

Federated Learning in Food Safety Surveillance: Architectures, Applications, and Open Challenges

Akinode John Lekan¹, Opeoluwa Oluyemisi Oyinola²

Department of Computer Science, Federal Polytechnic Ilaro, Ogun State, Nigeria¹

Department of Science Laboratory Technology, Federal Polytechnic Ilaro, Ogun State, Nigeria²

Abstract: The world food supply chains are increasingly being challenged to ensure safety and quality. As the chains break up and become geographically dispersed, traditional centralized surveillance systems fall behind. They rely on data processing in pools, which leads to conflicts with data sovereignty, business confidentiality, and privacy laws. This review explores Federated Learning, a decentralised solution that allows multiple stakeholders to collaboratively learn a model without sharing raw data. The study looks at the advancements that have been published in recent years and assesses the potential of edge-enabled and blockchain-backed federated systems in enhancing food safety monitoring. A systematic review of the new federated architectures and their application in food safety was carried out. Potential application scenarios for pathogen detection, spoilage prediction, fraud detection and traceability along the supply chain are explored. Various horizontal and vertical data partitioning options were evaluated to determine their suitability for different scenarios of operations. Decentralised federated frameworks always performed better than isolated local models. Their ability to balance privacy and performance was fairly similar to in the case of fully centralised systems. Blockchain-integrated and edge-deployed configurations proved to be especially promising for real-time cold-chain and contamination monitoring. Federated Learning is a promising approach to achieving privacy-preserving and high-performance food safety intelligence. However, there are still challenges that need to be addressed with respect to non-independent and identically distributed (non-IID) data and communication overhead in resource-constrained environments. To protect public health in a digitalised global food system, the path is going towards zero-trust federated frameworks.

Keywords—Federated learning; food safety; data privacy; distributed systems; pathogen detection; blockchain; edge computing; supply chain surveillance.

1 Introduction

Foodborne diseases continue to impose a significant burden on global public health. The World Health Organization (2020) indicates that annually, approximately 1 in 6 people in the world get a food borne illness. They are not isolated health crises, but events that produce market volatility and affect trade and consumer trust in food systems (Garcia, Osburn & Jay-Russell, 2020). The spread of geographically dispersed supply chains, and the increasing complexity of their traceability, makes it more difficult to monitor the biological and chemical hazards of such chains at several jurisdictions. The traditional regulatory approach (reactive inspection regimes and localised manual testing) is no longer sufficient in today's context of fast-changing and large-scale safety threats (Kanicheril Ambikalekshmi et al., 2025).

A plethora of intelligent tools, including the use of biosensors, high throughput genome sequencing, and automated production monitoring systems are new technologies that have added to the surveillance capabilities of the farm. However, it is a challenge to transform all this information into actionable intelligence. Safety data is usually dispersed amongst competing commercial organisations and private research institutes, and stakeholders are generally unwilling to share safety data. In addition, concerns about liability, reputation and the loss of competitive advantage are all disincentives (Qian et al., 2022). There is an additional layer of complexity with legal frameworks. For example, the European Union's General Data Protection Regulation (GDPR) limits the consolidation of data on the central level,

relating to consumer behavior or hygiene records from the facilities. This creates a landscape of isolated "data islands" that reduce the predictive power of traditional machine learning models (Fodor et al., 2025).

Centralized machine learning attempts to overcome this by gathering into a single repository all the training data. Yet, the high communication cost in the presence of high-dimensional data and the single-point-of-failure nature of the system as well as high privacy risk (Aledhari, Razzak, Parizi, & Saeed, 2020) are potential drawbacks of this approach. There is another way: federated learning. It does not collect data from all participants but gives them a shared model to train locally and only sends the model parameters (gradients or weights) to the aggregator. The raw data stays where it's from. This architecture will enable privacy to be preserved while benefiting from the collective knowledge found in various, distributed data sets (McMahan et al., 2023).

The use of federated approaches has started to be introduced in various food safety applications, including pathogen detection on food preparation surfaces or shelf-life prediction for perishable food during transport. The aim of this review is to explore and discuss the emerging architectures and real-world applications of federated learning in food safety surveillance, in a systematic manner. It documents the process of conceptual prototypes being moved to real-world implementation, assesses the different system configurations to address the unique needs of the food industry, and outlines the technical and regulatory challenges that still need to be

overcome. Its intent is to help shape future monitoring tools that will safeguard public health and data privacy.

2. Fundamentals of Federated Learning

One of the key features of federated learning is to separate the training process of the model from direct access to the data. A standard configuration for federated is the distribution of a global model from a central server to participating clients. The model is trained by each client; this can be a food processing plant, a mobile sensor unit, or a regional laboratory. The client then only sends the modified model parameters back to the aggregating server, and not the data itself. These updates are then aggregated by the server, typically using the Federated Averaging (FedAvg) approach, to create a consolidated global model that is subsequently shared for subsequent local training (McMahan et al., 2023).

Learning is a process of repetition in this paradigm. The number of local client operations in each round of communication is kept small, so that there is a limited amount of data that needs to be sent over the network. This property is particularly important in the food industry where sensors are placed in remote agricultural locations that are often subject to limited bandwidth conditions (Park et al., 2021). Most importantly, sensitive information from the organisation—such as a particular processing plant's hygienic status or proprietary formulation information—is kept behind the organisation's firewall all along the way.

Distributed systems are usually categorized based on the way data is shared among the participants. In horizontal federated learning, all clients have access to the same features but the data held by each client is from a different population segment. A common configuration of this type of poultry farm consists of several farms that monitor multiple environmental indicators of Salmonella prevalence using the same type of sensors (Fendor et al., 2025). Vertical federated learning is used when clients have several features for the same entities. This structure works well for a scenario of traceability of a supply chain: a producer, a logistics provider and a retailer have different parts of information concerning the same product batch. A third type is federated transfer learning, which is a method to transfer knowledge from one domain (e.g., a controlled laboratory environment) to another, but related, domain (e.g., an industrial processing line) (Saha & Ahmad, 2021).

This framework has several benefits, such as increased data variety, enhanced privacy assurance, and better scalability. Models developed from multiple data sets tend to be more accurate and are better generalised than those developed from a single organisation's data. That said, there are some limitations to that approach. Model convergence is complicated by one feature of statistical heterogeneity — significant variation in the data distributions across clients — which is still an active topic of research (Zhao et al., 2018). Although raw data is never shared, the updates to the models themselves can contain sensitive information. However, there

are some advanced techniques like gradient inversion and membership inference techniques that can be used by sophisticated adversaries, so additional protection like differential privacy is often needed (Wei et al., 2020).

3. Emerging Architectures for Food Safety Surveillance

Within the food sector, the implementation of federated intelligence has led to the emergence of a number of different system designs, depending on the logistical and regulatory requirements of the context in which they are used.

In the cross-organisational federated learning, or cross-silo learning, usually there are many large institutional participants. Large food manufacturers, national food agencies or academic research groups may work together to establish the trend of the occurrence of chemical hazards. These configurations are able to support complex deep learning models such as detecting patterns of pesticide residue in multiple provincial surveillance datasets (Yu et al., 2022), because each silo has significant computational resources.

Edge-architectures solve another priority: real-time responsiveness at the point of production and/or retail. These systems can detect safety violations without needing to make a round trip to a central cloud, moving computation to the edge of the network — smart sensors, IoT gateways, or mobile devices. Stephan et al. (2025) has shown this with a smart vending cart system which applied edge intelligence to locally classify food freshness. Vendors were alerted to worsening storage conditions with near instantaneous notification — even without ongoing cloud connectivity.

Hierarchical architectures offer a hierarchical aggregation structure that is ideal for managing a global supply chain. In these schemes, the sensor models from neighbors are aggregated at a regional level and then an updated ensemble is sent to a global coordinator. This is a design that minimizes the number of top level interactions and enables the regional models to include localised patterns for safety. Blockchain-based systems are often used in such an architecture to improve accountability. It was demonstrated by Zhang et al. (2023) that the introduction of a blockchain ledger to record the model parameters guarantees that the learning process is not manipulated by any single participant. The immutability of the record also ensures it offers the audit trail consumers and regulatory compliance require.

The next generation of frameworks, called multimodal, combines visual, textual and sensor data streams into a single safety platform, and is a further evolution of IoT-based multimodal frameworks. Zhang et al. (2025) proposed a system integrating computer vision to detect surface defects with real-time sensor monitoring of the cold chain's integrity. These platforms often include virtual representations of physical assets or supply chains known as "digital twins," enabling safety professionals to explore contamination

scenarios and simulate interventions before they are used in real-world situations (Nasirahmadi & Hensel, 2022).

4. Applications in Food Safety

In recent years, various federated implementations have been developed to address specific safety and quality issues along the farm-to-fork continuum. Below, some of the most important advances in four fields are discussed.

4.1 Pathogen Detection and Contamination Monitoring

In a federated system, Taheri Gorji et al. (2023) created a system based on fluorescence imaging to identify organic residues and biofilms on food preparation surfaces. The deep learning models were trained at each of the participating facilities, allowing all the images to be classified for contamination at more than 94% accuracy, without images being transferred from one facility to the other. The results showed that it is not necessary to sacrifice detection performance when collaborating in a way that respects privacy.

Jia et al. (2024) explored the use of paper chromogenic array sensors coupled with neural networks for detection of viable foodborne pathogens, such as *Salmonella* and *Escherichia coli*, in the complex background of microbial communities in cheese and poultry products. The results suggest that the low cost sensor technologies can be used as federatable data sources in environments with poor lab facilities.

4.2 Spoilage Prediction and Cold-Chain Monitoring

For perishables, cold-chain integrity is vital, and it is a frequent proximate cause of losses due to spoilage. The FAAS-Chain framework was developed by Alrashdi and Taloba (2026) for the detection of spoilage in smart packaging. Temporal autoencoders at the edge were used to track gas emissions and environmental fluctuations to identify anomalies signaling early bacterial growth at 98.5% accuracy. Stacking ensemble methods in a federated configuration worked to near perfection to classify freshness of a wide variety of fruit types at the street-level food vending level, allowing for automatic environmental changes in resource-limited environments, as shown by Stephan et al. (2025).

4.3 Food Authenticity and Fraud Detection

Large and complex global supply chains foster opportunities for fraudulent activity, such as mislabelling, adulteration, and substitution, to go undetected for many years. Gavai et al. (2023) modeled a scenario where the confirmed fraud cases were not shared among data custodians either. They were able to achieve this by federating their Bayesian Network, which learned fraud patterns by combining model updates from various country-based data stations, while maintaining data sovereignty at each station with the same performance as a fully centralised model. Federated spectral analysis has been applied in related work to identify adulteration events from various production sites in a variety of food types and lab environments (Gulati et al., 2023).

4.4 Risk Assessment and Disease Surveillance

A risk assessment and disease surveillance process is in place. There is a risk assessment and disease surveillance process in place.

Federated approaches have also bolstered quantitative approaches to chemical risk assessment. A risk entropy model was proposed by Yu et al. (2022) to combine the pesticide detection information from provincial regulatory departments. Local entropy calculations generated a more complex picture of safety risk than a simple pass or fail rating, and were used to identify geographic hotspots that required targeted intervention. In addition to chemical hazards, federated learning is being investigated with regards to zoonotic source attribution. This approach could help public health authorities identify and trace back strains of pathogens to individual producers while maintaining the protection of individual producer anonymity (Zhang et al., 2019).

5. Challenges and Limitations

However, there are several remaining technical and regulatory challenges to translating federated learning into practical applications for food safety.

5.1 Data Heterogeneity and Resource Constraints

Food safety data has the characteristic of being variable. Data distributions vary significantly by client due to environmental differences, management practices, seasonal changes and production norms. This inevitably results in a well-known non-IID issue, with the statistical skew between client sites resulting in poor modelling convergence, or even prediction bias (Saeed, 2025). The challenge is added by resource limitations. The food sector is the sector with the greatest deployment of edge devices, especially in lower income production areas, and they do not have enough computational power nor energy supply to train complex model architectures in the local area. The design of efficient, lightweight systems while maintaining accuracy with limited computational load is still a challenge for large-scale deployment in agriculture (Stephan et al., 2025).

5.2 Communication Overhead and Scalability

When many edge nodes need to keep up to date copies of the model on the central server, a significant amount of bandwidth is required. This is a known limitation especially for multimodal systems with many parameters (Singh, Tripathi & Adhikari, 2023). The more clients that join in, the more challenging it is to handle synchronisation rounds, and even accept a drop-out of clients, which is something you're sure to experience if you have a large network spread across multiple continents. The lightweight network architectures and efficient parameter quantisation techniques are the prerequisite for real-world deployment in agricultural networks (Nachiappan, 2024).

5.3 Privacy leakages and Security vulnerabilities

Sensitive information can be included in the gradient updates sent during the training process, although the raw data is never sent out of the client. Under specific conditions, adversarial methods like model inversion attacks and membership inference attacks can partially reconstruct the training data or determine if a particular record was part of the training set (Rigaki & Garcia, 2023). Poisoning attacks are another attack type that can compromise federated systems, and prevent the global model from working well. The adoption of secure aggregation protocols and anomaly detection mechanisms for identifying suspicious updates, are thus not considered as mere features but as integral components to the integrity of any federated food safety network (Zeng et al., 2024).

The ability to comply with regulations and understand models.5.4 Regulatory Compliance and Model Interpretability.

Regulatory acceptance of federated decision-support tools requires that those tools be explainable. A lot of deep learning architectures are black-box models, with very little explanation of what aspects of the input were responsible for a specific prediction. This transparency erodes the confidence of food safety practitioners and regulators (Ranasinghe, Ramanan, & Fernando, 2022). Leveraging existing EU policies, such as the AI Act and GDPR, at the policy level, there is currently no specific guidance for decentralised learning systems, with compliance uncertainty. Another challenge is that there are no standardised benchmark datasets available for food safety tasks, which makes it difficult to compare federated algorithms objectively and hinders external validation (Fendor et al., 2025).

6. Future Research Directions

Several converging technological developments are likely to shape the next phase of federated food safety research.

Explainable AI (XAI) will be critical for building the regulatory confidence that these systems currently lack. Future frameworks should incorporate interpretability mechanisms that reveal how specific inputs — a temperature excursion, an anomalous gas reading — contribute to a risk prediction. Embedding visualisation tools within federated architectures would give safety professionals the means to interrogate and validate model outputs in operationally meaningful ways (Nachiappan, 2024).

Transfer learning and self-supervised pre-training offer a practical route around the label scarcity that affects many food safety datasets. Models pre-trained on large publicly available datasets could be fine-tuned through federated protocols, enabling accurate performance even where local labelled data is sparse (Hsieh et al., 2023). The emergence of large-scale foundation models — trained across diverse data modalities — may accelerate this further, producing adaptable safety platforms that can be rapidly repurposed for newly identified hazards.

Integrating zero-trust security principles with blockchain verification would substantially strengthen federated networks. In a zero-trust configuration, every model update is treated as potentially adversarial and requires continuous cryptographic verification before aggregation. This is a meaningful step beyond current secure aggregation methods. Alongside this, synthetic data generation offers a way to augment training sets for rare or low-frequency hazard events — specific chemical contaminants, for instance — without exposing actual facility or consumer records (Little, Elliot, & Allmendinger, 2023).

Finally, the development of interoperable, internationally harmonised food safety intelligence platforms would enable cross-border collaboration at a scale not currently achievable. Such platforms would allow countries to share safety insights and early warning signals while maintaining strict data sovereignty — addressing a gap that is particularly acute during transnational outbreak investigations.

7. Conclusion

Federated learning directly addresses the tension between the demand for data-intensive safety analytics and the necessity of protecting proprietary and personal information. This paradigm will enable the food industry to optimise the model collectively without having to centralize the raw data, thus using the distributed, isolated datasets to construct early warning systems that could not otherwise be built. This review has been able to show that progress has been made in the area of pathogen detection, spoilage prediction, fraud detection and pesticide risk assessment, each of which has shown that decentralised models can perform with the same effectiveness as centralised models under many practical conditions.

Several challenges remain. Limited acceptance of widespread use remains due to data distributions that are not iid and communication efficiency. The lack of federated-specific regulatory guidance, as well as the limited interpretability of deep learning models, create additional challenges to embedding deep learning into official safety frameworks. These challenges can only be overcome by developing a standardised, interoperable system design that prioritises interpretability and security, rather than addressing them as an afterthought.

Edge computing and trust mechanisms based on blockchain are a promising direction for achieving real-time and scalable food safety monitoring within complex supply chains. These technologies are ripe to become key infrastructure for the global food safety governance system as these technologies mature. Implementing that potential will require decades of interdisciplinary research, a synergy of computer scientists, food microbiologists, regulatory scientists, and public health professionals, built on a shared vision of a secure, data-driven food system, and a commitment to data sovereignty.

References

- Adedeji, A. A., et al. (2024). Artificial intelligence in diagnosing allergens in food. *Applied Sciences*.
- Aledhari, M., Razzak, R., Parizi, R. M., & Saeed, F. (2020). Federated learning: A survey on enabling technologies, protocols, and applications. *IEEE Access*, 8, 140699–140725.
- Alrashdi, I., & Taloba, A. I. (2026). AI-driven smart packaging: enhancing sustainability and cybersecurity in food supply chains. *Frontiers in Sustainable Food Systems*, 9, 1712080.
- Bera, S., Dey, T., Mukherjee, A., & Buyya, R. (2023). E-CropReco: a dew-edge-based multi-parametric crop recommendation framework for internet of agricultural things. *Journal of Supercomputing*, 79(11), 11965–11999.
- Fendor, Z., van der Velden, B. H. M., Wang, X., Carnoli, A. J., Mutlu, O., & Hürriyetoğlu, A. (2025). Federated learning in food research. *Journal of Agriculture and Food Research*, 23, 102238.
- Garcia, S. N., Osburn, B. I., & Jay-Russell, M. T. (2020). One health for food safety, food security, and sustainable food production. *Frontiers in Sustainable Food Systems*, 4, 1.
- Gavai, A., Bouzembrak, Y., Mu, W., Martin, F., Kaliyaperumal, R., van Soest, J., Choudhury, A., Heringa, J., Dekker, A., & Marvin, H. J. P. (2023). Applying federated learning to combat food fraud in food supply chains. *npj Science of Food*, 7(1), 46.
- Gulati, M., Dadkhah, N., Groß, B., Wunder, G., Glavonjic, J., Pavlovic, A., & Tomcic, A. (2023). BETA-FL: blockchain-event triggered asynchronous federated learning in supply chains. *EasyChair*.
- Hsieh, K., Moore, E., Ramage, D., & Arcas, B. A. (2023). Federated learning with transfer learning for target space optimization. *arXiv*.
- Jia, Y., et al. (2024). Deep feed-forward neural networks for paper chromogenic array sensors. *Food Control*.
- Kanicheril Ambikalekshmi, A., Pushparaj, P., Cherian, E., Rajan, R., & Mohan, L. (2025). Advancing food safety through artificial intelligence: A detailed review of monitoring and detection technologies. *Food Safety and Health*.
- Little, C., Elliot, M., & Allmendinger, R. (2023). Federated learning for generating synthetic data: a scoping review. *International Journal of Population Data Science*, 8(1), 2158.
- McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. (2023). Communication-efficient learning of deep networks from decentralised data. *arXiv*.
- Nachiappan, B. (2024). Enhancing system efficiency through AI, edge computing, and resource optimisation in modern infrastructure. *International Journal of Innovative Science and Research Technology*, 9(10), 1107–1112.
- Nasirahmadi, A., & Hensel, O. (2022). Digital twins in the agri-food sector. *Journal of Agriculture and Food Research*.
- Park, S., Jung, S., Lee, J., Kim, K., & Kim, Y. (2021). Strategies for water quality monitoring using federated learning. *Applied Sciences*.
- Qian, C., Liu, Y., Barnett-Neefs, C., Salgia, S., Serbetci, O., Adalja, A., Acharya, J., Zhao, Q., Ivanek, R., & Wiedmann, M. (2022). A perspective on data sharing in digital food safety systems. *Critical Reviews in Food Science and Nutrition*.
- Ranasinghe, N., Ramanan, A., & Fernando, S. L. (2022). Interpretability and accessibility of machine learning in food processing. *arXiv*.
- Rigaki, M., & Garcia, S. (2023). A survey of privacy attacks in machine learning. *ACM Computing Surveys*, 56(4), 1–34.
- Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., & Maier-Hein, K. (2020). The future of digital health with federated learning. *npj Digital Medicine*, 3(1), 1–7.

- Saeed, F. (2025). Challenges in federated learning for heterogeneous environments.
- Saha, S., & Ahmad, T. (2021). Federated transfer learning: concept and applications. *Intelligent Artificial*, 15(1), 35–44.
- Singh, N., Tripathi, T., & Adhikari, M. (2023). Communication-efficient federated learning for real-time applications. *IEEE International Conference on Advanced Networks*.
- Stephan, T., Paramana, P. P. D., Lin, C. C., Agarwal, S., & Verma, R. (2025). Federated learning-driven IoT system for automated freshness monitoring in resource-constrained vending carts. *Journal of Big Data*, 12, 33.
- Taheri Gorji, H., Saeedi, M., Mushtaq, E., Kashani Zadeh, H., Husarik, K., Shahabi, S. M., Qin, J., Chan, D. E., Baek, I., Kim, M. S., Akhbardeh, A., Sokolov, S., Avestimehr, S., MacKinnon, N., & Vasefi, F. (2023). Federated learning for clients' data privacy assurance in food service industry. *Applied Sciences*, 13(16), 9330.
- Wei, K., Li, J., Ding, M., Ma, C., Yang, H. H., Farokhi, F., Jin, S., Quek, T. Q. S., & Vincent Poor, H. (2020). Federated learning with differential privacy: algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security*, 15, 3454–3469.
- World Health Organization. (2020). Food safety: key facts. WHO.
- Yu, J., Chen, Y., Wang, Z., Liu, J., & Huang, B. (2022). Food risk entropy model based on federated learning. *Applied Sciences*, 12(10), 5174.
- Zeng, G. Q., Shao, J. M., Lu, K. D., Geng, G. G., & Weng, J. (2024). Automated federated learning-based adversarial attack and defence in industrial control systems. *IET Cyber-Systems and Robotics*, 6(2).
- Zhang, Q., Lu, Z., Liu, Z., Li, J., Chang, M., & Zuo, M. (2025). Application of machine learning in food safety risk assessment. *Foods*, 14, 4005.
- Zhang, W., et al. (2025). Multimodal deep learning-based intelligent food safety detection and traceability system. *International Journal of Management Science Research*, 8(3), 71–76.
- Zhang, X., Tan, X., Xu, J., Luo, S., & Qi, Z. (2023). A trusted sharing model for risk information based on blockchain and federated learning. *IEEE International Conference on Blockchain*.
- Zhang, Y., Tyner, K., Post, L., & Wiedmann, M. (2019). Machine learning identifies signatures of host adaptation in *Salmonella enterica*. *PLOS Genetics*, 14.
- Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., & Chandra, V. (2018). Federated learning with non-IID data. *arXiv*.