

ORIGINAL RESEARCH

Open access

Privacy-Preserving Precision Agriculture: A Federated Learning Architecture for Cross-Farm Model Development Without Centralized Raw Data Sharing

Charlotte Dupuis^{1*}, Hugo Perrin¹, Elise Morel², Thomas Martin³

Abstract

Precision agriculture depends on distributed data from sensors, remote sensing, machinery, crop and livestock monitoring, weather services, and farm-management systems. Collaborative model development can use these resources, yet central collection may intensify privacy, ownership, access, control, cybersecurity, participation, and secondary-use concerns. This article proposes a conceptual federated learning architecture in which farms retain raw data locally while contributing model updates to coordinated cross-farm learning. The architecture integrates use-case definition, stakeholder and governance mapping, farm-data sensitivity assessment, provenance and quality review, local node configuration, local model training, protected update exchange, secure aggregation, global and local evaluation, cross-farm validation, heterogeneity assessment, privacy safeguards, security controls, drift monitoring, auditability, and farmer feedback. It distinguishes horizontal federated learning from conditional vertical and transfer configurations and connects these principles to bounded agricultural use cases, including crop monitoring, disease detection, yield prediction, remote-sensing interpretation, livestock monitoring, and human-reviewed decision support. Secure aggregation, differential privacy, homomorphic encryption, and trusted execution environments are treated as context-dependent safeguards rather than guarantees. The contribution links decentralized model development with governance, accountability, communication constraints, non-identically distributed data, validation, monitoring, and claim boundaries. It supports research and design but does not establish certification, legal compliance, superiority, farmer approval, operational safety, or deployment readiness.

Keywords Precision agriculture, Farm data governance, Federated learning, Privacy-preserving AI, Cross-farm learning, Secure aggregation

*Correspondence:

Charlotte Dupuis
charlotte.dupuis@univ-lille.fr

¹ Department of Privacy-Preserving AI and Federated Learning, Faculty of Agriculture, University of Lille, Lille, France

² Department of Cross-Farm Model Development, Faculty of Agriculture, University of Montpellier, Montpellier, France

³ Department of Agricultural Data Privacy and Security, Faculty of Agriculture, University of Nantes, Nantes, France

Introduction

Precision agriculture combines sensing, imaging, machinery telemetry, management records, and computational models distributed across farms, devices, cooperatives, researchers, providers, and platforms. Federated learning enables local training and shared model development without routinely pooling raw datasets [1]. Decentralized computation nevertheless retains heterogeneity, communication, optimization, privacy, and coordination problems [2].

Federated learning separates local data and training from update exchange and aggregation. Unlike generic distributed computing, participating data are not necessarily centrally controlled. Farm nodes receive a model, train locally, generate

parameters, gradients, or deltas, and send approved updates to a coordinator that aggregates and redistributes a revised model [3].

Agricultural data may reveal production conditions, management choices, crop performance, resource use, machinery activity, land characteristics, and commercial patterns. Farms also vary in scale, climate, soil, equipment, connectivity, data quality, and capacity. Federated learning may reduce centralized pooling while leaving privacy, governance, communication, heterogeneity, and validation unresolved [4].

This article proposes a conceptual cross-farm architecture connecting use-case definition, stakeholder mapping, data sensitivity, local retention and training, update sharing, secure communication, aggregation, safeguards, validation, monitoring, traceability, accountability, and farmer participation. It rejects claims that federation automatically guarantees privacy, security, fairness, accuracy, trust, compliance, or readiness.

Privacy-Preserving Precision Agriculture

Privacy-preserving precision agriculture seeks useful intelligence while limiting unnecessary disclosure, concentration, transfer, or reuse. It spans agricultural AI, smart farming, sensor networks, remote sensing, decision support, and connected equipment. Federated learning can coordinate distributed holders while raw records remain local, but agricultural applications, clients, privacy assumptions, and evaluations vary [5].

Privacy depends on collected data and possible inferences. Sensors may reveal environmental, crop, livestock, and operational conditions; imagery exposes spatial patterns; management records describe yields, inputs, and practices. Federated crop classification supports decentralized collaboration, but one bounded experiment cannot establish universal suitability across crops, regions, devices, or objectives [6].

Data governance defines who authorizes processing, accesses artefacts, restricts use, and challenges decisions. A farm may retain raw data yet lose control over updates, metadata, or global models. Federated yield modelling supports shared prediction without conventional pooling, but usefulness still depends on quality, representativeness, validation, and site relationships [7]. Decentralization does not resolve authority or responsibility.

Confidentiality and cybersecurity differ from privacy. Local processing may reduce central-repository exposure, but compromised endpoints, insecure channels, malicious updates, and inference attacks remain possible. Computation, memory, connectivity, and energy also constrain safeguards [8]. The architecture makes no privacy or compliance guarantee and requires independent security, governance, and validation assessment.

Farm Data Sensitivity and Cross-Farm Learning

Farm-data sensitivity is contextual. Sensors may reveal schedules; imagery, field conditions; soil and yield records, productive capacity; machinery data, routes; and management records, commercial decisions. Cross-farm learning requires an inventory of data, purposes, harms, and restrictions. Trust also depends on transparency, benefit-sharing, accountability, and influence [9].

Ownership, access, and control are distinct. Storage does not confer unrestricted reuse authority, and technical access does not justify every purpose. Uncertainty about rights, uses, access, and control can discourage participation [10]. The architecture should record authorization, artefacts, retention, secondary-use limits, and change procedures.

Provenance, documentation, and quality matter because farms differ in collection, labels, calibration, missingness, timing, and context. Errors, biased samples, outdated observations, or incompatible definitions can distort updates and aggregation. Ethical and stakeholder concerns should remain integral to design [11], making provenance and quality review continuing gates.

Consent may be relevant but is neither universal nor sufficient. Where used, it should be understandable, specific, documented, and linked to realistic choices. Farmer participation should include feedback on relevance, limits, governance, and risk. Digital agriculture can reproduce unequal control and benefit [12]. **Table 1** summarises farm data sensitivity, cross-farm learning needs, privacy risks, governance concerns, and architecture roles in privacy-preserving precision agriculture.

Table 1. Farm Data Sensitivity and Cross-Farm Learning Requirements in Privacy-Preserving Precision Agriculture: Sensor Data, Remote Sensing, Soil Data, Crop Data, Weather Data, Machinery Data, Yield Data, Farm-Management Records, Data Ownership, Access, Control, Provenance, Farmer Trust, Consent, and Raw Data Sharing Boundaries

Farm data or learning component	Core privacy question	Required context or condition	Cross-farm learning function	Potential output	Privacy or governance risk	Safeguard or evaluation need	Architecture role
Sensor, weather, soil, crop, imagery, machinery, yield, management, and livestock data	What can records reveal about identity, operations, land, performance, or welfare?	Valid source, calibration, labels, time, location, granularity, and purpose	Environmental, image, forecasting, operational, or livestock learning	Local update	Re-identification, surveillance, bias, competitive or welfare harm	Minimisation, access control, provenance, domain review, local validation	Governance: local input
Ownership, access, and control	Who may authorize processing and future use?	Documented roles, rights, permissions, and withdrawal rules	Determines participation and permitted actions	Governance decision	Conflicting authority or loss of control	Governance review and escalation	Governance: gatekeeper
Provenance, documentation, quality, trust, and consent where supported	Are data traceable, reliable, and participation conditions understandable?	Versioned sources, limitations, practical choices, feedback channels	Supports reliable training and legitimate participation	Quality and participation record	Contamination, token consent, or trust-washing	Screening, audit logging, periodic review	Assurance and so: governance layer
Cross-farm development and raw-data boundary	What purpose justifies collaboration, and do raw records remain local?	Compatible task, governance agreement, threat model, data-flow map	Coordinates local training and aggregation	Candidate shared model	Hidden replicas, inference leakage, unequal benefit	Cross-farm validation, privacy testing, data-flow audit	Central function: mandate boundary

Federated Learning Principles and Agricultural Use Cases

Federated learning divides responsibilities between local nodes and coordinators. Farm or edge nodes retain approved data, train locally, and generate updates; a coordinator distributes models, schedules participants, aggregates updates, and returns revisions. Agricultural IoT environments include heterogeneous sensors, gateways, machinery, and devices with differing security and processing capabilities [13].

Edge federation may reduce high-volume data transfer, but it depends on processing, memory, power, bandwidth, latency, connectivity, and round reliability. Mobile-edge research identifies communication efficiency, client selection, resource allocation, and device heterogeneity as central [14]. Remote agriculture may amplify these constraints, so participation requires local resource and security assessment.

Horizontal federation fits farms with similar variables and different observations. Vertical federation may combine different features for aligned entities but adds linkage and governance risks. Transfer federation may assist related tasks with limited overlap. Non-IID conditions require explicit treatment because farms differ in crops, climate, labels, practices, hardware, and sampling [15].

Secure aggregation can hide individual plaintext updates but cannot prove accuracy, integrity, representativeness, or safety. The architecture must separate confidentiality from authentication, anomaly detection, robust aggregation, privacy protection, and validation. Different heterogeneity forms may require personalization, clustering, alternative aggregation, or redesigned evaluation [16].

Potential uses include crop monitoring, disease classification, yield prediction, remote-sensing interpretation, livestock monitoring, and human-reviewed decision support. Each requires assessment of sensitivity, labels, local relevance, uncertainty, and error consequences. Federated transfer learning illustrates bounded potential, but outcomes depend on data and assumptions [17]; no universal use or performance claim follows.

Model Development Without Centralized Raw Data Sharing

Cross-farm development separates raw data from permitted outbound artefacts. Each node retains approved records, trains a model, and generates parameters, deltas, or gradient-derived updates. A coordinator schedules nodes, distributes versions, screens updates, and aggregates accepted contributions. Clients, topology, aggregation assumptions, privacy mechanisms, and boundaries must be explicit [18]. **Figure 1** illustrates the workflow.




Figure 1. Federated Learning Workflow for Cross-Farm Model Development Without Centralized Raw Data Sharing

Alt-text description: Multiple farm nodes keep raw data local, train models, send protected updates to a coordinator, receive an aggregated model, validate locally, monitor drift, and return feedback.

The no-centralized-raw-data-sharing boundary requires end-to-end data-flow verification. Raw records may still appear in caches, backups, logs, diagnostics, or centralized preprocessing. Cross-silo agricultural learning permits shared modelling without raw-data exchange, but compatible objectives, controlled coordination, and validation remain necessary [19].

Agricultural heterogeneity affects training and aggregation. Differences in sample size, production, climate, machinery, sensors, resolution, labels, and practices create imbalance and conflicting local optima. Adaptation and pruning may help under site-specific and communication constraints, but remain task dependent [20]. Local, client-level, and cross-farm evaluation are required.

Bandwidth, connectivity, computation, memory, and energy determine participation. On-device smart-irrigation research makes resource measurement central [21]. Unequal connectivity can exclude farms and bias learning. Model and data drift should be monitored as seasons, sensors, management, and biological conditions change. Aggregation alone does not establish quality.

Privacy, Security, and Governance Requirements

Federated learning changes training location but not attack surfaces. Compromised participants may alter data, code, or updates; adversaries may intercept communications, steal credentials, exploit coordinators, or infer information. Model poisoning, data poisoning, inference attacks, malicious aggregation, and disruption remain concerns [22]. The architecture should define threat assumptions, rejection criteria, response, and recovery.

Privacy assessment should examine updates, gradients, metadata, participation records, outputs, and logs. Confidentiality requires authentication, authorization, encryption, protected storage, and secure communication. Secure aggregation, differential privacy, homomorphic encryption, or trusted execution provide conditional protection with different assumptions, costs, and privacy-utility trade-offs [23]. No safeguard is universal.

Differential privacy requires a specified protection objective, privacy unit, accounting method, and utility assessment. Noise may reduce selected inference risks, but repeated participation accumulates loss and stronger perturbation may weaken convergence or utility. Privacy interacts with participation, communication rounds, and performance [24], requiring documented parameters and local and global evaluation.

Homomorphic encryption supports computation over protected values, while trusted execution may isolate operations. Suitability depends on keys, hardware trust, computation, latency, dropout, attestation, and side-channel risk. Additive homomorphic encryption demonstrates protected computation with implementation and performance costs [25]. Neither mechanism establishes correctness, integrity, complete privacy, or certification.

Governance determines who authorizes participation, operates coordination, defines aggregation, inspects logs, modifies models, suspends nodes, and changes purpose. Auditability preserves records; traceability links outputs to data, code, models, and decisions; accountability assigns responsibility. Defence performance depends on assumptions and implementation [26]. **Table 2** maps these requirements.

Table 2. Federated Learning Principles, Decentralized Model Development, Privacy Safeguards, Security Risks, and Governance Requirements for Cross-Farm Agricultural AI: Local Data Retention, Local Training, Model Updates, Global Aggregation, Secure Aggregation, Differential Privacy, Data Heterogeneity, Non-IID Farm Data, Model Poisoning, Inference Attacks, Auditability, Traceability, and Claim Boundaries

Principle, component, requirement, or risk	Core architecture question	Required input or procedure	Federated or safeguard function	Potential output	Residual or failure risk	Validation, monitoring, or governance need	Boundary statement
Local retention, training, update sharing, and aggregation	What remains local, what leaves, and how is it combined?	Data-flow map, valid data, update format, aggregation rule	Produces local and candidate global models	Model states	Leakage, poisoning, dominance, instability	Local, client-level, and convergence analysis	Locality and aggregation do not provide privacy or quality
Secure communication, secure aggregation, differential privacy, homomorphic encryption, or TEE	Which safeguards fit the threat, utility, and resources?	Authentication, authorization, encryption, accounting, keys, or attestation	Protects channels, updates, or selected computation	Protected contribution	Endpoint compromise, collusion, utility loss, overhead	Privacy-utility and security assessment	Conditional safeguards not guaranteed
Horizontal, vertical, or transfer	Which relationship applies, and	Feature alignment, task relation,	Organises collaboration	Shared or adapted model	Linkage exposure, negative	Configuration-specific and	No universal configuration or

federation under heterogeneous farm conditions	how do clients differ?	distributions, devices, connectivity	and participation		transfer, imbalance, exclusion	per-client validation	performance claim
Poisoning, inference, auditability, traceability, accountability, and claim boundaries	Can learning be corrupted or information inferred, and can decisions be reconstructed?	Provenance, anomaly detection, attack modelling, versioned logs, role map	Flags risk and supports review	Alert, risk report, audit trail	Adaptive attacks, incomplete testing, missing lineage	Incident response, rollback, independent review	No privacy security, compliance or readiness guarantee

Proposed Federated Learning Architecture

The architecture begins with use-case and stakeholder definition: the problem, intended users, output, contexts, and need for collaboration. Relevant actors include farmers, operators, advisors, stewards, developers, providers, validators, and governance decision-makers. Sensitivity, quality, provenance, documentation, and control are assessed before training. Smart-agriculture research supports separating local modelling, privacy processing, communication, and aggregation [27].

The local layer contains approved data, preprocessing, training, evaluation, update generation, and resource monitoring. Updates pass through authentication, authorization, encryption, format checks, anomaly screening, and justified privacy mechanisms. Cybersecurity remains distinct from privacy because confidential systems may still contain compromised devices or malicious updates. Agricultural federated intrusion detection supports separate security monitoring [28].

The coordination layer manages eligibility, model distribution, scheduling, update receipt, aggregation, version control, and candidate generation. Participation policies should consider availability, connectivity, staleness, capability, and exclusion of poorly connected farms. Changing agricultural networks may require adaptive coordination [29], which should remain transparent, monitored, and fairness-reviewed.

The assurance layer evaluates global and local performance, cross-farm generalization, heterogeneity, bias, fairness, robustness, uncertainty, and drift. Candidate models should return to farms rather than be judged only by averages. Auditability links documentation, versions, updates, tests, incidents, and approvals. Blockchain-assisted federation may support integrity and provenance [30] but does not prove accuracy, trustworthiness, or legitimate governance.

The final layer links monitoring, farmer participation, governance review, and updating. Drift, failures, anomalous updates, incidents, local degradation, or farmer concerns can trigger reassessment, suspension, rollback, renewed validation, or termination. Remote settings may require hierarchical or satellite-assisted coordination with added infrastructure dependencies [31]. Table 3 presents the architecture.

Table 3. Proposed Federated Learning Architecture for Privacy-Preserving Cross-Farm Model Development in Precision Agriculture: Use-Case Definition, Governance Mapping, Data Sensitivity, Local Farm Nodes, Edge Training, Model Updates, Secure Communication, Privacy Safeguards, Security Safeguards, Federated Aggregation, Validation, Heterogeneity Assessment, Drift Monitoring, Auditability, Farmer Feedback, and Deployment Boundaries

Architecture stage	Core purpose	Required input	Architectural function	Potential output	Validation or assessment need	Failure risk	Feedback mechanism

Use-case, stakeholders, governance, and data review	Define purpose, authority, sensitivity, quality, and provenance	Agricultural problem, stakeholder map, data inventory and documentation	Sets scope and screens inputs	Approved use case and data reports	Necessity, representation, privacy, quality	Function creep, responsibility gaps, contaminated data	Stakeholder and steward revision
Farm-node configuration, local training, and update generation	Prepare controlled nodes and produce minimum necessary updates	Hardware, software, credentials, valid data, training rules	Hosts data, training, local evaluation, and update creation	Eligible node and local update	Misconfiguration, overfitting, leakage	Node testing, local validation, update screening	Unverified nodes exclude raw gradient require justification
Secure communication, safeguards, coordination, and aggregation	Protect transfer and combine eligible contributions	Threat model, channels, participation and aggregation rules	Transmits protected updates and produces a candidate model	Protected aggregate and global model	Compromise, false assurance, poisoning, dominance, staleness	Privacy-utility, security, convergence, rollback monitoring	No universal safeguard aggregation is not validated
Validation, heterogeneity, fairness, drift, auditability, and traceability	Test context-specific behaviour and preserve evidence	Structured holdouts, client metrics, longitudinal signals, logs	Assesses generalization, distributional effects, change, and lineage	Validation report, alert, audit trail	Hidden local failure, exclusion, missed drift, sensitive logging	Redesign, suspension, rollback, remediation	No reading claim; traceability does not prove validity
Farmer feedback, governance updating, and claim boundaries	Maintain participation and reassess purpose, risk, and evidence	Feedback, incidents, validation, residual-risk registers	Revises, restricts, or ends the federation	Updated decision	Token consultation or premature use	Formal change control and independent review	No privacy security performance compliance or deployment guarantee

Figure 2 presents the proposed federated learning architecture for privacy-preserving precision agriculture, integrating farm data sensitivity, local training, secure communication, aggregation, privacy safeguards, security safeguards, validation, governance, farmer feedback, and architecture updating.




Figure 2. Federated Learning Architecture for Privacy-Preserving Precision Agriculture

Alt-text description: A layered architecture shows governance, local farm data, local training, protected update exchange, aggregation, privacy and security controls, validation, monitoring, auditability, farmer feedback, and explicit decision boundaries.

Practical Implications and Future Research

For farmers, operators, advisors, cooperatives, stewards, and public-interest organisations, the architecture clarifies whether collaboration is necessary, what remains local, what is exchanged, who controls development, and how concerns are handled. Participation requires understandable information about purpose, risk, access, benefit, limits, and escalation. These implications are conceptual, not professional advice.

For providers, developers, model teams, and cybersecurity specialists, future research should examine agricultural non-IID data, sparse participation, unequal resources, intermittent connectivity, seasonal change, personalization, cross-farm validation, drift, update leakage, poisoning resistance, and privacy-mechanism trade-offs. Client-level outcomes should be reported because global averages may conceal local failure, exclusion, or unequal benefit.

For researchers, policymakers, public institutions, and governance bodies, longitudinal evaluation should test whether privacy risks re-emerge through metadata, repeated updates, model access, or linkage; whether auditability remains accurate; whether drift affects farms unequally; and whether responsibility remains clear when shared models fail. The architecture can identify questions for qualified review but does not establish compliance, prescribe regulation, endorse vendors, or authorize deployment.

Conclusion

This article proposes an architecture linking farm-data sensitivity, cross-farm learning, local retention, local training, update sharing, secure communication, aggregation, privacy safeguards, cybersecurity controls, governance, validation, monitoring, auditability, farmer participation, and architecture updating. Reducing centralized raw-data sharing may support collaboration but cannot establish privacy, security, fairness, accuracy, trust, compliance, reliability, or deployment readiness. Responsible development requires bounded use cases, stakeholder mapping, quality review, heterogeneity assessment, cross-farm validation, privacy and threat assessment, drift monitoring, traceability, and governance feedback.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 15 Aug 2025

Revised: 29 Sep 2025

Accepted: 18 Nov 2025

Published online: 25 May 2026

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

Yang Q, Liu Y, Chen T, Tong Y. Federated machine learning: Concept and applications. *ACM Trans Intell Syst Technol.* 2019;10(2):12:1-19.
<https://doi.org/10.1145/3298981>.

Li T, Sahu AK, Talwalkar A, Smith V. Federated learning: Challenges, methods, and future directions. *IEEE Signal Process Mag.* 2020;37(3):50-60.
<https://doi.org/10.1109/MSP.2020.2975749>.

Zhang C, Xie Y, Bai H, Yu B, Li W, Gao Y. A survey on federated learning. *Knowl Based Syst.* 2021;216:106775.
<https://doi.org/10.1016/j.knosys.2021.106775>.

Dembani R, Karvelas I, Akbar NA, Rizou S, Tegolo D, Fountas S. Agricultural data privacy and federated learning: A review of challenges and opportunities. *Comput Electron Agric.* 2025;232:110048.
<https://doi.org/10.1016/j.compag.2025.110048>.

Žalik KR, Žalik M. A review of federated learning in agriculture. *Sensors (Basel).* 2023;23(23):9566.
<https://doi.org/10.3390/s23239566>.

Idoje G, Dagiuklas T, Iqbal M. Federated learning: Crop classification in a smart farm decentralised network. *Smart Agric Technol.* 2023;5:100277.
<https://doi.org/10.1016/j.atech.2023.100277>.

Zhang Q, Zhao X, Han Y, Yang F, Pan S, Liu Z, et al. Maize yield prediction using federated random forest. *Comput Electron Agric.* 2023;210:107930.
<https://doi.org/10.1016/j.compag.2023.107930>.

Aggarwal M, Khullar V, Goyal N, Prola TA. Resource-efficient federated learning over IoAT for rice leaf disease classification. *Comput Electron Agric.* 2024;221:109001.
<https://doi.org/10.1016/j.compag.2024.109001>.

Jakku E, Taylor BM, Fleming A, Mason C, Fielke S, Sounness C, et al. "If they don't tell us what they do with it, why would we trust them?" Trust, transparency and benefit-sharing in smart farming. *NJAS Wagening J Life Sci.* 2019;90-91:100285.
<https://doi.org/10.1016/j.njas.2018.11.002>.

Wiseman L, Sanderson J, Zhang A, Jakku E. Farmers and their data: An examination of farmers' reluctance to share their data through the lens of the laws impacting smart farming. *NJAS Wagening J Life Sci.* 2019;90-91:100301.
<https://doi.org/10.1016/j.njas.2019.04.007>.

van der Burg S, Bogaardt MJ, Wolfert S. Ethics of smart farming: Current questions and directions for responsible innovation towards the future. *NJAS Wagening J Life Sci.* 2019;90-91:100289.
<https://doi.org/10.1016/j.njas.2019.01.001>.

Rotz S, Duncan E, Small M, Botschner J, Dara R, Mosby I, et al. The politics of digital agricultural technologies: A preliminary review. *Sociol Ruralis.* 2019;59(2):203-29.
<https://doi.org/10.1111/soru.12233>.

Nguyen DC, Ding M, Pathirana PN, Seneviratne A, Li J, Poor HV. Federated learning for Internet of Things: A comprehensive survey. *IEEE Commun Surv Tutor.* 2021;23(3):1622-58.
<https://doi.org/10.1109/COMST.2021.3075439>.

Lim WYB, Luong NC, Hoang DT, Jiao Y, Liang YC, Yang Q, et al. Federated learning in mobile edge networks: A comprehensive survey. *IEEE Commun Surv Tutor.* 2020;22(3):2031-63.
<https://doi.org/10.1109/COMST.2020.2986024>.

Sattler F, Wiedemann S, Müller KR, Samek W. Robust and communication-efficient federated learning from non-IID data. *IEEE Trans Neural Netw Learn Syst.* 2020;31(9):3400-13.
<https://doi.org/10.1109/TNNLS.2019.2944481>.

Zhu H, Xu J, Liu S, Jin Y. Federated learning on non-IID data: A survey. *Neurocomputing.* 2021;465:371-90.
<https://doi.org/10.1016/j.neucom.2021.07.098>.

Aggarwal M, Khullar V, Goyal N, Gautam R, Alblehai F, Elghatwary M, et al. Federated transfer learning for rice-leaf disease classification across multiclient cross-silo datasets. *Agronomy.* 2023;13(10):2483.

<https://doi.org/10.3390/agronomy13102483>.

Li Q, Wen Z, Wu Z, Hu S, Wang N, Li Y, et al. A survey on federated learning systems: Vision, hype and reality for data privacy and protection. *IEEE Trans Knowl Data Eng.* 2023;35(4):3347-66.

<https://doi.org/10.1109/TKDE.2021.3124599>.

Durrant A, Markovic M, Matthews D, May D, Enright J, Leontidis G. The role of cross-silo federated learning in facilitating data sharing in the agri-food sector. *Comput Electron Agric.* 2022;193:106648.

<https://doi.org/10.1016/j.compag.2021.106648>.

Li AH, Markovic M, Edwards P, Leontidis G. Model pruning enables localized and efficient federated learning for yield forecasting and data sharing. *Expert Syst Appl.* 2024;242:122847.

<https://doi.org/10.1016/j.eswa.2023.122847>.

Dakhia Z, Lazzaro A, Sebti MR, Russo M, Merenda M. On-device federated learning for energy-efficient smart irrigation. *Electronics.* 2025;14(21):4311.

<https://doi.org/10.3390/electronics14214311>.

Mothukuri V, Parizi RM, Pouriyeh S, Huang Y, Dehghantanha A, Srivastava G. A survey on security and privacy of federated learning. *Future Gener Comput Syst.* 2021;115:619-40.

<https://doi.org/10.1016/j.future.2020.10.007>.

Yin X, Zhu Y, Hu J. A comprehensive survey of privacy-preserving federated learning: A taxonomy, review, and future directions. *ACM Comput Surv.* 2021;54(6):131:1-36.

<https://doi.org/10.1145/3460427>.

Wei K, Li J, Ding M, Ma C, Yang HH, Farokhi F, et al. Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Trans Inf Forensics Secur.* 2020;15:3454-69.

<https://doi.org/10.1109/TIFS.2020.2988575>.

Phong LT, Aono Y, Hayashi T, Wang L, Moriai S. Privacy-preserving deep learning via additively homomorphic encryption. *IEEE Trans Inf Forensics Secur.* 2018;13(5):1333-45.

<https://doi.org/10.1109/TIFS.2017.2787987>.

Rodríguez-Barroso N, Jiménez-López D, Luzón MV, Herrera F, Martínez-Cámara E. Survey on federated learning threats: Concepts, taxonomy on attacks and defences, experimental study and challenges. *Inf Fusion.* 2023;90:148-73.

<https://doi.org/10.1016/j.inffus.2022.09.011>.

Kumar P, Gupta GP, Tripathi R. PEFL: Deep privacy-encoding-based federated learning framework for smart agriculture. *IEEE Micro.* 2022;42(1):33-40.

<https://doi.org/10.1109/MM.2021.3112476>.

Friha O, Ferrag MA, Shu L, Maglaras L, Choo KKR, Nafaa M. FELIDS: Federated learning-based intrusion detection system for agricultural Internet of Things. *J Parallel Distrib Comput.* 2022;165:17-31.

<https://doi.org/10.1016/j.jpdc.2022.03.003>.

Puppala S, Sinha K. Towards secure and efficient farming using self-regulating heterogeneous federated learning in dynamic network conditions. *Agriculture.* 2025;15(9):934.

<https://doi.org/10.3390/agriculture15090934>.

Manoj T, Makkithaya K, Narendra VG. A blockchain-assisted trusted federated learning for smart agriculture. *SN Comput Sci.* 2025;6(3):221.

<https://doi.org/10.1007/s42979-025-03672-4>.

Puppala S, Sinha K. Scalable satellite-assisted adaptive federated learning for robust precision farming. *Agronomy*. 2026;16(2):229.
<https://doi.org/10.3390/agronomy16020229>.