behavior but a common shortcoming is that the client heterogeneity is treated as a technical nuisance rather than as domain property with agronomic meaning. This framing provides conclusions that underestimate the magnitude and tenacity of real-life variance among farms [84]. The third group deals with privacy preserving and secure aggregation, providing conceptually clean solutions with high confidentiality or integrity assurances. Nonetheless, its implementation bottleneck frequently lies in coordination overhead, latency, and complexity of operation especially in agricultural alliances where infrastructure maturity and incentives are highly disparate. The fourth cluster centers on explainable AI in agricultural decision making as well as agricultural accountability and is conceptually broad and uneven in methods. Older contributions take interpretability as a reliability criterion, which assert that explanations are not meaningless in domain shift and also that explanations are auditable, whereas weaker work downgrades explainability to a visualization auxiliary that fails exactly when one needs uncertainty, disagreement, or high-stakes intervention. The recent merging of these clusters can be summed up in the form of a triangular constraint: federated learning is associated with the introduction of decentralization and heterogeneity, privacy is associated with perturbation and limited visibility, and explainability applies with the requirements of stable reasoning. These requirements are not natural to be harmonized and the literature more

clearly shows this uneasiness than attempts to cover it.

C. Geographic and Institutional Distribution

An undercover bibliometric approach may consider what areas are published most and a more informative approach would be to consider which institutions are structurally placed to do this sort of research and how that placement is used to construct a discursive of viability. The problem has institutional asymmetry. Organizations with well relationships, well developed computing infrastructure, and well developed governance structures find it much easier to struggle with research that requires multi-site data, federate deployment and privacy sensitive collaboration. Such asymmetry creates a material bias on the literature: systems-level federated research is more likely to be associated with contexts where multi-participant coordination is institutionally viable and smaller or less-resourced groups are more likely to make model-centric contributions that are measured on publicly available data.

This is the same kind of imbalance as the decentralization issue in agriculture. Similar to farms, research groups have unequal access to multi-farm data and deployment environments due to the unequal infrastructure of the farms. This leads to numerous assertions of decentralization being based on simulated federation and not sustained multi-stakeholder deployment, not due to carelessness in the methods, but due to institutional cost. This fact cannot be a secondary fact; it directly influences what the literature over- and under-represents. The corpus overemphasizes disease classification using imagery that can be easily retrieved, single-season testing and centralized or synthetic training, and underemphasizes multi-horizon robustness, multi-farm control, and responsibility in the face of conflicting stakeholder incentives. This disparity assists in the description of why some of the solutions seem mature on paper but are still uncommon to find in operating agricultural systems.

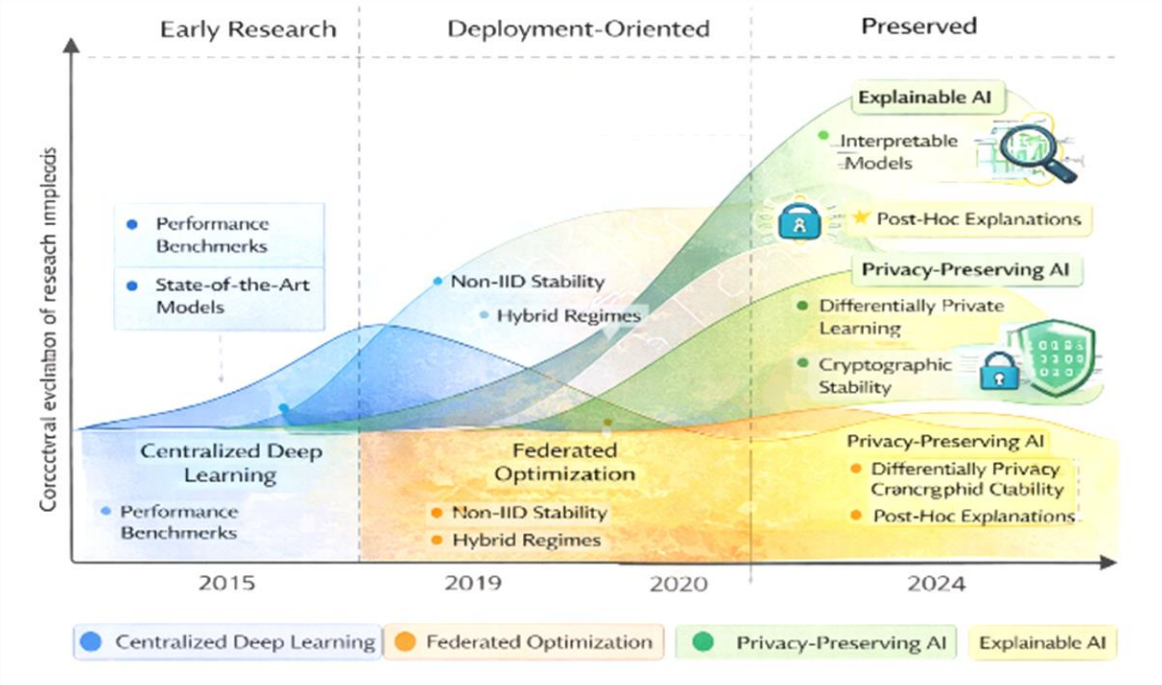

FIGURE 8. Conceptual temporal evolution showing shifts in emphasis from centralized deep learning towards federated, privacy-preserving, and explainable agricultural intelligence.

Figure 8 is provided to ward off a typical interpretive error: understanding the field as a linear advance of the more desirable models. The figure 8 shows discontinuities - especially the inflection of governance post-2020, as research motivation is no longer based on the need to demonstrate capability but on the need to deploy in a constrained manner. The next chapter of the roadmap can make shifts visible, thus making its argument based on direction, as opposed to opinion, and explaining which themes are emerging and which are simply enduring.

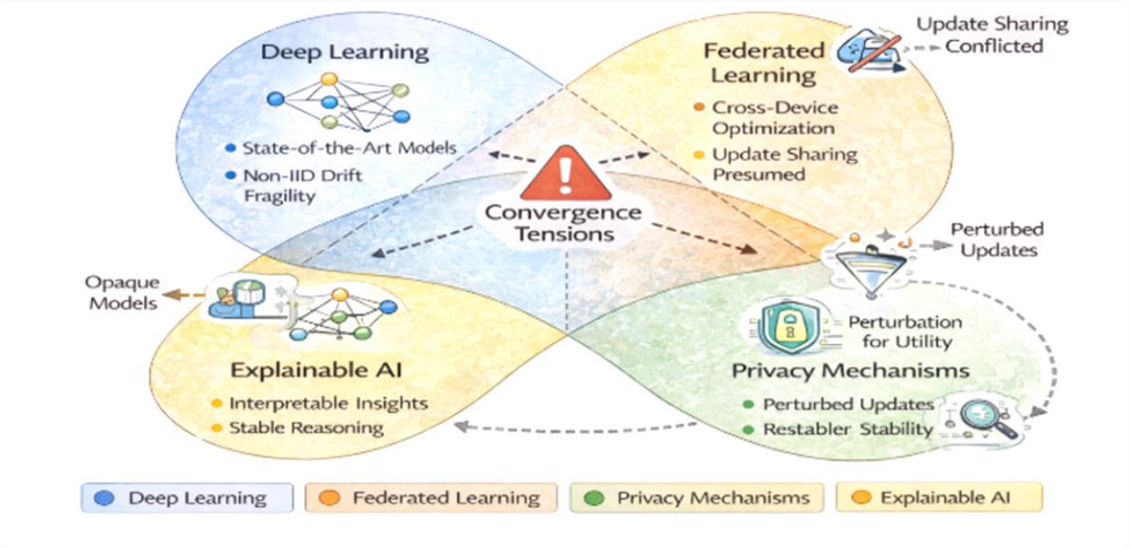

FIGURE 9. Cluster map of deep learning, federated learning, privacy mechanisms, and explainable AI, illustrating the convergence tensions that define decentralized precision agriculture.

Figure 9 has been added to depict the fact that integration is not additive. The figure 9 directly shows that the different clusters have assumptions which can potentially be in conflict with those of the other clusters: federated optimization assumes updates are shared, privacy constraints restrict visibility and inject perturbation, and explainability requires stable reasoning. This conceptual map provides a brief explanation of why many of the papers on the topic of variables presented as frameworks seem convincing and yet they do not outline how reliability in the explanation and privacy assurances can be simultaneously maintained within non-IID federated environment.

VII. Challenges and Open Research Gaps

The following section is not an attempt to count papers per se. Instead, its goal is more backward-moving and advantageous: to shed light on the development of the research agenda, why some ideas have become more popular and others have not, and how this trend would impact the validity of the existing assertions. As corpus cuts across a variety of disciplines, including computer vision, privacy engineering, distributed optimization, and agricultural informatics, bibliometric indicators are used here as measures of maturity and coordination, not instead of a technical evaluation.

A. Data Realism and Representational Fragility

One continuing weakness of the literature reviewed is the utilization of visually clean and insufficiently contextualized datasets that fail to capture the agronomic variability that is experienced in real farms. Although nominally agricultural datasets often exclude temporal dynamics, compound stressors and field scale noise, thus, models are biased to internalize stereotypes of symptoms instead of agronomic processes. In federated models, this weakness is compounded by the fact that data realism interacts

directly with non-IID distributions, causing local models to deviate in manners that never existed when realism is sacrificed to convenience in benchmarks [36]. The unresolved common gap is that there is no generally accepted protocol to determine that an agricultural dataset is deployment-representative and not only task-relevant and so the claim of robustness fails regardless of the learning paradigm used.

B. Domain Shift as a First-Order Constraint

The domain shift has been implicitly viewed as a secondary effect in most of the decentralized agricultural pipelines and can be countered by aggregation or personalization. In reality, though, processes of performance degradation are dominated by changes with cultivar variation, seasonal change, sensor heterogeneity, and management practices especially in cases where training is spread across farms. Explainable tools tend to conceal these changes, with attribution maps able to display a visual plausible encoding of radically different causal characteristics across fields. A gap thus remains open in the absence of mechanisms in existing federated learning and explainability systems to diagnose domain shift in the field and practitioners remain oblivious to the fact that deployed models have left their validity envelope without their knowledge.

C. Privacy–Utility Collapse in Federated Agriculture

Privacy-preserving mechanisms are often considered isolated, utility loss as an acceptable and minor trade-off, and not as a system failure mode [72]. Under real-world farming applications, privacy budgets are often bargained in an institutional or regulatory context, and this compels models into regimes where predictive accuracy and explanatory coherence are both compromised [77]. This breakdown is particularly severe when gradient level noise or cryptographic resources are engaged in connection with small, non-uniform client groups. However, there is no official standard of determining at which point privacy protection is no longer agronomically viable and decisions to deploy rely on ad-hoc decision-making as opposed to principled decision-making [80].

D. Explainability Inconsistency and Governance Risk

It is now understood in the literature that stability of explanation is as significant as the quality of explanation, and very little is done to systematically assess the consistency in clients, training rounds, or privacy settings. With the lack of consistency in decentralized agricultural situations, the loss of trust is not only at the individual user level, but also at the institutional level of advisory bodies, which should explain their recommendations at the level of regions and stakeholders. The instability of post-hoc explainability approaches is especially high when the training objective is not tied to them and consequently they adopt all the upstream distributional distortions. One defining open gap that is still critically unsolved is the lack of a governance conscious definition of what acceptable variance of explanation is in automated decision systems, though increasingly regulation demands transparency and accountability of such.

E. Absence of Auditable Benchmarks

The agricultural AI benchmarking practices are very fragmented, and the majority of the assessments are specific to a crop, sensor, or learning setup. Where they do exist, federated benchmarks often lack explicit privacy limitations and hardly have measures of explainability, making cross-study comparison practically impossible. This fragmentation stimulates local optimization in place of cumulative science, especially in review-based sub disciplines. The field thus has no lowest auditable benchmark collection that can collaboratively assess the accuracy, stability, privacy leakage, and explanation reliability in a realistic non-IID setting.

F. Sustainability and Long-Term Cost Structures

Computational sustainability is consequently regarded as an engineering consideration, but energy usage, communication waste, and hardware obsolescence have a direct effect on the sustainability of decentralized agricultural AI. Federated pipelines can minimize the raw data flow and at the same time, they raise the cost of coordination, switching and not eradicating the

environmental and economic costs. Solutions that are characterized as scalable might not be sustainable without explicit cost accounting when implemented in geographically dispersed and resource-constrained farming communities. One open gap is that there is no standardized model of assessing the life-cycle sustainability of privacy-preserving and explainable federated systems in agriculture, although they are often marketed as socially responsible technologies.

G. Synthesis

Collectively, these issues suggest that the enhancement of decentralized precision agriculture will not be realized by placing incremental improvements in architectural design. The existing gaps are structural, which cut across data realism, domain awareness, privacy governance, interpretability reliability, benchmarking discipline, and sustainability economics. To tackle them, they need a change in focus to develop more integrated evaluation protocols to discourage the idea that agriculture is the convenient application domain of AI, but rather a stress test of how much modern AI paradigms can actually operate when faced with fundamental limits.

VIII. Future Directions and Research Roadmap

The analysis above shows that the development of decentralized precision agriculture will rely less on ongoing improvements in accuracy and more on the re-conceptualization of the processes of learning, explanation and governance. This chapter does not give any predictions but suggests a research agenda, which gives a direction of the research that is technically grounded, institutionally plausible, and auditable over time [7].

A. Agricultural Foundation Models

General and large-scale models that are pre-trained on a wide range of agricultural cues provide a bright future in order to mitigate task-specific brittle behavior, yet enable local adaptation via federated learning. As opposed to generic vision or language foundation models, agricultural foundation models need to explicitly represent phenological cycles, biotic and abiotic stress interactions, and sensor-specific quirks, or large-scale pretraining will only increase bias instead of alleviating it. Even though consistent with agricultural realities, federated pretraining among farms, regions, or institutions has been understudied despite its natural ability to meet data-sovereignty considerations restricting the construction of centralized corpus. Going forward, it is desirable in the roadmap to consider federated protocols of pretraining that directly measure transfer stability across crops, seasons, and sensing conditions instead of just focusing on downstream fine-tuning accuracy.

B. Explainable Federated Optimization

The vast majority of explainability methods are post-hoc by nature, which inherently limits their capacity to affect the learning dynamics within federated systems. A more encouraging avenue is to directly incorporate explanation consistency in to the optimization process and consider interpretability reliability a formal constraint instead of a visualization artifact. This would allow training processes to punish explanation drift with a client, by directly addressing one of the failure modes that occur regularly in non-IID agricultural contexts. Its roadmap implication is that explanation-conscious loss functions should be formalized in the future, and their cost robustness under privacy restrictions ought to be assessed in order to establish the consistency of interpretability at the expense of unacceptable losses in predictive utility.

C. Edge-First Federated Architectures

Most agricultural applications are characterized by hard bandwidth, energy, and hardware limits that nullify cloud-based assumptions that are often embraced in the field of machine learning studies. In feature extraction, partial training, or explaining feature generation, edge-first federated architectures provide a means of decreasing communication overhead and maintaining responsiveness and on-farm autonomy. Moving computation to the edge does carry with it, though, moving complexity to the management and maintenance of devices and lifecycle sustainability-dimensions which are not commonly reflected in existing assessments. The

research in architecture has to then step out of the latency-based benchmarks to incorporate energy-conscious and maintenance-conscious metrics that will better represent the viability of the long term in rural and resource-restrained conditions.

D. Digital Twins for Federated Agriculture

Digital twins offer a guided process of emulating crop dynamics, agricultural intervention, and environmental variability without presenting uncooked farm data. Digital twins, when combined with federated learning, can also serve as a tool to validate models by being subjected to synthetic and agronomically reasonable scenarios in an intermediate state before going into the real world. Explainability becomes particularly important in this context, because twin-based simulations require explainable causal mechanisms instead of black box forecasting [94]. The implication of the roadmap is that future studies ought to seek federated digital twin ecosystems whereby models are being trained, audited, stress-tested against changing virtual counterparts as a condition to field-level adoption.

E. Governance-Aware AI Systems

To ensure adoption in agriculture, AI-driven decisions overlap regulation, liability, and farmer trust, and only technical robustness may ensure adoption. Governance-conscious AI re- conceptualizes privacy, explainability, and accountability as intrinsic properties of the system that need to be evidenced to non-practitioners of machine learning seeking to understand the system. Federated systems add to this landscape further by spreading responsibility between data owners, model aggregators, and technology providers, and thereby blurring conventional lines of accountability. Based on this, future work must outfits evaluation protocols with governance criteria, including auditability, traceability, and explanation accountability, as opposed to considering such criteria as secondary compliance inspections.

F. Integrated Research Trajectory

Combined, these recommendations imply the change of model-centric to pipeline-centric design where learning, explanation, privacy, and governance are evaluated together. Instead of leaps of faith, the roadmap focuses on controlled growth, which includes foundation models, explainable optimization, edge-first deployment, digital twins, and governance alignment. The continuation to this path will entail interdisciplinary cooperation and mutual standards that compensate consistency, transparency and sustainability and predictive performance.

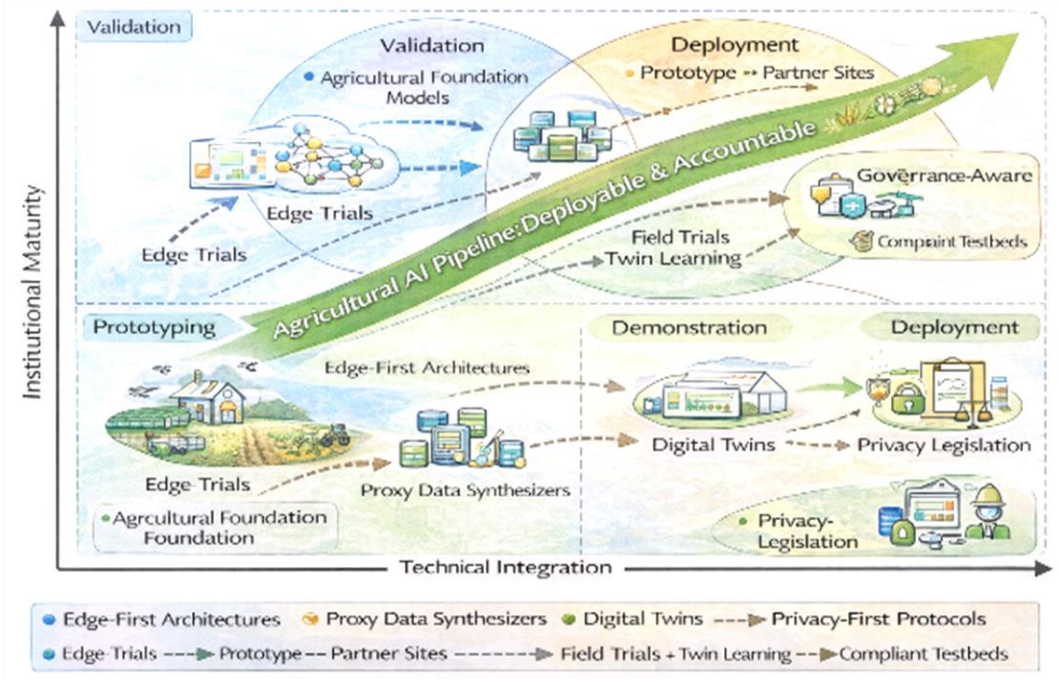

FIGURE 10. Research roadmap for decentralized precision agriculture, illustrating the interaction between federated learning paradigms, explainability integration, privacy governance, and deployment contexts over increasing system maturity.

The synthesis of the agenda of the chapter is presented in Figure 10, which places the future research directions in two axes: technical integration and institutional maturity. The figure 10 explains that agricultural foundation models, explainable federated optimization, edge-first architecture, digital twin, and governance-conscious AI are not isolated trends, but reinforcing elements of a feasible long-term pipeline. Having this roadmap included, the chapter does not read as a list of unrelated opportunities and rather as a logical progression towards implementable, responsible farming AI.

IX. Conclusion

This review critically evaluates the interaction between federated learning, privacy protection, and explainable artificial intelligence in the context of decentralized precision agriculture. Instead of considering these paradigms as independently mature solutions that can simply be pulled together, the analysis demonstrated how their intersection reveals structural tensions that are often concealed by task-level performance signals. Empirical evidence indicates that improvements in predictive accuracy, regardless of a variety of architectural designs, learning methods, and deployment conditions, are inadequate proxy measures of reliability, trust, or institutional viability in agriculture. A few salient points have become apparent. First, decentralization in agriculture goes beyond just design choice and is a structural necessity due to the existence of data ownership, heterogeneity, and governance realities. Second, although federated learning reduces the risks of centralization, it introduces a state of instability that manifests only when the conditions are non independent and identically distributed (non-IID) and under resource-constrained circumstances. Third, explainability cannot be an afterthought; it loses its credibility at the very point of federation, where the challenges are most severe. Lastly, privacy mechanisms redefine both optimization dynamics and explanatory behavior, and are typically not well described by formal guarantees alone. Of no lesser importance are the unresolved problems: there is still no common standard of dataset realism; there is no auditable standard that simultaneously judges privacy, robustness, and

interpretability; and no standard definition of what explanation stability is, to be used either in regulation or in advice. These missing pillars stifle not just technical progress, but also the ability of institutions to make righteous justifications for adoption. Therefore, convergence and non dominance were the main thesis. Individual paradigms, such as deep learning, federated optimization, privacy engineering, and explainable AI, cannot assert their lead in isolation without causing the failures in other parts of the pipeline. Decentralized precision agriculture will require further development that is based on designs clearly identifying trade-offs, prioritizing accountability over innovation, and aligning technical choices with agronomic, social, and governance limits. It is only on the convergence of such convergence that agricultural AI can be transformed into an experimental promise into a sustainable practice. These findings are based on a comparative synthesis as opposed to new empirical assessment, which is the existing evidentiary framework of the discipline and the purpose of the review to preempt systemic limitation as opposed to algorithmic delivery assertions.

REFERENCES

[1] R. Karthik, A. Ajay, A. Singh Bisht, T. Illakiya, and K. Suganthi, "A Deep Learning Approach for Crop Disease and Pest Classification Using Swin Transformer and Dual-Attention Multi-Scale Fusion Network," IEEE Access, vol. 12, pp. 152639–152655, 2024, doi: 10.1109/ACCESS.2024.3481675.

[2] Y. Borhani, J. Khoramdel, and E. Najafi, "A deep learning based approach for automated plant disease classification using vision transformer," Sci Rep, vol. 12, no. 1, p. 11554, July 2022, doi: 10.1038/s41598-022-15163-0.

[3] H. Chen, H.-Y. Hsu, J.-Y. Hsieh, and H.-E. Hung, "A differential privacy-preserving federated learning scheme with predictive maintenance of wind turbines based on deep learning for feature compression and anomaly detection with state assessment," J Mech Sci Technol, vol. 38, no. 7, pp. 3413–3429, July 2024, doi: 10.1007/s12206-024-0616-9.

[4] H. A. Tahir, W. Alayed, and W. U. Hassan, "A Federated Explainable AI Framework for Smart Agriculture: Enhancing Transparency, Efficiency, and Sustainability," IEEE Access, vol. 13, pp. 97567–97584, 2025, doi: 10.1109/ACCESS.2025.3571340.

[5] H. M. Albarakati et al., "A Novel Deep Learning Architecture for Agriculture Land Cover and Land Use Classification from Remote Sensing Images Based on Network-Level Fusion of Self-Attention Architecture," IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, vol. 17, pp. 6338–6353, 2024, doi: 10.1109/JSTARS.2024.3369950.

[6] S. Basudan, "A privacy-preserving federated learning protocol with a secure data aggregation for the Internet of Everything," Computer Communications, vol. 223, pp. 1–14, July 2024, doi: 10.1016/j.comcom.2024.05.005.

[7] K. R. Žalik and M. Žalik, "A Review of Federated Learning in Agriculture," Sensors, vol. 23, no. 23, p. 9566, Jan. 2023, doi: 10.3390/s23239566.

[8] X. Zhang, Y. Luo, and T. Li, "A Review of Research on Secure Aggregation for Federated Learning," Future Internet, vol. 17, no. 7, p. 308, July 2025, doi: 10.3390/fi17070308.

[9] C. Zhang, Y. Xie, H. Bai, B. Yu, W. Li, and Y. Gao, "A survey on federated learning," Knowledge-Based Systems, vol. 216, p. 106775, Mar. 2021, doi: 10.1016/j.knosys.2021.106775.

[10] V. Mothukuri, R. M. Parizi, S. Pouriyeh, Y. Huang, A. Dehghantanha, and G. Srivastava, "A survey on security and privacy of federated learning," Future Generation Computer Systems, vol. 115, pp. 619–640, Feb. 2021, doi: 10.1016/j.future.2020.10.007.

[11] R. J. Garro, C. S. Wilson, D. L. Swain, A. J. Pordomingo, and S. Wibowo, "A systematic literature review on the applications of federated learning and enabling technologies for livestock management," Computers and Electronics in Agriculture, vol. 234, p. 110180, July 2025, doi: 10.1016/j.compag.2025.110180.

[12] Md. T. Ahmed, M. W. Ahmed, and M. Kamruzzaman, "A systematic review of explainable artificial intelligence for spectroscopic agricultural quality assessment," Computers and Electronics in Agriculture, vol. 235, p. 110354, Aug. 2025, doi: 10.1016/j.compag.2025.110354.

[13] F. Galli, K. Jung, S. Biswas, C. Palamidessi, and T. Cucinotta, "Advancing Personalized Federated Learning: Group Privacy, Fairness, and Beyond," SN COMPUT. SCI., vol. 4, no. 6, p. 831, Oct. 2023, doi: 10.1007/s42979-023-02292-0.

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

[15] M. K. Saini, N. Goel, M. Miguez, and D. Singh, Agricultural-Centric Computation: Second International Conference, ICA 2024, Delhi, India, May 21–24, 2024, Revised Selected Papers. Springer Nature, 2025.

[16] S. Vats, M. S. Khanna, V. Kukreja, and S. Mehta, "AI in Agriculture: A Federated Learning CNN Approach to Detecting Almond Leaf Disease," in 2024 IEEE International Conference on Interdisciplinary Approaches in Technology and Management for Social Innovation (IATMSI), Mar. 2024, pp. 1–6. doi: 10.1109/IATMSI60426.2024.10503271.

[17] A. Alsuwaidi, M. A. Talib, Q. Nasir, F. Lamghari, S. Zerisenay, and F. M. Tsombou, "AI-Driven Early Detection of Crop Diseases Using Computer Vision and Federated Learning," in 2025 IEEE 4th International Conference on Computing and Machine Intelligence (ICMI), Apr. 2025, pp. 1–5. doi: 10.1109/ICMI65310.2025.11141296.

[18] Q. Gao, Y. Sun, X. Chen, F. Yang, and Y. Wang, "An Efficient Multi-Party Secure Aggregation Method Based on Multi-Homomorphic Attributes," Electronics, vol. 13, no. 4, p. 671, Jan. 2024, doi: 10.3390/electronics13040671.

[19] Y. Alqahtani, M. Nawaz, T. Nazir, A. Javed, F. Jeribi, and A. Tahir, "An improved deep learning approach for localization and recognition of plant leaf diseases," Expert Systems with Applications, vol. 230, p. 120717, Nov. 2023, doi: 10.1016/j.eswa.2023.120717.

[20] K. Hu, S. Gong, Q. Zhang, C. Seng, M. Xia, and S. Jiang, "An overview of implementing security and privacy in federated learning," Artif Intell Rev, vol. 57, no. 8, p. 204, July 2024, doi: 10.1007/s10462-024-10846-8.

[21] I. Siniosoglou et al., "Applying Federated Learning on Decentralized Smart Farming: A Case Study," in 2023 IEEE International Conference on Communications Workshops (ICC Workshops), May 2023, pp. 1295–1300. doi: 10.1109/ICCWorkshops57953.2023.10283681.

[22] C. G. Simhadri and H. K. Kondaveeti, "Automatic Recognition of Rice Leaf Diseases Using Transfer Learning," Agronomy, vol. 13, no. 4, p. 961, Apr. 2023, doi: 10.3390/agronomy13040961.

[23] S. K. Singh, M. Kumar, A. Khanna, and B. Virdee, "Blockchain and FL-Based Secure Architecture for Enhanced External Intrusion Detection in Smart Farming," IEEE Internet of Things Journal, vol. 12, no. 3, pp. 3297–3304, Feb. 2025, doi: 10.1109/JIOT.2024.3478820.

[24] M. Bhandari, T. B. Shahi, A. Neupane, and K. B. Walsh, "BotanicX-AI: Identification of Tomato Leaf Diseases Using an Explanation-Driven Deep-Learning Model," Journal of Imaging, vol. 9, no. 2, p. 53, Feb. 2023, doi: 10.3390/jimaging9020053.

[25] P. V. Caminha and H. M. N. da S. Oliveira, "Cloud-Driven Federated Learning for Plant Disease Detection in Agriculture," in 2024 IEEE 13th International Conference on Cloud Networking (CloudNet), Nov. 2024, pp. 1–4. doi: 10.1109/CloudNet62863.2024.10815857.

[26] J. Kim, G. Park, M. Kim, and S. Park, "Cluster-Based Secure Aggregation for Federated Learning," Electronics, vol. 12, no. 4, p. 870, Jan. 2023, doi: 10.3390/electronics12040870.

[27] K. Peng, X. Shen, L. Gao, B. Wang, and Y. Lu, "Communication-Efficient and Privacy-Preserving Verifiable Aggregation for Federated Learning," Entropy, vol. 25, no. 8, p. 1125, Aug. 2023, doi: 10.3390/e25081125.

[28] M. Ouhami, A. Hafiane, Y. Es-Saady, M. El Hajji, and R. Canals, "Computer Vision, IoT and Data Fusion for Crop Disease Detection Using Machine Learning: A Survey and Ongoing Research," Remote Sensing, vol. 13, no. 13, p. 2486, Jan. 2021, doi: 10.3390/rs13132486.

[29] Y. Wang, Y. Chen, and D. Wang, "Convolution Network Enlightened Transformer for Regional Crop Disease Classification," Electronics, vol. 11, no. 19, p. 3174, Jan. 2022, doi: 10.3390/electronics11193174.

[30] B. Tugrul, E. Elfatimi, and R. Eryigit, "Convolutional Neural Networks in Detection of Plant Leaf Diseases: A Review," Agriculture, vol. 12, no. 8, p. 1192, Aug. 2022, doi: 10.3390/agriculture12081192.

[31] S. P. Singh, K. Pritamdas, K. J. Devi, and S. D. Devi, "Custom Convolutional Neural Network for Detection and Classification of Rice Plant Diseases," Procedia Computer Science, vol. 218, pp. 2026–2040, Jan. 2023, doi: 10.1016/j.procs.2023.01.179.

[32] D. Mamba Kabala, A. Hafiane, L. Bobelin, and R. Canals, "Decentralized federated learning using validation loss for model sharing in crop disease classification," Ecological Informatics, vol. 90, p. 103205, Dec. 2025, doi: 10.1016/j.ecoinf.2025.103205.

[33] J. You, X. Li, M. Low, D. Lobell, and S. Ermon, "Deep Gaussian Process for Crop Yield Prediction Based on Remote Sensing Data," Proceedings of the AAAI Conference on Artificial Intelligence, vol. 31, no. 1, Feb. 2017, doi: 10.1609/aaai.v31i1.11172.

[34] K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," Computers and Electronics in Agriculture, vol. 145, pp. 311–318, Feb. 2018, doi: 10.1016/j.compag.2018.01.009.

[35] A. Haridasan, J. Thomas, and E. D. Raj, "Deep learning system for paddy plant disease detection and classification," Environ Monit Assess, vol. 195, no. 1, p. 120, Nov. 2022, doi: 10.1007/s10661-022-10656-x.

[36] W. Shafik, A. Tufail, L. C. De Silva, R. A. A. Haji Mohd Apong, and K.-H. Kim, "Deep learning technique for plant disease classification and pest detection and model explainability elevating agricultural sustainability," BMC Plant Biol, vol. 25, no. 1, p. 1491, Nov. 2025, doi: 10.1186/s12870-025-07377-x.

[37] G. Batchuluun, S. H. Nam, and K. R. Park, "Deep learning-based plant classification and crop disease classification by thermal camera," Journal of King Saud University - Computer and Information Sciences, vol. 34, no. 10, Part B, pp. 10474–10486, Nov. 2022, doi: 10.1016/j.jksuci.2022.11.003.

[38] R. H. Hridoy, Md. Tarek Habib, Md. Sadekur Rahman, and M. S. Uddin, "Deep Neural Networks-Based Recognition of Betel Plant Diseases by Leaf Image Classification," in Evolutionary Computing and Mobile
Sustainable Networks, V. Suma, X. Fernando, K.-L. Du, and H. Wang, Eds., Singapore: Springer, 2022, pp. 227–241. doi: 10.1007/978-981-16-9605-3_16.

[39] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs," IEEE
Transactions on Pattern Analysis and Machine Intelligence, vol. 40, no. 4, pp. 834–848, Apr. 2018, doi: 10.1109/TPAMI.2017.2699184.

[40] K. Khairnar and R. Dagade, "Disease Detection and Diagnosis on Plant using Image Processing A Review," IJCA, vol. 108, no. 13, pp. 36–38, Dec. 2014, doi: 10.5120/18973-0445.

[41] S. Mehta, V. Kukreja, and S. Vats, "Empowering Farmers with AI: Federated Learning of CNNs for Wheat Diseases Multi-Classification," in 2023 4th International Conference for Emerging Technology (INCET), May 2023, pp. 1–6. doi: 10.1109/INCET57972.2023.10170091.

[42] S. Mehta, V. Kukreja, and R. Gupta, "Empowering Precision Agriculture: Detecting Apple Leaf Diseases and Severity Levels with Federated Learning CNN," in 2023 3rd International Conference on Intelligent Technologies (CONIT), June 2023, pp. 1–6. doi: 10.1109/CONIT59222.2023.10205784.

[43] M. Y. Shams, S. A. Gamel, and F. M. Talaat, "Enhancing crop recommendation systems with explainable artificial intelligence: a study on agricultural decision-making," Neural Comput & Applic, vol. 36, no. 11, pp. 5695–5714, Apr. 2024, doi: 10.1007/s00521-023-09391-2.

[44] P. Chitra, P. Raghuraman, K. S. Varsha, and M. Mallavaram, "Enhancing Precision Agriculture through Stacked Ensemble Model and Interpretability," in 2025 1st International Conference on Secure IoT, Assured and Trusted Computing (SATC), Feb. 2025, pp. 1–8. doi: 10.1109/SATC65530.2025.11137185.

[45] M. M. Hasan et al., "Enhancing Rice Crop Management: Disease Classification Using Convolutional Neural Networks and Mobile Application Integration," Agriculture, vol. 13, no. 8, p. 1549, Aug. 2023, doi: 10.3390/agriculture13081549.

[46] A. Kumar, A. S. Gill, A. Sharma, and V. Kumar, "Enhancing the Explainability and Interpretability of Crop Yield Prediction Models Through Precision Agriculture," in Agricultural-Centric Computation, M. K. Saini, N. Goel, M. Miguez, and D. Singh, Eds., Cham: Springer Nature Switzerland, 2025, pp. 218–226. doi: 10.1007/978-3-031-74440-2_20.

[47] H. K. Kondaveeti and C. G. Simhadri, "Evaluation of deep learning models using explainable AI with qualitative and quantitative analysis for rice leaf disease detection," Sci Rep, vol. 15, no. 1, p. 31850, Aug. 2025, doi: 10.1038/s41598-025-14306-3.

[48] V. Jindal, V. Kukreja, S. Mehta, R. Yadav, and N. Mohd, "Evolving Agritech: Implementing Federated Learning & CNN for Parsley Leaf Disease Detection," in 2023 3rd Asian Conference on Innovation in Technology (ASIANCON), Aug. 2023, pp. 1–6. doi: 10.1109/ASIANCON58793.2023.10270228.

[49] S. Mostafa, D. Mondal, K. Panjvani, L. Kochian, and I. Stavness, "Explainable deep learning in plant phenotyping," Front. Artif. Intell., vol. 6, Sept. 2023, doi: 10.3389/frai.2023.1203546.

[50] L.-D. Quach, K. N. Quoc, A. N. Quynh, N. Thai-Nghe, and T. G. Nguyen, "Explainable Deep Learning Models With Gradient-Weighted Class Activation Mapping for Smart Agriculture," IEEE Access, vol. 11, pp. 83752–83762, 2023, doi: 10.1109/ACCESS.2023.3296792.

[51] K. Wei et al., "Explainable Deep Learning Study for Leaf Disease Classification," Agronomy, vol. 12, no. 5, p. 1035, May 2022, doi: 10.3390/agronomy12051035.

[52] B. Ben Youssef, L. Alhmidi, Y. Bazi, and M. Zuair, "Federated Learning Approach for Remote Sensing Scene Classification," Remote Sensing, vol. 16, no. 12, p. 2194, Jan. 2024, doi: 10.3390/rs16122194.

[53] V. Hiremani et al., "Federated learning for crop yield prediction: A comprehensive review of techniques and applications," MethodsX, vol. 14, p. 103408, June 2025, doi: 10.1016/j.mex.2025.103408.

[54] S. Moreno-Álvarez, M. E. Paoletti, A. J. Sanchez-Fernandez, J. A. Rico-Gallego, L. Han, and J. M. Haut, "Federated learning meets remote sensing," Expert Systems with Applications, vol. 255, p. 124583, Dec. 2024, doi: 10.1016/j.eswa.2024.124583.

[55] F. S. Khan et al., "Federated learning-based UAVs for the diagnosis of Plant Diseases," in 2022 International Conference on Engineering and Emerging Technologies (ICEET), Oct. 2022, pp. 1–6. doi: 10.1109/ICEET56468.2022.10007133.

[56] M. Aledhari, R. Razzak, R. M. Parizi, and F. Saeed, "Federated Learning: A Survey on Enabling Technologies, Protocols, and Applications," IEEE Access, vol. 8, pp. 140699–140725, 2020, doi: 10.1109/ACCESS.2020.3013541.

[57] G. Idoje, T. Dagiuklas, and M. Iqbal, "Federated Learning: Crop classification in a smart farm decentralised network," Smart Agricultural Technology, vol. 5, p. 100277, Oct. 2023, doi: 10.1016/j.atech.2023.100277.

[58] M. Aldossary, J. Almutairi, and I. Alzamil, "Federated LeViT-ResUNet for Scalable and Privacy-Preserving Agricultural Monitoring Using Drone and Internet of Things Data," Agronomy, vol. 15, no. 4, p. 928, Apr. 2025, doi: 10.3390/agronomy15040928.

[59] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," in 2017 IEEE International Conference on Computer Vision (ICCV), Oct. 2017, pp. 618–626. doi: 10.1109/ICCV.2017.74.

[60] R. Chen, H. Qi, Y. Liang, and M. Yang, "Identification of plant leaf diseases by deep learning based on channel attention and channel pruning," Front. Plant Sci., vol. 13, Nov. 2022, doi: 10.3389/fpls.2022.1023515.

[61] G. G. and A. P. J., "Identification of plant leaf diseases using a nine-layer deep convolutional neural network," Computers & Electrical Engineering, vol. 76, pp. 323–338, June 2019, doi: 10.1016/j.compeleceng.2019.04.011.

[62] D. Mamba Kabala, A. Hafiane, L. Bobelin, and R. Canals, "Image-based crop disease detection with federated learning," Sci Rep, vol. 13, no. 1, p. 19220, Nov. 2023, doi: 10.1038/s41598-023-46218-5.

[63] J. G. A. Barbedo, "Impact of dataset size and variety on the effectiveness of deep learning and transfer learning for plant disease classification," Computers and Electronics in Agriculture, vol. 153, pp. 46–53, Oct. 2018, doi: 10.1016/j.compag.2018.08.013.

[64] A. Bilal, X. Liu, H. Long, M. Shafiq, and M. Waqar, "Increasing Crop Quality and Yield with a Machine Learning-Based Crop Monitoring System," CMC, vol. 76, no. 2, pp. 2401–2426, 2023, doi: 10.32604/cmc.2023.037857.

[65] H.-Y. Chen, K. Sharma, C. Sharma, and S. Sharma, "Integrating explainable artificial intelligence and blockchain to smart agriculture: Research prospects for decision making and improved security," Smart Agricultural Technology, vol. 6, p. 100350, Dec. 2023, doi: 10.1016/j.atech.2023.100350.

[66] W. Khuen Cheng, J. Cheng Khor, W. Zheng Liew, K. Thye Bea, and Y.-L. Chen, "Integration of Federated Learning and Edge-Cloud Platform for Precision Aquaculture," IEEE Access, vol. 12, pp. 124974–124989, 2024, doi: 10.1109/ACCESS.2024.3454057.

[67] M. Aggarwal, V. Khullar, N. Goyal, A. Alammari, M. A. Albahar, and A. Singh, "Lightweight Federated Learning for Rice Leaf Disease Classification Using Non Independent and Identically Distributed Images," Sustainability, vol. 15, no. 16, p. 12149, Jan. 2023, doi: 10.3390/su151612149.

[68] M. K. Uppalapati, R. Vaddi, and R. Kata, "Machine Learning and Deep Learning Architectures for Enhanced Crop Yield Prediction and Prescriptive Precision Agriculture," International Journal of Science and Technology, vol. 1, no. 1, pp. 15–20, Dec. 2025.

[69] F. Deng, W. Mao, Z. Zeng, H. Zeng, and B. Wei, "Multiple Diseases and Pests Detection Based on Federated Learning and Improved Faster R-CNN," IEEE Transactions on Instrumentation and Measurement, vol. 71, pp. 1–11, 2022, doi: 10.1109/TIM.2022.3201937.

[70] Q. Li, H. Yuan, T. Fu, Z. Yu, B. Zheng, and S. Chen, "Multispectral Semantic Segmentation for UAVs: A Benchmark Dataset and Baseline," IEEE Transactions on Geoscience and Remote Sensing, vol. 62, pp. 1–17, 2024, doi: 10.1109/TGRS.2024.3457674.

[71] R. N. V. J. Mohan, P. S. Rayanoothala, and R. P. Sree, "Next-gen agriculture: integrating AI and XAI for precision crop yield predictions," Front. Plant Sci., vol. 15, Jan. 2025, doi: 10.3389/fpls.2024.1451607.

[72] P. Kumar, G. P. Gupta, and R. Tripathi, "PEFL: Deep Privacy-Encoding-Based Federated Learning Framework for Smart Agriculture," IEEE Micro, vol. 42, no. 1, pp. 33–40, Jan. 2022, doi: 10.1109/MM.2021.3112476.

[73] M. Majdalawieh, C. Martins, M. Radi, M. Alaraj, and S. Khan, "Precision agriculture in the age of AI: A systematic review of machine learning methods for crop disease detection," Smart Agricultural Technology, vol. 12, p. 101491, Dec. 2025, doi: 10.1016/j.atech.2025.101491.

[74] Md. A. R. Nishad, M. A. Mitu, and N. Jahan, "Predicting and Classifying Potato Leaf Disease using K-means Segmentation Techniques and Deep Learning Networks," Procedia Computer Science, vol. 212, pp. 220–229, Jan. 2022, doi: 10.1016/j.procs.2022.11.006.

[75] L. Lyu et al., "Privacy and Robustness in Federated Learning: Attacks and Defenses," IEEE Transactions on Neural Networks and Learning Systems, vol. 35, no. 7, pp. 8726–8746, July 2024, doi: 10.1109/TNNLS.2022.3216981.

[76] L. Zhu and X. Chen, "Privacy protection in federated learning: a study on the combined strategy of local and global differential privacy," J Supercomput, vol. 81, no. 1, p. 326, Dec. 2024, doi: 10.1007/s11227-024-06845-9.

[77] S. A. Mahmud, N. Islam, Z. Islam, Z. Rahman, and S. T. Mehedi, "Privacy-Preserving Federated Learning-Based Intrusion Detection Technique for Cyber-Physical Systems," Mathematics, vol. 12, no. 20, p. 3194, Jan. 2024, doi: 10.3390/math12203194.

[78] A. Bergstrom et al., "Protecting farm privacy while researching large-scale unmanned aircraft systems platforms for agricultural applications," Agronomy Journal, vol. 114, no. 5, pp. 2700–2714, 2022, doi: 10.1002/agj2.21054.

[79] J. Kaur, S. M. Hazrati Fard, M. Amiri-Zarandi, and R. Dara, "Protecting farmers' data privacy and confidentiality: Recommendations and considerations," Front. Sustain. Food Syst., vol. 6, Oct. 2022, doi: 10.3389/fsufs.2022.903230.

[80] T. Markovic, M. Leon, D. Buffoni, and S. Punnekkat, "Random forest with differential privacy in federated learning framework for network attack detection and classification," Appl Intell, vol. 54, no. 17, pp. 8132–8153, Sept. 2024, doi: 10.1007/s10489-024-05589-6.

[81] A. Milioto, P. Lottes, and C. Stachniss, "Real-Time Semantic Segmentation of Crop and Weed for Precision Agriculture Robots Leveraging Background Knowledge in CNNs," in 2018 IEEE International Conference on Robotics and Automation (ICRA), May 2018, pp. 2229–2235. doi: 10.1109/ICRA.2018.8460962.

[82] T. Kattenborn, J. Leitloff, F. Schiefer, and S. Hinz, "Review on Convolutional Neural Networks (CNN) in vegetation remote sensing," ISPRS Journal of Photogrammetry and Remote Sensing, vol. 173, pp. 24–49, Mar. 2021, doi: 10.1016/j.isprsjprs.2020.12.010.

[83] M. Moshawrab, M. Adda, A. Bouzouane, H. Ibrahim, and A. Raad, "Reviewing Federated Learning Aggregation Algorithms; Strategies, Contributions, Limitations and Future Perspectives," Electronics, vol. 12, no. 10, p. 2287, Jan. 2023, doi: 10.3390/electronics12102287.

[84] S. Mehta, V. Kukreja, and A. Gupta, "Revolutionizing Maize Disease Management with Federated Learning CNNs: A Decentralized and Privacy-Sensitive Approach," in 2023 4th International Conference for Emerging Technology (INCET), May 2023, pp. 1–6. doi: 10.1109/INCET57972.2023.10170499.

[85] S. Natarajan, P. Chakrabarti, and M. Margala, "Robust diagnosis and meta visualizations of plant diseases through deep neural architecture with explainable AI," Sci Rep, vol. 14, no. 1, p. 13695, June 2024, doi: 10.1038/s41598-024-64601-8.

[86] Z. Li et al., "RR-LADP: A Privacy-Enhanced Federated Learning Scheme for Internet of Everything," IEEE Consumer Electronics Magazine, vol. 10, no. 5, pp. 93–101, Sept. 2021, doi: 10.1109/MCE.2021.3059958.

[87] H. Devaraj et al., "RuralAI in Tomato Farming: Integrated Sensor System, Distributed Computing, and Hierarchical Federated Learning for Crop Health Monitoring," IEEE Sensors Letters, vol. 8, no. 5, pp. 1–4, May 2024, doi: 10.1109/LSENS.2024.3384935.

[88] Y. Li, J. Lai, R. Zhang, and M. Sun, "Secure and efficient multi-key aggregation for federated learning," Information Sciences, vol. 654, p. 119830, Jan. 2024, doi: 10.1016/j.ins.2023.119830.

[89] J. Shen, Y. Zhao, S. Huang, and Y. Ren, "Secure and Flexible Privacy-Preserving Federated Learning Based on Multi-Key Fully Homomorphic Encryption," Electronics, vol. 13, no. 22, p. 4478, Jan. 2024, doi: 10.3390/electronics13224478.

[90] V. Badrinarayanan, A. Kendall, and R. Cipolla, "SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 39, no. 12, pp. 2481–2495, Dec. 2017, doi: 10.1109/TPAMI.2016.2644615.

[91] A. Rajini, P. Parthasarathy, G. Shyam, B. Harshit kumar, and S. Shawn, "Soil Classification Using Deep Learning Techniques," in 2024 2nd International Conference on Recent Trends in Microelectronics, Automation, Computing and Communications Systems (ICMACC), Dec. 2024, pp. 400–404. doi: 10.1109/ICMACC62921.2024.10894675.

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

[93] A. K. Rangarajan, R. Purushothaman, and A. Ramesh, "Tomato crop disease classification using pre-trained deep learning algorithm," Procedia Computer Science, vol. 133, pp. 1040–1047, Jan. 2018, doi: 10.1016/j.procs.2018.07.070.

[94] P. K. Gupta, B. D. Mazumdar, S. N. Pillai, and R. S. Komaragiri, "Towards the Development of eXplainable Digital Twins for Precision Agriculture," in 2024 First International Conference on Pioneering Developments in Computer Science & Digital Technologies (IC2SDT), Aug. 2024, pp. 64–69. doi: 10.1109/IC2SDT62152.2024.10696477.

[95] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds., Cham: Springer International Publishing, 2015, pp. 234–241. doi: 10.1007/978-3-319-24574-4_28.