ZIBELINE INTERNATIONAL™ PUBLISHING

ISSN: 2521-0874 (Print)
ISSN: 2521-0505 (Online)
CODEN: AIMCCO

Acta Informatica Malaysia (AIM)

DOI: http://doi.org/10.26480/aim.01.2026.18.22

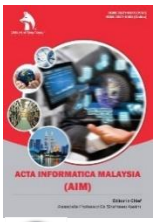

RESEARCH ARTICLE

BEYOND DATA SILOS: A PROPOSED FEDERATED LEARNING FRAMEWORK FOR INCLUSIVE GOVERNANCE

Tinashe Magara*, Mampilo Phahlane

Department of Information Systems, College of Science, Engineering, and Technology, University of South Africa, Pretoria, 0003, South Africa.
Corresponding Author Email: tinashenosh@gmail.com

ARTICLE DETAILS

Article History:

Received 13 March 2026
Revised 18 April 2026
Accepted 23 May 2026
Available online 11 June 2026

ABSTRACT

Fragmented data infrastructures, limited interoperability and persistent privacy concerns continue to hinder evidence based governance and weaken public trust. This study proposes a federated learning (FL) framework that enables institutions to collaboratively train machine learning models without sharing raw data, thereby preserving data sovereignty and reducing privacy risks while enhancing cross sectoral coordination. Grounded in a global review of FL applications, the framework incorporates four principles policy relevance, scientific robustness, trust and transparency, and feasibility tailored to the realities of low- and middle-income contexts. Its architecture integrates a four-pillar trust model using privacy-enhancing technologies, blockchain-based auditability, and participatory co-design, alongside application pathways in healthcare, education and economic inclusion. Aligned with national development priorities, the framework offers a scalable, ethical approach to modernizing public data ecosystems and transforming fragmented datasets into a trusted foundation for inclusive governance.

KEYWORDS

Federated Learning, Big Data, e-Governance, Policy Innovation, Data Sovereignty

1. INTRODUCTION

1.1 The Data Paradox in Modern Governance

Effective governance in the 21st century is predicated upon the timely and accurate analysis of comprehensive data. However, in many national and sub-national contexts, particularly in low and middle-income countries, the promise of data-driven policy remains unrealised due to the phenomenon of data silos (Richardson, 2021; Bibri, 2022). Data silos are structural and cultural impediments where institutional boundaries restrict the flow and interoperability of critical information (Saberi et al., 2025). These fragmentation challenges are pervasive, affecting essential services like healthcare where patient data remains isolated in clinics, education where school performance metrics are not linked to socio-economic indicators, and public finance where revenue and expenditure tracking lack real-time integration (Engin et al., 2020; Al-Assaf K et al., 2024; Zang et al., 2023). The resultant policy analysis is often partial, backward-looking, and vulnerable to systemic bias. Beyond technical interoperability, the single greatest constraint is the imperative for privacy protection and data sovereignty (Hellmeier et al., 2023; Hummel et al., 2021). Governments and citizens rightly resist centralised data models, given the heightened risks of massive data breaches, misuse of personal information, and the erosion of public trust (Li, 2025). This creates a critical data paradox: policymakers need integrated, cross-sectoral insights to solve complex national challenges, but the mechanisms traditionally used to achieve integration violate privacy and security norms (Magara and Zhou, 2024).

1.2 Federated Learning as an Enabling Technology

This study proposes that federated learning (FL) offers the necessary technical and ethical resolution to this paradox. FL is a distributed machine learning paradigm where the algorithm is sent to the data, rather than the data being sent to a central server (Liu et al., 2022; Ma et al., 2022). Participating institutions clients train a shared global model using their local, proprietary datasets, and only transmit the resulting model updates gradients or weights back to a central orchestrator (Zhong et al., 2022). This process inherently preserves data locality, ensuring that raw, sensitive datasets never leave the secure confines of the original data custodian (Chen et al., 2025). FL represents a strategic pivot from the collect all data to one place model. Which is institutionally difficult and ethically precarious, to a collaborate on insights-in-place model, which respects data custodianship and privacy by design principles (Chiaro et al., 2025; Zhu et al., 2021).

1.3 Research Objectives and Framework Alignment

The primary objective of this research is to design a concrete, adaptable FL framework that can enable inclusive, evidence-based governance. This is achieved through three key objectives:

- Systematically assess the global technical and institutional feasibility of FL, particularly in contexts with fragmented data infrastructure.
- Design and elaborate a novel governance structure, the four-pillar trust model, that embeds security and accountability mechanisms into the FL pipeline.
- Illustrate the framework's utility through high-impact, scenario-based applications aligned with the national strategic priorities for inclusive growth phase III such as climate adaptation, poverty monitoring.

The proposed framework is not merely a theoretical exercise; it is an adapted technical and governance architecture aligned with socio-technical environment, acknowledging resource constraints, the need

Quick Response Code	Access this article online	
	Website: www.actainformaticamalaysia.com	**DOI:** 10.26480/aim.01.2026.18.22

Cite The Article: Tinashe Magara, Mampilo Phahlane (2026). Beyond Data Silos: A Proposed Federated Learning Framework For Inclusive Governance. *Acta Informatica Malaysia, (AIM)*, 10, (1): 18-22.

for cost-efficiency, and the complexity of multi-ministerial coordination. The remainder of this paper is organized as follows. **Section II** presents the Materials and Methods, outlining the overall research design and methodological approach. **Section III** discusses the Results and Discussion, evaluating the technical feasibility of the proposed framework and demonstrating its practical applicability across three key domains Healthcare Access, Education Equity, and Economic Inclusion followed by an analysis of governance and institutional adoption. Finally, **Section IV** provides the conclusion, summarizing the core contributions of the study and emphasizing how the proposed framework addresses the data fragmentation paradox while advancing alignment with national development and policy priorities.

2. Materials and Methods

2.1 Systematic Evidence Review and Scoping

The study began with a systematic evidence review (SER) to establish the global state of federated learning (FL) implementation across governance, healthcare, and finance. Following PRISMA guidelines, the review systematically searched major electronic databases Scopus, Web of Science, IEEE Xplore, and ACM Digital Library and grey literature sources from international policy institutions. Search terms combined "Federated Learning," "Distributed AI," "privacy-preserving computation," "e-governance," and "low-resource setting." (Kwao et al., 2023).

Extracted data focused on three dimensions:

- Technical Architectures, including aggregation strategies (FedAvg), horizontal and vertical configurations, and the use of privacy-enhancing technologies (PETs);
- Governance models, covering client selection, incentive design, and compliance mechanisms; and
- Performance Metrics, assessing communication efficiency, resilience under Non-IID data, and model accuracy. Synthesized evidence informed the adaptation of FL to environments with heterogeneous data sources, intermittent connectivity, and limited computational capacity typical of low- and middle-income contexts.

2.2 Contextual Adaptation and Design Principles

The framework design was adapted to the socio-technical and institutional realities of resource-constrained settings, emphasizing scalability, cost-efficiency, and policy alignment (Bahna et al., 2025). Four core design principles guided the architectural and operational decisions:

- Policy Relevance – Ensuring direct alignment with national development priorities such as poverty reduction, climate resilience, and service delivery improvement to foster institutional adoption.
- Scientific Robustness – Embedding reproducibility, model provenance tracking, and transparent validation to enhance the credibility of data-driven insights.
- Trust and Transparency – Integrating accountability and audit mechanisms to ensure ethical governance and public legitimacy.
- Feasibility – Employing asynchronous FL protocols, lightweight communication schemes, and open-source frameworks (TensorFlow Federated, PySyft) to mitigate infrastructural constraints.

2.3 Architectural Design: The Four-Pillar Trust Model

The proposed architecture adopts a centralized federated learning (FL) topology anchored in a four-pillar trust model that embeds privacy, accountability, inclusivity, and transparency throughout the training lifecycle. The model ensures that governance objectives are integrated directly into the technical design (Figure 1).

- Data Security: Model updates are encrypted using secure multi-party computation (SMPC) or homomorphic encryption (HE), ensuring that gradients transmitted between institutions remain inaccessible to aggregators, while client-level differential privacy (DP) introduces calibrated noise to prevent inference attacks and safeguard sensitive information.
- Auditability: A permissioned blockchain like the hyperledger fabric records each training round's cryptographic hashes, client participation metadata, and model versioning, creating an immutable audit trail for regulatory verification and institutional accountability.
- Inclusivity: Participatory co-design workshops engage policymakers, data custodians, ethicists, and civil society in defining ethical boundaries, relevant feature spaces, and acceptable model performance thresholds, strengthening legitimacy and ensuring policy relevance.
- Transparency: Public-facing dashboards disseminate aggregated, non-sensitive performance indicators such as coverage metrics and model validation summaries reinforcing public trust and enabling external scrutiny without compromising confidentiality.

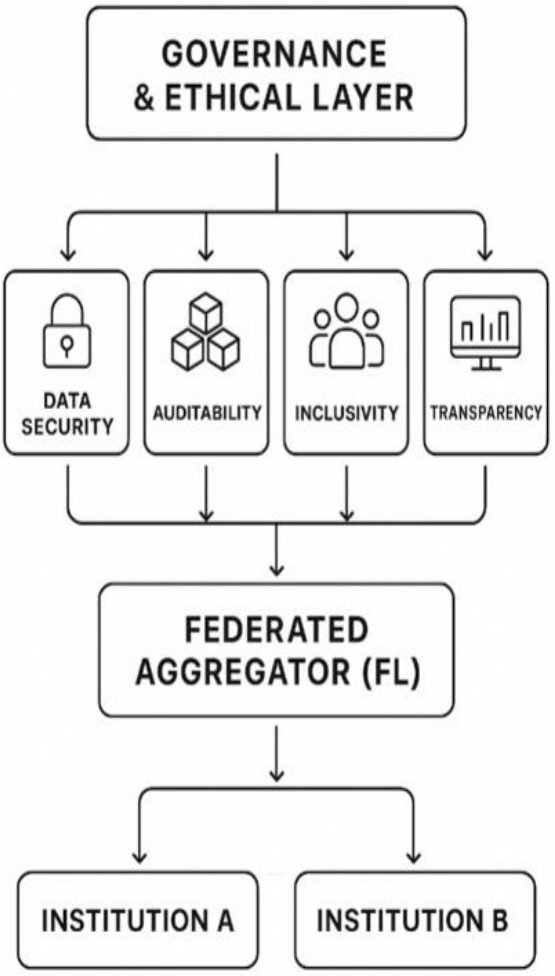

Figure 1: Proposed Centralized Federated Learning Architecture for Inclusive Governance

Building upon this foundation, the federated learning framework for inclusive governance (FLIG) integrates privacy-preserving distributed learning with governance primitives to enable inclusive, transparent, robust, and accountable decision-making systems. In this framework, heterogeneous governance nodes (government agencies, NGOs, and citizen data hubs) collaboratively train a global model w_t without centralizing raw data. At each round t, a stratified subset S_t of clients is selected to ensure representational fairness, and each client $k \in S_t$ performs local optimization $w_t^k = w_t - \eta \Delta L_K(w_t)$, where $L_K(w_t)$ denotes the local objective. The differentially private, clipped update $\hat{U}_k = Clip(w_t^k - w_t, C) + N(0, \sigma^2 I)$ is securely aggregated using a Byzantine-robust fairness-aware function (Zheng et al., 2025; Bowen et al., 2024).

$$w_{t+1}^k = w_t - \eta \sum_{k \in S_t} \rho_k A_{robust}(\hat{U}_k),$$

where ρ_k are reweighting coefficients derived from group-level fairness constraints that minimize loss disparities across protected groups,

$$\min_w \sum_k \frac{n_k}{N} L_k(w) + \alpha \sum_{g \in G} (L_g(w) - \hat{L}(w))^2$$

A governance module (GM) manages policy voting, audit trails, and transparency logs (G_{log}), while a reputation and incentive mechanism quantifies stakeholder contributions through approximate Shapley values, ensuring equitable participation and accountability. This unified algorithmic framework, denoted as $FPR - FedAvg - G$, operationalizes the four-pillar trust architecture through secure aggregation, differential privacy (ε, δ)-DP guarantees, and fairness-constrained optimization thus

aligning federated intelligence with ethical, human-centric, and transparent governance principles.

2.4 Application Pathways: Scenario-Based Validation

To demonstrate applied feasibility, the framework is evaluated through scenario-based analyses in three high-impact policy domains: healthcare access, education equity, and economic inclusion. Each scenario models institutional collaboration under horizontal or vertical FL configurations, testing interoperability, communication efficiency, and model generalization across distributed data silos. These scenarios serve as empirical pathways for assessing the framework's technical, governance, and ethical viability in advancing evidence-based and inclusive policymaking.

3. Results and Discussion

3.1 Technical Feasibility and Heterogeneity Management

The technical feasibility of the proposed federated learning (FL) framework is demonstrated through its use of established, communication-efficient algorithms such as FedAvg, validated in distributed machine learning environments. The primary technical challenge in the target context lies in handling non-independent and identically distributed (Non-IID) data, reflecting the heterogeneity of datasets across ministries and agencies—for instance, clinical data from health systems versus economic records from fiscal institutions.

To mitigate model drift and bias arising from this heterogeneity, the framework incorporates advanced FL variants such as FedProx and FedNova, which stabilize gradient updates and enhance model generalization across disparate domains.

Additionally, asynchronous aggregation strategies ensure resilience against client dropout caused by intermittent network connectivity, enabling continuous global model updates even in bandwidth-constrained environments. As illustrated in Figure 2, these technical adaptations underpin diverse application pathways of federated learning in governance spanning public health analytics, fiscal transparency, urban management, and social service optimization demonstrating the framework's integrative potential across sectors.

Collectively, these adaptations confirm the framework's capacity to operate effectively in low-resource governance settings where infrastructural variability and data fragmentation are prevalent.

3.2 Healthcare Access and Disease Surveillance

Within the healthcare domain, the FL framework transforms fragmented health data into a unified, privacy-preserving intelligence system for service delivery improvement. National disease surveillance and resource allocation have traditionally been constrained by the distributed nature of health data across hospitals, clinics, and ministries. Through a horizontal FL configuration, each participating healthcare facility trains a local model on its patient data to forecast resource demands and disease trajectories. Only the model parameters are shared with the central Ministry of Health, which aggregates them into a national predictive model. This configuration enables accurate forecasting of resource needs, such as ventilators or essential medicines, without compromising patient confidentiality. The integrated four-pillar trust model reinforces this process by maintaining blockchain-based audit logs, guaranteeing accountability to both policymakers and citizens, and ensuring that data-driven interventions remain ethical and transparent.

3.3 Education Equity and Dropout Prediction

In education, the framework addresses equity and inclusion challenges through secure, privacy-preserving integration of education and socio-economic datasets. Conventional data sharing between ministries, such as linking student performance data with household income, poses serious privacy risks and regulatory barriers. The proposed vertical FL approach overcomes these challenges by enabling joint model training across institutions without exposing raw data. Through secure multi-party computation (SMPC), sensitive variables such as student income or academic grades remain encrypted while the model learns from their correlations. The resulting predictive model identifies students at risk of dropping out based on both academic and socio-economic indicators, facilitating targeted interventions such as conditional cash transfers or community-based support programs. The approach directly enhances evidence-based policymaking for poverty reduction and educational inclusion while preserving family and student privacy.

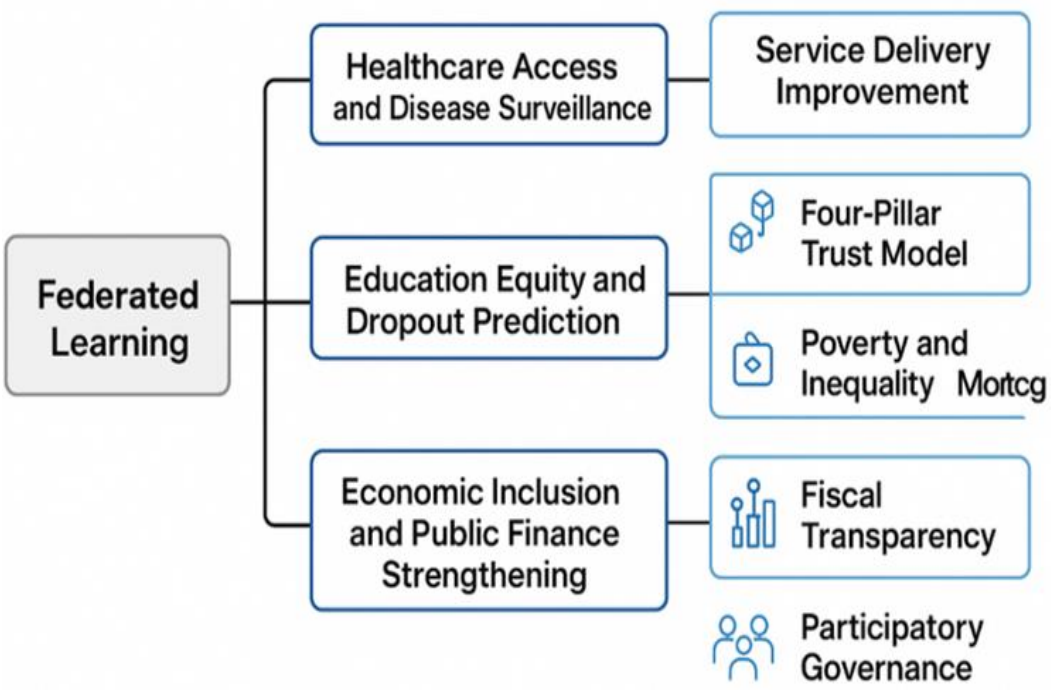

Figure 2: Application Pathways of Federated Learning in Governance

3.4 Economic Inclusion and Public Finance Strengthening

In the financial governance context, the framework supports advanced fraud detection, tax compliance monitoring, and social protection targeting. Detecting large-scale financial irregularities traditionally requires linking transaction-level data from the ministry of finance with macroeconomic indicators from the central bank—an operation limited by strict data protection laws. The FL framework overcomes this by employing horizontal FL across decentralized tax offices to train a unified compliance model, while vertical FL links banking and social protection datasets to refine eligibility models for social welfare programs. The inclusion of blockchain-based audit trails guarantees transparency in model development and deployment, offering immutable proof of analytical integrity. This combination of privacy-preserving analytics and verifiable accountability enhances fiscal governance, enabling more efficient resource targeting and strengthening public trust in financial oversight.Table 1 provides a cross-sector comparative summary of the three application pathways, detailing the FL configuration employed, the data sources integrated, the privacy mechanisms activated, and the key policy outcomes enabled, thereby highlighting the framework's versatility and transferability across heterogeneous governance contexts.

Table 1: Cross-Sector Application of the Federated Learning Framework

Sector	FL Configuration	Data Sources Integrated	Privacy Mechanisms Employed	Key Policy Outcomes Enabled
Healthcare Access	Horizontal FL	Hospital EHRs; clinic-level patient records; Ministry of Health resource inventories; disease surveillance registers.	Client-level Differential Privacy (DP); FedAvg aggregation; blockchain audit trail for training rounds.	Predictive resource allocation (ventilators, medicines); national disease trajectory forecasting; privacy-compliant cross-facility epidemiological intelligence.
Education Equity	Vertical FL	Ministry of Education student performance records; Ministry of Social Affairs household income data; school attendance registers.	Secure Multi-Party Computation (SMPC); encrypted joint model training across ministries; fairness-constrained aggregation.	Early dropout risk identification; targeted conditional cash transfer eligibility; evidence-based educational inclusion policy without exposing family financial data.
Economic Inclusion	Horizontal FL + Vertical FL (hybrid)	Decentralised tax office transaction data; Central Bank macroeconomic indicators; Social Protection Ministry welfare eligibility datasets.	Blockchain audit trails; differential privacy; Byzantine-robust aggregation (FPR-FedAvg-G); SMPC for cross-institution linkage.	Unified tax compliance model; enhanced fraud detection; refined social welfare eligibility targeting; strengthened fiscal transparency and public trust in financial governance.

3.5 Governance and Institutional Adoption

The success of the FL framework is determined not only by its technical soundness but by its institutional adaptability and governance legitimacy. The four-pillar trust model, anchored in privacy, auditability, inclusivity, and transparency, creates the socio-technical foundation necessary for multi-agency collaboration.

By delivering quantifiable privacy guarantees through differential privacy and SMPC, and immutable accountability through blockchain, the framework addresses the principal deterrents to cross-sectoral data sharing: legal liability, institutional distrust, and reputational risk. The co-design process, central to the Inclusivity Pillar, ensures that policy actors, data custodians, and civil society jointly define the ethical and operational parameters of the system.

This participatory model shifts the governance paradigm from data control to data stewardship, fostering a cooperative ecosystem where shared analytical value replaces institutional competition. Ultimately, the framework demonstrates that privacy-preserving collaboration can serve as both a technical and governance innovation, transforming fragmented data infrastructures into a trusted, policy-relevant foundation for inclusive national development.

4. Conclusion

The endemic challenge of data fragmentation and legitimate privacy concerns has historically crippled evidence-based governance, creating a critical hurdle for inclusive national development.

This study proposes the federated learning (FL) framework, beyond data silos, as a strategic, ethical, and scalable blueprint to resolve this paradox. By enforcing data sovereignty through decentralised computation, the framework ensures that advanced, cross-sectoral predictive analytics, vital for addressing the nation's most complex challenges are achievable without compromising citizen rights or institutional control.

The framework's core innovation is the integration of the four-pillar trust model with robust technical FL architecture. This combination provides a system that is not only scientifically sound (scientific robustness and non-IID data handling) but also legally and politically viable (trust and transparency via PETs and Blockchain Auditability).

The demonstrated application pathways in healthcare, education, and economic inclusion confirm its direct alignment with the inclusive growth phase III priorities, enabling concrete actions in climate adaptation, service delivery, and poverty monitoring.

While the framework is conceptual and requires real-world pilots to validate performance against infrastructural and socio-political friction, its theoretical grounding in over a hundred global FL case studies in comparable settings suggests high transferability and feasibility.

The adoption of this framework would mark a pivotal shift towards a model of governance that is inherently privacy-conscious, participatory, and analytically powerful. By enabling institutions to collaborate on insights while retaining control over their data, the FL framework offers the most viable pathway yet to transform fragmented data assets into a trusted, collective foundation for equitable and sustainable national development.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. Funding for this research was covered by the authors of the article.

References

Al-Assaf, K., Bahroun, Z., Ahmed, V. 2024. Transforming service quality in healthcare: A comprehensive review of healthcare 4.0 and its impact on healthcare service quality. In Informatics, 11 (4), 96. MDPI.

Bahna, M., and Szalay, Z.G., 2025. Barriers and enablers of cloud computing adoption in Syria: a literature review grounded in the TOE framework. Journal of Agricultural Informatics/Agrárinformatika Folyóirat, 16(2).

Bibri, S. E. 2022. The social shaping of the metaverse as an alternative to the imaginaries of data-driven smart Cities: A study in science, technology, and society. Smart Cities, 5(3), 832-874.

Bowen, D., Haiquan, W., Yuxuan, L., Zhao, J., Ma, Y., and Runhe, H. 2024. Fair and robust federated learning via decentralized and adaptive aggregation based on blockchain. ACM Transactions on Sensor Networks.

Chen, C., Liu, J., Tan, H., Li, X., Wang, K. I. K., Li, P., Dou, D. 2024. Trustworthy federated learning: privacy, security, and beyond. arXiv preprint arXiv:2411.01583.

Chiaro, D., Qi, P., Mele, V., and Piccialli, F. 2025. FLAIR: Federated Learning for Augmented Industrial Retrieval. IEEE Internet of Things Journal.

Engin, Z., Van Dijk, J., Lan, T., Longley, P. A., Treleaven, P., Batty, M., and Penn, A. 2020. Data-driven urban management: Mapping the landscape. Journal of Urban Management, 9(2), 140-150.

Hellmeier, M., Pampus, J., Qarawlus, H., and Howar, F. 2023. Implementing data sovereignty: Requirements & challenges from practice. In Proceedings of the 18th international conference on availability, reliability and security (pp. 1-9).

Hummel, P., Braun, M., Tretter, M., and Dabrock, P. 2021. Data sovereignty: A review. Big Data & Society, 8(1), 2053951720982012.

Kwao Dawson, J., Twum, F., Hayfron Acquah, J. B., and Missah, Y. M. 2023. PRISMA Archetype-Based Systematic Literature Review of Security Algorithms in the Cloud. Security and Communication Networks,

2023(1), 9210803.

Li, L. 2025. The Construction of the Relationship between the Chinese Government and Public Trust: A Case Study of Online Personal Information Security. The Frontiers of Society, Science and Technology, 7(4).

Liu, J., Huang, J., Zhou, Y., Li, X., Ji, S., Xiong, H., and Dou, D. 2021. From distributed machine learning to federated learning: A survey. arXiv preprint arXiv:2104.14362.

Ma, C., Li, J., Shi, L., Ding, M., Wang, T., Han, Z., and Poor, H. V. 2022. When federated learning meets blockchain: A new distributed learning paradigm. IEEE Computational Intelligence Magazine, 17(3), 26-33.

Magara T. and Zhou Y. 2024. Internet of things (IoT) of smart homes: privacy and security. Journal of Electrical and Computer Engineering, 2024(1), 7716956.

Richardson, S. 2021. Against generalisation: Data-driven decisions need context to be human-compatible. Business Information Review, 38(4), 162-169.

Saberi, M.A., Mcheick, H., and Adda, M. 2025. From data silos to health records without borders: a systematic survey on patient-centered data interoperability. Information, 16(2), 106.

Zang, C., Zhang, Y., Xu, J., Bian, J., Morozyuk, D., Schenck, E. J., and Kaushal, R. 2023. Data-driven analysis to understand long COVID using electronic health records from the RECOVER initiative. Nature Communications, 14(1), 1948.

Zheng, H., Wang, R., Wang, X., and Ling, Q. 2025. Can fairness and robustness be simultaneously achieved under byzantine attacks?. In ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (pp. 1-5). IEEE.

Zhong, Z., Xu, M., Rodriguez, M. A., Xu, C., and Buyya, R. 2022. Machine learning-based orchestration of containers: A taxonomy and future directions. ACM Computing Surveys (CSUR), 54(10s), 1-35.

Zhu, H., Wang, R., Jin, Y., and Liang, K. 2021. PIVODL: Privacy-preserving vertical federated learning over distributed labels. IEEE Transactions on Artificial Intelligence, 4(5), 988-1001.