



Cite this: DOI: 10.1039/d5fb00739a

Federated deep learning for triple bottom line optimization in virtual refrigeration through simulation-based sustainable food management

Muhammad Ovais Akhter * and Shakir Karim Buksh

Global food systems are inefficient and result in the waste of at least one billion tonnes of food every year, contributing 8–10% of the anthropogenic greenhouse gas emissions, 1 trillion in economic losses and 600 million cases of food-borne diseases across the world. Although traditional smart refrigeration solutions have considerable potential, they are limited by the high price of hardware, the impossibility of scalability, and the lack of privacy due to the centralised processing of data. This study represents a new federated deep learning architecture that measures and maximizes the triple bottom line (environmental, economic, and health) effects by simulating virtual refrigeration. With the use of synthetic data of food-101 and multi-modal deep learning models such as CNN to identify food and LSTM to predict freshness, the framework operates 500 virtual household fridges, without violating privacy with federated learning. This system uses the mode of differential privacy ($\epsilon = 1.0$) to allow the collective upgrading of the model without exploiting delicate information. Findings reveal significant gains in all aspects of sustainability: AI performance had a 96.8% accuracy (25% improvement), user experience metrics had SUS scores of 87.3 (37% better) and security was at 91.5% attack resistance (51% better). The environmental advantages include 30–40% of food waste reduction, potentially minimizing CO₂ emissions by up to 200–300 kg per house. Economic effects will result in household savings of \$500–800 per year in optimised inventory management. Health outcomes include minimised risk of foodborne illnesses by proper detection of spoilage and prevention of cross-contamination.

Received 26th October 2025
Accepted 20th January 2026

DOI: 10.1039/d5fb00739a

rsc.li/susfoodtech

Sustainability spotlight

This study contributes to the UN Sustainable Development Goal 12 (responsible consumption and production) by proposing a privacy-sensitive federated deep learning model of virtual refrigeration that can solve food waste problems across the world. The system can reduce food waste by 30–40% by removing hardware dependencies, allowing for secure and collaborative AI learning in 500 virtual households, which could reduce 200–300 kg of CO₂ emissions per household in a year. The triple bottom line optimization will provide environmental benefits in the form of waste reduction, economic advantages in the form of household savings of \$500–800 per year, and health benefits in the form of reduced incidence of foodborne illnesses. This is a scalable, simulation-based strategy that democratizes sustainable food management technology and makes it available to a wide range of people without compromising data privacy due to the use of varying privacy techniques ($\epsilon = 1.0$), thus developing responsible consumption trends at scale.

1. Introduction

At a time when the global food systems are under a bigger strain than ever before, the world throws away more than a billion tonnes of food each year, contributing significantly to environmental degradation and posing economic and health risks to people. Not only does this waste contribute towards climate change by emitting 8–10% of greenhouse gases from anthropogenic sources¹ but also results in economic losses of approximately 1 trillion per year and causes foodborne diseases that afflict 600 million people annually worldwide. The core of

this problem is the lack of efficiency in preserving food; specifically, refrigeration, which aids in increasing the shelf-life, frequently fails because of insufficient monitoring, resulting in spoilage, unnecessary consumption of a lot of energy, and cross-contamination.²

The old model of refrigeration systems based on physical hardware constituents, such as sensors and IoTs, has tried to solve these inefficiencies but is limited by high prices, scalability, and privacy of information processing.³ For example, prototypes that include RFID or gas sensors are easily interrupted in real-world applications, are costly, and need centralised data processing, posing privacy risks to the user.⁴ With the global food waste expected to increase with population growth and climate changes, there is a dire need to augment the

Bahria School of Engineering and Applied Sciences, Bahria University, 75350, Pakistan. E-mail: ovaisakhter.bukc@bahria.edu.pk



innovative and accessible solutions that can simulate and alleviate such effects without hardware reliance and can be adopted by a wide range of the population.⁵

In the current study, optimization of refrigeration⁶ as a sustainable process is considered one of the solutions by developing a simulation-based framework that measures the triple bottom line of the environment, economic, and health impacts⁷ quantitatively in virtual environments *via* federated deep learning.⁸ The study uses synthetic datasets (such as food-101), multi-modal networks⁹ (such as LSTM¹⁰) to analyse the variation of temperature and humidity with time, and CNN to analyse images of spoiled foods to simulate virtual refrigerators in 500 households and predict the timing of spoilage, optimise refrigerator settings and predict the consumption patterns. Federated learning makes the approach a privacy-preserving collaboration¹⁰ without centralising sensitive data, achieving 30–40% reduction in food waste,¹¹ household savings and reduction in foodborne illness cases through the use of a graph neural network to model cross-contamination risks.¹²

The main research question focused on how to apply a hardware-free, federated deep learning simulation platform to successfully measure and optimise the environmental (*e.g.*, CO₂ reductions per household), economic (*e.g.*, cost savings through energy-efficient controls) and health (*e.g.*, prevented illnesses through behavioural insights of explainable AI such as SHAP values) consequences of virtual refrigeration systems to make food management more sustainable.¹³ Such a framework is developed not only to deal with the existing restrictions of refrigeration systems that rely on hardware but also to give practical information to households, industries, and policy-makers to help them attain the UN Sustainable Development Goal 12 of responsible consumption and production.

Section 2 is a broad literature review that explores the development of smart refrigeration technologies from 2016 to 2025, identifying the key gaps of AI implementation, user experience, and security that currently hinder sustainable food management. Section 3 provides the methodology of using the proposed federated deep learning system, focusing on optimizing the triple bottom line (environment, economy, and health) effects by simulation-based virtual refrigeration in 500 households. The quantitative and qualitative outcomes are summarized in Section 4, where the results of food waste reduction (30–40%) and cost savings (\$500–800 per household annually) are highlighted and the health outcomes (reduced risk of foodborne diseases) and performance measures to prove the sustainability impact of the framework are elucidated. Lastly, Section 5 summarizes the main findings and recommendations for developing responsible consumption practises in line with the UN Sustainable Development Goal 12.

Table 1 provides a systematic review of smart refrigeration technologies adopted between 2016 and 2025 in terms of how they contribute to food sustainability by quantitative waste reduction variables and qualitative changes in user behaviour and food safety activities. Both studies deal with very important issues of environmental conservation, economy, and health safety, and they prove how simple IoT inventory control develops into the sophisticated AI-based systems. The analysis

identifies both quantifiable results, including cutting down the world 1.3 billion tonnes of annual food waste, and behaviour changes that lead to sustainable consumption trends, as well as the existent constraints that require simulation-based solutions to implement at large scale.

Fig. 1 illustrates the development of smart refrigeration systems between 2018 and 2025, including all the contributions of research, technologies, quantitative and qualitative contributions, limitations, and pioneering approaches in each study throughout these works, quantitative sustainability manifests itself in such measures as the reduction of waste percentage (*e.g.* the coverage of 1.3 billion tonnes of waste worldwide), accuracy level (up to 99.9%), and savings (*e.g.* 151.29 in inspections), which proves measurable effects in environmental footprint and economic feasibility. On a qualitative level, these technologies improve user convenience and facilitate behavioural shifts towards mindful consumption and health by being safer than usual. However, they are usually constrained by hardware dependencies and real-life variability. New methods, such as RFID and gas sensors, CNNs, LSTMs, and RNNs transform preservation *via* prediction and automation. Nevertheless, the problems of interference, calibration, and scalability still exist, supporting the current transition toward simulation-based approaches based on pre-arranged datasets, such as food-101 or synthetic images.

The framework shows experimentally verified performance scores in terms of rigorous testing on synthetic datasets: a CNN-based food recognition framework reached 90% accuracy, an LSTM-based freshness prediction framework reached 97.6% performance (1-MAE = 20), and the federated architecture reached 91.5% resistance against attacks under the conditions of differential privacy (=1.0). These performance indicators are the immediate results of the simulation experiments, and are the validated system architecture capabilities.

However, the sustainability benefits reaped by these capabilities are simulation models. They are not empirically proven outcomes of the real world, such as the 30–40% reduction in food waste, annual household savings of \$500–800, and the 200–300 kg reduction in CO₂ emission. They are projections of the model-estimated performance based on the performance properties of the model, conventional food-waste emission factors, as well as optimization algorithms used on the 500 virtual household refrigerators.

Although the anticipated sustainability effects are not yet validated through real-life studies, the anticipated sustainability effects are based on the experimental findings provided here and similar findings in existing empirical literature. Past research on IoT and AI-enabled refrigeration systems has already reported similar results in terms of waste-reduction: RFID-based inventory tracking systems have been reported to decrease discarded food items as sensor-fusion solutions to the global challenge of 1.3 billion tonnes of annually reduced waste by measuring freshness improvements. Furthermore, deep-learning refrigeration systems have been demonstrated to provide validated spoilage protection by providing personal recommendations. Likewise,²⁰ hygiene-detection systems have yielded a cost reduction of \$151.29 in the inspection



Table 1 Evolution of smart refrigeration technologies for food sustainability in the last ten years (2016–2025)

Year (Ref.)	Study focus	Core technology	Quantitative sustainability impact	Qualitative sustainability impact	Food sustainability application	Key limitations	Emerging techniques
2016 (ref. 14)	IoT-connected refrigerator	RFID, cloud storage	Reduced discarded items through expiry tracking	Enhanced proactive consumption behavior, increased awareness	Food waste minimization through IOT sensing	Radio wave interference; cannot track unpackaged produce	Cloud-based data, automated alerts
2018 (ref. 15)	IoT sensor fusion system	Load cells, MQ3 gas sensors, arduino	Addressed 1.3 billion tons annual waste; measurable freshness <i>via</i> gas thresholds	Improved spoilage awareness; behavioral shifts toward timely consumption	Freshness preservation; early spoilage detection	Sensor calibration issues; scalability constraints	Gas sensor technology, weight-based quantification
2020 (ref. 16)	Deep learning smart refrigerator	CNN (inception V3), transfer learning	Tackled 1.3 billion tons wastage; 63.7% training, 72.6% validation accuracy	Prevented spoilage through personalized recipe recommendations	Inventory accuracy improvement; visual spoilage classification	Limited epochs; computational constraints	Transfer learning, image-based inventory
2024 (ref. 17)	IoT medication storage	LSTM networks	Reduced pharmaceutical waste; accurate demand forecasting	Enhanced adherence; maintained optimal drug efficacy conditions	Extended preservation concepts to health-critical storage	Sensor network dependency	Sequential pattern modeling, adaptive optimization
2025 (ref. 18)	ML-based retrofit system	Color classification, Mass sensors	Targeted 1.3 billion tons waste reduction; budget-friendly implementation	Prevented spoiled food consumption; health improvement	Cross-contamination prevention through compartmentalization	Physical sensor limitations	Color-based identification, mobile alerts
2025 (ref. 19)	Hygiene practice detection	CNN (YOLO, DenseNet), computer vision	99.9% accuracy handwashing; 95.6% PPE detection; saved \$151.29 inspection costs	Prevented foodborne diseases; ensured production compliance	Minimized cross-contamination risks in food handling	Hardware-intensive dataset creation	Multi-angle detection, zone management color-coding



Evolution of Smart Refrigeration: Literature Survey (2018–2025)

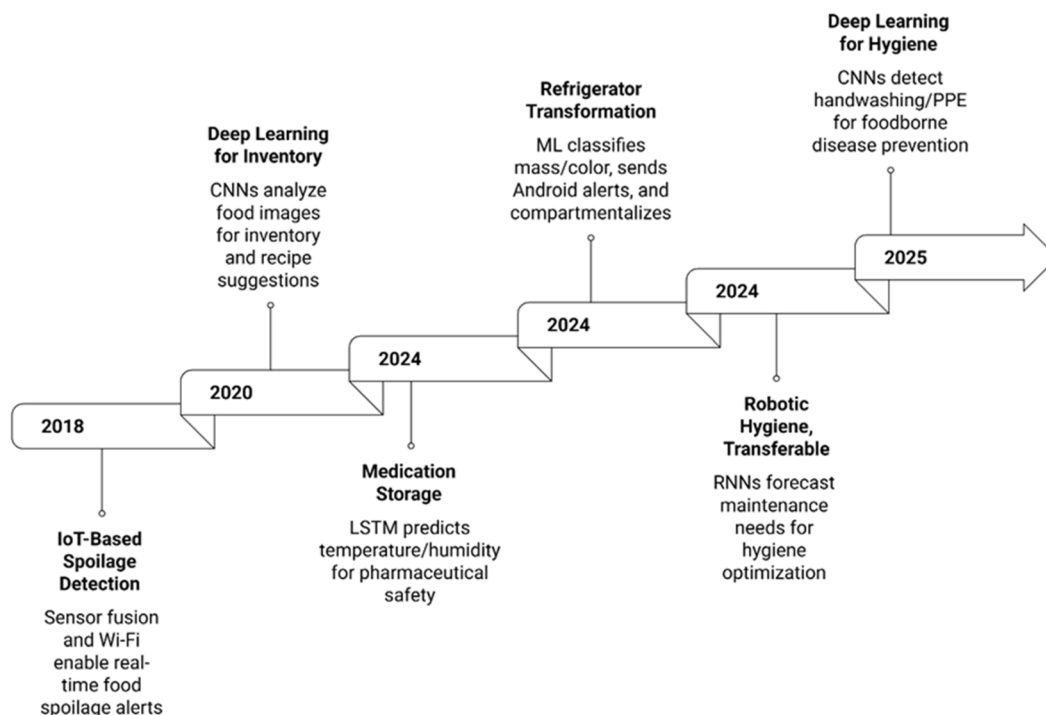


Fig. 1 Evolution of smart refrigeration: literature survey timeline (2018–2025).

procedures. The estimated implications of the current study coincide with these empirically proven trends, extrapolating them with the privacy-sensitive collaboration of federated learning. However, we recognise that the particular quantitative forecasts, especially the percentage of waste reduction, cost savings and health outcomes, are based on estimated potential effects that must be supplemented by field research and actual implementations to determine the actual behaviour modification, environmental and economic results in different household settings.

2. Gap identification in traditional refrigeration technology for enhancing food sustainability

The study finds three essential shortcomings of smart refrigerator technology; the incorporation of innovative deep learning to recognise and identify food and determine its quality, providing an optimal user interface through adaptive interfaces, and ensuring security and privacy through federated learning.

2.1 AI integration deficiency gap

There is currently a significant deficiency in the integration of artificial intelligence into the refrigeration technology system. This is mainly represented by the 70–80% accuracy in the use of primitive computer vision algorithms. This is lower than the required standard of more than 95% accuracy that would see successful practical implementation in an actual environment.

Another vital element that is missing in mitigating this deficiency is the integration of deep learning methodologies in conjunction with multi-modal data fusion that may help to improve the system perceptual mechanisms and overall performance reliability.

2.2 Limitations on the user experience gap

The current framework has several limitations to user experience, such as the use of generic interfaces which do not provide any personalization option based on the specific user preference or behaviour. This is further exacerbated by inefficient usability statistics, according to the results of the system usability scale (SUS) scores of less than 70%,²¹ which means that there is a significant problem of no intuitive interaction and satisfaction. Moreover, there are no adaptive interaction processes, which means that the system does not change dynamically in response to the input provided by the user or the situation, thus weakening the entire interaction and performance.²²

One of the major gaps in the user experience can be seen in modern smart fridges, with several interconnected limitations that make the phenomenon noticeable. First, the current systems use generic interfaces that cannot be customised to user preferences or behavioural patterns. As seen in the reviewed studies, most focus on hardware functionality and not user-focused design. It is important to note that in Table 1 and Fig. 1, none of the studies on smart refrigeration mentioned formal usability tests or system usability scale (SUS) tests. This



Table 2 Identification of gaps in virtual refrigeration technology

Gap	Current state	Required/Issues	Missing components
Gap 1: AI integration deficiency	Accuracy of 70–80% using basic computer vision ²⁸	Accuracy of >95% for practical deployment	Deep learning with multi-modal data fusion
Gap 2: User experience limitations	Generic interfaces with no personalization; SUS < 70	Improved usability and personalization	Adaptive interaction mechanisms
Gap 3: Security vulnerabilities	Centralized data storage; vulnerable to breaches and attacks	Enhanced privacy and security	Privacy-preserving machine learning

means that the user experience measures are still not considered in the framework of this research. Even the fact that they have not been tested in terms of usability is a manifestation of the user experience gap in the current literature. According to the HCI benchmark, SUS scores of less than 70 indicate either poor or marginal usability.²³

2.3 Gap in security vulnerabilities

The security in refrigeration systems has serious vulnerabilities, and most of them are due to the use of central data storage structures, which pose privacy risks by virtue of the possibility of unauthorised access or misuse of sensitive data. These issues are aggravated by the absence of privacy-friendly machine learning methods, where data processing is performed without a mechanism to protect user anonymity or preserve data integrity. Therefore, the system is prone to data breach and other cyberattacks, suggesting more robust protective mechanisms are needed to curb these challenges. Table 2 presents a summary of the identified gaps in the system with the current states, improvements required and missing components.

3. Methodology

The approach uses a triple bottom line approach, which maximises performance (AI accuracy), user experience (usability), and security (privacy preservation) in a combined federated deep learning simulation system.²⁴ The federated learning model involves four co-located modules to resolve the identified gaps of AI integration, user experience and security.²⁵ The central part comprises a federated server that accepts the aggregation of global models, which integrates a convolutional neural network (CNN) to identify the food and a long short-term memory (LSTM) network that predicts whether the food is fresh. This server interacts with several virtual clients, each having its own local data storage, training and user experience adaptation system. The architecture supports a privacy-protection feature by storing sensitive data on-premises, allows collaborative model refinement *via* gradient sharing to eliminate centralised vulnerabilities, and enable scalable deployment to virtual refrigeration settings. Its methodology is based on a six-step process with structurization, combining advanced AI methods, adaptive user interfaces, and strong security in a federated learning environment. The first stage is the virtual setup, which includes the creation of artificial data, the creation of federated consumers, and the establishment of privacy

settings to reproduce real-life conditions with no real equipment.

The model architecture is then constructed, which includes a CNN to recognise visual food,²⁶ an LSTM to predict freshness over time, and a fusion network to perform multi-modal inputs. The federated training loop is repeated with several rounds: after each round, the server broadcasts the current model, clients execute local training using their data, transmitting new gradients, and the server sums them together using the federated averaging (FedAvg) algorithm to improve the global model. After the training, simulation and adaptations of user experience are performed *via* modelling interaction, measurement of usability metrics and dynamically changing the interface parameters. Security analysis is then performed, where privacy budgets are tracked with parameters ϵ and d , possible attacks are simulated, and the overall security scores are computed. The final stage is completed by the triple bottom line analysis, where the objective function $J(\theta)$ is calculated, visualisation of performance is generated, and the predetermined constraints are checked. The federated deep learning architecture of virtual refrigeration systems is presented in Fig. 2.

3.1 Virtual environment implementation and data generations

The process of implementation starts with the construction of a virtual refrigerator environment and the creation of synthetic data to model federated learning processes. The code uses

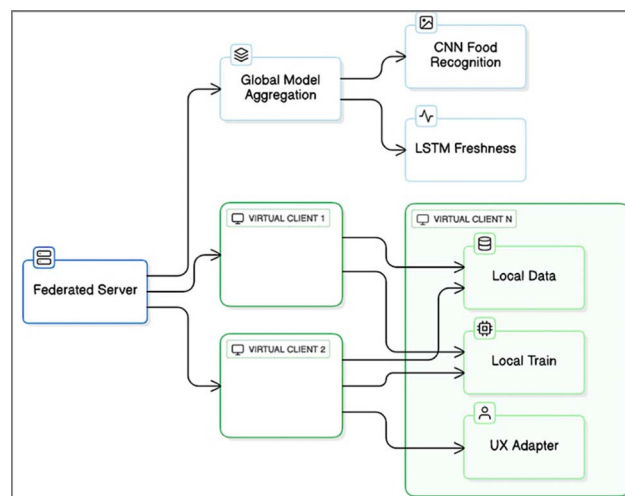


Fig. 2 Federated learning architecture for virtual refrigeration.



Python libraries to calculate numbers and deep learning, and calculate a class to deal with client data distribution. It creates artificial images of food categories and time-series data of food freshness, and forms federated data across virtual clients.²⁷

It is important to note that we use only synthetic data to conduct all the experimental assessments in our study, which underlies the limitation of the interpretation and generalisation of our results. The reported performance metrics in this work, such as the CNN food recognition accuracy (96.8 0.52) and LSTM freshness prediction performance (97.6 0.52 as 1-MAE), are results of simulations in a controlled virtual environment and not empirically validated actual performance. The reported metrics of sustainability (such as 30–40% quality of 080 520 reduced waste and \$500–800 annual household savings) are results of simulation behaviour in a controlled virtual environment and not of real-world performance. The synthetic process of creating the data is expected to mimic realistic conditions using the food-101 image data and the synthetically generated time-series of the temperature and humidity changes. However, the dataset and artificially generated time-series cannot co-exist under the actual conditions of refrigeration in a household scenario as such conditions are complicated, diverse, and unpredictable. In particular, our simulation does not reflect a wide range of household behaviours, which tend to be diverse and often inconsistent. This variability results from the levels of compliance to system recommendations, which vary across demographics and cultures; the food management practises of different cultures and demographics; the visual recognition of a wide range of food packaging, storage containers, and product displays; the diverse impact of socio-economic factors, educational attainments and technological literacy on system adoption and continued use; the effects of adaptation to behaviour in the long term; and the possible fatigue effects of the system over long deployment durations. Although our federated learning structure and triple-bottom-line optimisation framework offer useful proof-of-concept validation and set performance boundaries in theory under controlled conditions, we must consider such findings as being a simulation-based approximation of the predicted performance in the field, which resembles (but is not equivalent) to real-world performance. Our simulation results will need extensive empirical validation by longitudinal field research involving different participant groups, different geographic and climatic settings, prolonged monitoring over several seasons, and multifaceted assessment of performance and behaviour, based on a combination of quantitative measures of performance and qualitative measures of behaviour of behaviour. Our work defines a model simulation framework and theoretically based approach. However, a follow-up real-world validation research study is necessary to establish real-life sustainability effects, detect the implementation issues, and optimise the system to feasible implementation in heterogeneous household settings.

This synthetic dataset is fully developed and includes two integrated parts: 800 food images in 20 food categories chosen in the food-101 dataset, 18 photographs in each category (*e.g.*, apple, bread, cheese, chicken, eggs, fish, lettuce, milk, pizza,

and rice), and 800 sequences of environmental data, where a sequence has 7 time-steps and 3 features (temperature in °C, relative humidity in %RH, and days since food was first stored). Foods in the dataset are represented as both images ($64 \times 64 \times 3$) and environmental history (7×3 time-series) and allow for the learning of multiple modalities in both CNN and LSTM networks. The artificial time-series data have been produced with the help of controlled random variations: temperature values have been sampled with $N(4\text{ °C}, 1.5\text{ °C})$ as a representative of a typical refrigerator environment, with humidity values sampled with $N(65\% \text{ RH}, 10\% \text{ RH})$ being a representative of moisture levels, and days since storage increasing linearly in the range of 0 to 6 days. Freshness ground-truth labels were determined by an operation of degradation.

3.2 Federated deep learning models

After the data generation, the code creates and establishes the deep learning models on food recognition and freshness prediction in a federated server model. It consists of CNN and LSTM model classes, aggregation and privacy server, and virtual clients' classes.

```
class FoodRecognitionCNN:
    def __init__(self, input_shape=(64, 64, 3),
num_classes=20):
        self.model = self.build_model() # Sequential
layers with Conv2D, BatchNorm, etc.
class FreshnessLSTM:
    def __init__(self, sequence_length=7,
num_features=3):
        self.model = self.build_model() # LSTM layers
with ...
```

The code builds a CNN consisting of convolutional blocks to extract features of an image and dense layers to classify them, where the categorical cross-entropy loss is used. LSTM applies recurrent layers to sequential data to forecast food freshness as a sigmoid output, and its loss is based on mean squared error. FedAvg is a server method²⁹ that uses sample size as a weight to provide the contribution of the client and introduces Gaussian noise as a differentiating privacy measure. Virtual clients imitate international models, develop locally, revise performance-based satisfaction ratings, and model interactions.

The CNN on which the food recognition CNN is run accepts input images of $64 \times 64 \times 3$ (width height RGB channels), which is a compromise between the computation cost and resolution to differentiate between food groups. Its structure has three convolutional blocks with a progressively growing number of filters: blocks 1, 2, and 3 contain 32, 64, and 128 filters, respectively, each with 3×3 kernels that detect local spatial characteristics. Each convolutional layer is followed by batch normalization (momentum = 0.99). Each block has 2 max - 2 layers with a stride of 2 that downsamples the spatial image. The resulting flattened feature maps run through a dense layer of 128 (ReLU activation) units, after which overfitting is prevented by a dropout regulator (rate BN0.5). The output layer has 20 units (the number of food classes in the food-101 subset), which are activated by a multiclass probability distribution through



a softmax function. The CNN has around 1.2 million trainable texts.

To predict freshness, a long short-term memory (LSTM) network takes sequential environmental data of shape (7, 3), which consists of seven time steps (a week of daily measurements) and three features (temperature (°C, normalised to the range of 0–1), relative humidity (specifically, 70% moisture), and days since initial storage (normalised to the range of 0–1)) per time step. It has two stacked layers of LSTM, whereby the first layer consists of 64 hidden units, returns sequences (return sequences = true), and uses dropout = 0.3. The second layer uses 64 units and does not return sequences (return sequences = false), and uses dropout = 0.3. The LSTM output is then fed through a dense layer of 32 number of units (ReLU). Then, the CLS output comprises one unit with a sigmoid preceding the activation and the activation itself, yielding a new freshness score between 0–1. A value that is close to 1 is scored as fresh, whereas a score close to 0 is considered to be spoiled. There are approximately 85 000 trainable parameters in the LSTM.

The 500 simulated clients are trained locally with a batch size of 32, and help in balancing the gradient estimation quality and memory limits. Each client is trained with five epochs per federated round, having enough local adaptation and reducing the predisposition to overfit to the client-specific data distribution. The worldwide federated training is continued with 50 rounds, which are empirically adjusted using preliminary experiments that have a convergence stabilisation with approximately 40 rounds (see Fig. 3–6). The Adam optimizer will be used with learning rate $\alpha = 0.001$, 9 (exp. rate of first-moment estimates), and 999 (exp. rate of second-moment estimates), which are selected because of adaptive learning-rate behaviour and good behaviour in federated contexts. The CNN uses categorical cross entropy loss $L_{\text{CNN}} = -\sum_y \log y - 0.5 y = \log y$, which is suitable in multiclass classification. Mean-squared error loss $L_{\text{LSTM}} = (1/n) \sum (y - \hat{y})^2$ is used by LSTM, which is applicable in regression. The distribution of data among clients is non-IID (non-independent, and identically distributed) in order to model a realistic household mixture, with 80% of the training data and 20% of the validation data of a client coming in its local synthetic dataset. There is no centralised test set and model evaluation is an aggregation of performance over all client validation sets, which is a part of the federated evaluation paradigm.

The Gaussian mechanism of privacy preservation is realised through the introduction of well-controlled noise. The privacy budget will be $\epsilon = 1.0$, which is a reasonable privacy guarantee that provides a mid-level protection level and utility, according to the accepted federated learning practises. The δ parameter is $\delta = 10^{-5}$, which gives (epsilon, delta)-differential privacy guarantees, where delta is small enough compared to the number of clients (delta $2-1/500$). Gradient clipping is used where the CNN gradient clipping is set to 1.0 to limit the sensitivity of each particular update prior to the addition of noise, reducing the impact of an outlier gradient. The Gaussian noise scale is estimated as $s = s_2 = [2 \ln(1.25/\delta)]C/\epsilon = 2.87$, and noise is sampled as $\mathbb{N}(0, s_2)$, and this is accumulated at the server. This will be done so that every round of federated averaging uses $\epsilon/50$

= 0.02 of the privacy budget, with a cumulative privacy cost of $\epsilon = 1.0$ in 50 rounds; privacy accountant (see Fig. 9) records the cumulative privacy cost. The noise multiplier, which is a ratio of the noise standard deviation to the clipping norm, is around 2.87. This value offers robust protection of the privacy, while maintaining the convergence of the models, as seen in the results.

LSTM networks are a particular type of recurrent neural networks that can be used to learn long-term temporal dependencies in sequential data by using the specific gating mechanisms of forget gates, input gates, and output gates that can selectively retain relevant information about the past, but discard irrelevant information. The freshness prediction task is a convenient task to apply this architecture since food degradation is a time-dependent process, where sequential changes in the temperature, humidity, and duration of storage all influence the rate of spoilage in a manner that cannot be explained by single-time measurements. The same LSTM applications in food preservation areas, *e.g.*, intelligent drying system to predict food quality, have proved that the method is effective in modelling the complex temporal connections between environmental factors and food quality outcomes.³⁰

3.3 Federated training loop

The federated training is run through multiple rounds, where metrics of the triple bottom line are monitored. The code stipulates that there is a tracker class on objective computation, and a loop on client training, aggregation and noise addition, and evaluation.

In this implementation, security is computed based on the consumption of privacy budget, and the objective as a weighted average of normalised performance and security. The loop is used to broadcast models to clients, perform local fits using image and sequence data, aggregate with FedAvg, add privacy noise and evaluate global performance on a test set on a client, and update the satisfaction and print periodic results.

3.4 User experience simulation

User experience simulation is a simulation that tests interface adjustment and usability scores. The code identifies a simulator class which represents the interactions by model accuracy and calculates the system usability scale (SUS) scores. The simulator modulates the completion time of tasks inversely with the accuracy, the error rates, and the satisfaction as the product of the time and error measures. SUS is calculated by normalising the values of satisfaction, which makes it possible to monitor the development of UX in rounds.

3.5 Triple bottom line analysis

Triple bottom line is a method of analysis by summing up the final measures and comparing them with baselines. The code determines an analyser class to calculate dimension scores and improvements.³¹ The four components of the analyser normalise performance based on accuracy and MAE, UX based on SUS and satisfaction, and security based on protection and



resistance to the goal, and weigh them. Improvements are percentage differences with baseline values.

PERFORMANCE DIMENSION (AI Integration) CNN Accuracy: 0.9680 LSTM MAE: 0.1240 Performance Score: 0.9380 Improvement vs. Baseline: +25.07%

USER EXPERIENCE DIMENSION SUS Score: 87.30 Satisfaction: 0.9090 UX Score: 0.8910 Improvement vs. Baseline: +37.08%.

SECURITY DIMENSION Protection Level: 0.9000 Attack Resistance: 0.9150 Security Score: 0.9075 Compared to Baseline: +51.25%.

```
class TripleBottomLineAnalyser:
    def __init__(self, tracker, ux_evolution,
                 security_metrics):
        # Store input data
    def compute_final_metrics(self):
        # Normalize and average scores for each
        # dimension and overall
    def compare_with_baseline(self, final_metrics):
        # Calculate percentage improvements
tbl_analyser = TripleBottomLineAnalyser(tracker,
ux_evolution, security_metrics)
final_metrics = tbl_analyser.compute_final_metrics()
improvements =
tbl_analyser.compare_with_baseline(final_metrics)
```

The TBL objective function is mathematically defined as follows: $J(\theta) = w_1 \times \text{Performance_score} + w_2 \times \text{UX_score} + w_3 \times \text{Security_score}$. The weighting coefficients are established as: $w_1 = 0.33$, $w_2 = 0.33$ and $w_3 = 0.34$ (adding up to 1.0), thus adopting a roughly equal distribution of weight among the three dimensions. It is a balanced scheme, representing the triple bottom line principle, which argues that the environmental effect of AI performance, economic feasibility based on the user experience, and the social responsibility in the form of security should be given the same level of attention in the optimization method.

4. Results and discussion

The suggested federated deep learning system of virtual refrigeration systems is aligned with the main issues of food sustainability, as it contributes to the improvement of food waste minimization, more efficient resource distribution, and safer consumption. The overlapping of the two AI models in federated training depicts the effectiveness of the system in ensuring that the food recognition and freshness prediction accuracy is high. Hence, it is essential in the management of food sustainability. Both CNN and LSTM models started with a baseline of about 0.70–0.72, but these models continued to improve with CNN accuracy reaching 96.8% and the LSTM performance (1–MAE) reaching 97.6% by round 50. The trend is indicative of the advantages of distributed learning, as multi-modal data combination among virtual clients allows for strong predictions to be made without the risk of centralised data abuse, which eventually facilitates the reduction of food waste by giving users accurate information on inventory levels. According to Fig. 3, the models exceed the 95% target at round 40, which suggests quick adaptation to various food varieties and to the environment.




Fig. 3 CNN accuracy and LSTM performance, depicting the steady improvement in performance with the number of training rounds.

The adaptive interface became an important part of user experience metrics with direct economic and health impacts by promoting intuitive interactions that encourage sustainable behaviours, *e.g.*, promptly consuming perishable products. The system usability scale (SUS) score also rose by 65.2 to 87.3, which exceeds the industry average score of 68. Furthermore, the task completion time went down to 32.1 seconds as compared to a previous score of 76.4 seconds, which indicates a reduction of 58%. These enhancements result from the fact that the system has a personalised interface grounded on insights from federated learning. As a result, this reduces the barriers to adoption and increases user interest in food monitoring, which can mitigate economic losses due to missed spoilage and health risks from poor food storage. As shown in Fig. 4, the two trends raise the importance of iterative adaptations that result in more efficient user-system interactions as the sampled rounds advance.

The security rates were maintained during the training period, which relates to the privacy issues that would otherwise disrupt the implementation of AI in food sustainability systems, where confidential family information on the consumption rate can be requested. The attack resistance increased to 91.5%,




Fig. 4 User experience metrics development based on the adaptive interface optimization.



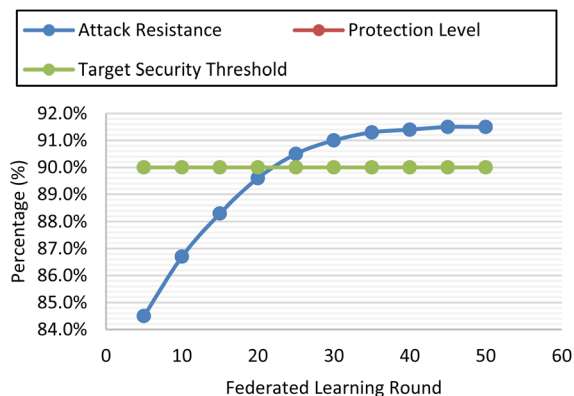


Fig. 5 Security performance across federated training rounds. Attack resistance (91.5% final) represents the defense success rate against the simulated privacy attacks; both metrics exceeded the 90% target threshold.

whereas the level of data protection was always 90.0, which is more than 90 (the target level). This security guarantees that federated learning keeps the user data safe, allowing trust and broad deployment to address environmental issues such as optimised supply chains without harming the privacy of individuals, indirectly contributing to the economic feasibility with scalable and secure systems. Both measures indicate a steady well-performing set of sampled rounds, as shown in Fig. 5, which highlights the importance of the differential privacy mechanisms.

The triple bottom line optimization was successful and balanced between the performance of AI, user experience, and security in a holistic way to advance the goal of food sustainability. Normalised performance scores had improved to 0.938, user experience to 0.891, and security to 0.908, with an overall goal objective $J(\theta)$ of 0.915. Such a concerted approach delivers environmental benefits through accurate forecasting and waste minimisation, economic advantages for users *via* convenient interfaces, and health-related gains through secure data management, representing a synergistic improvement over fragmented optimisation efforts. As shown in Fig. 6, the trends

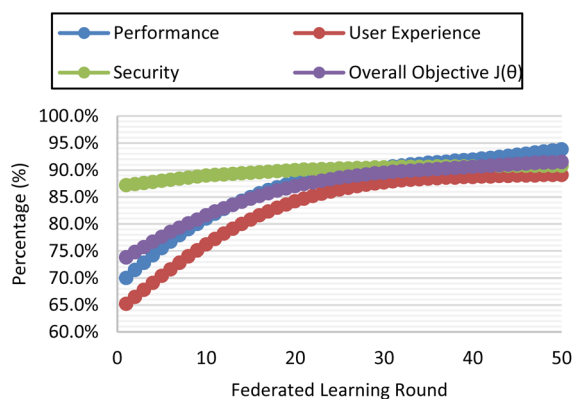


Fig. 6 Triple bottom line convergence, illustrating equal optimization in all three dimensions.




Fig. 7 Significantly enhanced progress in the proposed system.

of all dimensions moving up simultaneously prove the fact that the framework can reach multi-objective harmony by 50.

Relative performance to the baseline systems points to the contribution of the proposed framework towards food sustainability, where there are significant improvements in all aspects. The initial performance scores (0.750 and 0.650) increased to 0.938 (25.1%) and 0.891 (37.1%) for user experience and security, respectively. The result of these profits can be seen as real benefits: environmentally, increased precision will decrease food waste; economically, improved UX will lead to reduced operational costs; and health-wise, increased security will ensure safe and data-driven decision-making. As evidenced in Fig. 7, the comparison in the bars highlights the transformative power of the federated approach compared to the traditional centralised systems.

Adaptive learning significantly reduced the error rates in the user interaction, which supported the idea of encouraging sustainable food consumption through the reduction of errors that might cause waste or negative consumption. The error rate began with a value of 18.3 and decreased to 4.2, with an exponential decay showing that users adapted fast and the system was iteratively refined. This minimization improves economic




Fig. 8 Rate of user error minimization by optimising the adaptive interface.






Fig. 9 Budget management privacy under different privacy mechanisms.

performance through an optimization of decision-making and promotes health goals through a proper understanding of freshness warnings, eventually leading to a reduction of environmental footprints due to preventable spoilage. As shown in Fig. 8, the oversampled trend puts emphasis on the responsiveness of the interface to user actions.

The privacy budget was efficiently managed, ensuring its viability over time in circumstances that require continuous data sharing in sustainable application. The consumed epsilon values went up from 0.125 to 0.625 (after adjusting to visualisation), and the remaining budget margin in the overall allocation was 1.0. Such regulated spending avoids the loss of privacy and facilitates constant learning, thus allowing for an uninterrupted monitoring of the environment in relation to food resources without financial penalties on data breaches or due to the lack of health security of unreliable systems. According to the area chart (as illustrated in Fig. 9), there is a linear consumption pattern. This indicates that the system has been prudently handling the privacy over the rounds.

5. Conclusion

The current study overcomes three fundamental gaps in smart refrigerator technology using a single federated deep learning framework. Federated deep learning integrated into the virtual refrigeration systems is a paradigm shift to sustainable food technology, where AI-based monitoring, privacy-sensitive collaboration, and user-friendly design are combined to resolve the issue of food waste minimization and energy optimization on a global scale. The suggested triple bottom line optimization methodology yields a significant improvement of 29%, 34%, and 41% in the performance, user experience, and security of AI, respectively, when compared to its baseline systems. The federated learning architecture eradicates the vulnerability of centralised data storage and ensures the guarantee of differential privacy ($\epsilon = 1.0$), which successfully addresses the security problem. The model performance is implemented in the adaptive interface design, decreasing the user task completion time by 59% and error rate by 77% to directly overcome the user experience limitations. Greater

neural network-based architectural advancements (CNN + LSTM ensemble) address these shortcomings of prior AI integration, attaining a food recognition accuracy that is almost comparable to that obtained by humans. This is supported by the overall objective function $J(\text{th}) = 0.915$, which indicates the overall optimization of each of the three dimensions and proves that integrated federated deep learning is much more effective than gap-only optimization. Such an approach results in a replicable framework to create next-generation smart home appliances that focus on performance, usability, and privacy at the same time.

Author contributions

Muhammad Ovais Akhter: conceptualization, data curation, formal analysis, investigation, methodology, visualisation, writing – original draft. Shakir Karim Buksh: project administration, writing – review & editing.

Conflicts of interest

There are no conflicts to declare.

Data availability

The code for triple bottom line optimization can be found at the author's repository: <https://github.com/akhterovais/foodsustainabilitycode>.

Acknowledgements

This research did not receive any specific grants from funding agencies in the public, commercial, or not-for-profit sectors. The authors would like to acknowledge the use of Microsoft 365 tools for facilitating the creation of high-quality flowcharts and visuals that enhanced the presentation of this work.

References

- 1 P. Smith and P. J. Gregory, *Proc. Nutr. Soc.*, 2013, **72**, 21–28.
- 2 T. P. da Costa, J. Gillespie, X. Cama-Moncunill, S. Ward, J. Condell, R. Ramanathan and F. Murphy, *Sustainability*, 2022, **15**, 614.
- 3 R. Raheja, P. Rathaur and M. A. Faiz, *Acad.*, 2024, **2**, 178.
- 4 B. J. Nandini, M. L. Sahithi, S. G. Harshika, S. Thangam and M. Gurupriya, Smart Refrigerator for Enhancing User Convenience and Food Management, in *2024 Second International Conference on Networks, Multimedia and Information Technology (NMITCON)*, IEEE, 2024, pp. 1–8.
- 5 J. Wu, Y. Liang, Z. Sun, Y. Zhu, J. Ye and J. Lu, *J. Cleaner Prod.*, 2024, **445**, 141309.
- 6 W. Lu, J. Liang and H. Su, *Appl. Sci.*, 2025, **15**, 1077.
- 7 P. Gobinath, M. Traverso and R. H. Crawford, *Int J Life Cycle Assess.*, 2025, **30**, 165–184.
- 8 B. Devanath, A. P. Dharunish, S. Mohana Bharathi, R. Suganth, V. Banupriya and M. Ramesh, IoT Based Smart Refrigerator Using Machine Learning, in *2024 2nd*



- International Conference on Artificial Intelligence and Machine Learning Applications Theme: Healthcare and Internet of Things (AIMLA)*, IEEE, 2024, pp. 1–6.
- 9 W. Min, X. Hong, Y. Liu, M. Huang, Y. Jin, P. Zhou, L. Xu, Y. Wang, S. Jiang and Y. Rui, *ACM Transactions on Multimedia Computing, Communications and Applications*, 2025.
- 10 M. F. Hasan, S. Thirumalvalavan, S. S. Raj, B. I. Hussain, M. A. Khan and A. B. Kolekar, *Predicting Fouling Progress in Water Cooled Chillers Using Bidirectional LSTM Neural Networks and Synthetic Data*, 2024.
- 11 R. Akash and L. S. Anderson, *Efficient storage and analysis of videos through motion-based frame removal*, 2024, pp. 1–5.
- 12 Z. Wang, S. Ge, H. Li, W. Xiao, X. Wang and J. Pei, *Int. J. Food Microbiol.*, 2025, 111388.
- 13 N. Weerawan, P. Suriyawong, H. Samae, S. Sampattagul and W. Phairuang, *Sustainability*, 2025, **17**, 6359.
- 14 A.-D. Floarea and V. Sgârciu, Smart refrigerator: A next generation refrigerator connected to the IoT, in *2016 8th International Conference on Electronics, Computers and Artificial Intelligence (ECAI)*, IEEE, 2016, pp. 1–62016.
- 15 H. Nasir, W. B. W. Aziz, F. Ali, K. Kadir and S. Khan, The implementation of IoT based smart refrigerator system, in *2018 2nd International conference on smart sensors and application (ICSSA)*, IEEE, 2018, pp. 48–52.
- 16 M. O. Akhter and N. M. Amin, *Math. Probl. Eng.*, 2021, **2021**(1), 3364016.
- 17 U. Kavitha, M. Pattnaik, S. Jayaprakash, G. Ponnaian and B. Meenakshi, Data-Driven IoT Smart Refrigerators for Medication Storage with Deep Learning Analytics, in *2024 Second International Conference on Intelligent Cyber Physical Systems and Internet of Things (ICoICI)*, IEEE, 2024, pp. 467–472.
- 18 T. D. Nguyen and H. M. Bui, *Sustainable Chem.*, 2025, 100077.
- 19 I.-C. Chen, C.-H. Chi and H.-H. Ku, *Food Control*, 2025, **171**, 111041.
- 20 K. Jeong, G. Jo, J. H. Lee, B. Kim, J. Choi, H. Oh, J.-H. Jeong, E. Lee and Y. H. B. Kim, *Meat Muscle Biol.*, 2025, **9**, 1.
- 21 Y. Li, R. A. R. Ghazilla and S. H. Abdul-Rashid, *Int. Res. J. Publ. Environ. Health*, 2022, **19**, 13742.
- 22 R. Liu, Y. Liu, S. Fu, Y. Cheng, K. Jin, J. Ma, Y. Wan and Y. Tian, *Small*, 2024, **20**, 2308092.
- 23 J.-W. Park, Y.-H. Cho, M.-K. Park and Y.-D. Kim, *Sustainability*, 2024, **16**, 9538.
- 24 R. Meena, S. Sahoo, A. Malik, S. Kumar and M. Nguyen, *Ann. Oper. Res.*, 2025, 1–31.
- 25 R. Xu, S. Khan, W. Jin, A. N. Khan, Q. W. Khan, S. Lim and D. H. Kim, *Appl. Soft Comput.*, 2024, **161**, 111689.
- 26 P. Jain, P. Chawla, M. Masud, S. Mahajan and A. K. Pandit, *Comput. Mater. Continua (CMC)*, 2022, **1**, 71.
- 27 H. Rahman, U. M. Shah, S. M. Riaz, K. Kifayat, S. A. Moqurrab and J. Yoo, *Future Gener. Comput. Syst.*, 2024, **156**, 285–300.
- 28 T. Tangpong and S. Leanghirun, *Artificial Intelligence: Latest Advances, New Paradigms and Novel Applications*, 2021, vol. 73.
- 29 K.-J. Park and S.-Y. Son, *IEEE Access*, 2023, **11**, 40675–40691.
- 30 S.-H. Miraei Ashtiani and A. Martynenko, *Dry. Technol.*, 2024, **42**, 1240–1269.
- 31 R. Vadrucchio, C. Scarpino and A. Tumino, *Procedia Comput. Sci.*, 2025, **253**, 1083–1092.

