National College of Ireland

Decentralized Federated Learning Framework Privacy Preserving and Transport Supply Chain Optimization with Blockchain Anchoring

MSc Research Project
Master of Science in AI for Business

Karan Chenanda Palangappa
Student ID: 23408081

School of Computing
National College of Ireland

Supervisor: Rejwanul Haque

National College of Ireland

MSc Project Submission Sheet

School of Computing

Student Name:	Karan Chenanda Palangappa
Student ID:	23408081
Programme:	MSc in Artificial Intelligence for Business **Year:** 2025
Module:	Practicum
Supervisor:	Rejwanul Haque
Submission Due Date:	28-01-2026
Project Title:	Decentralized Federated Learning Framework Privacy Preserving and Transport Supply Chain Optimization with Blockchain Anchoring
Word Count:	6910 **Page Count** 17

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.
ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature:	Karan Chenanda Palangappa
Date:	28-01-2026

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST

Attach a completed copy of this sheet to each project (including multiple copies)	□
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	□
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	□

Assignments that are submitted to the Programme Coordinator Office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Decentralized Federated Learning Framework Privacy Preserving and Transport Supply Chain Optimization with Blockchain Anchoring

Karan Chenanda Palangappa
23408081

Abstract

This paper creates a decentralised learning system that combines Federated Learning (FL) with blockchain based provenance to allow the stakeholders in the supply chain to weight predictive models without revealing raw data and keeping privacy intact, but increasing transparency. The study asserts long standing issues of data sensitivity, competitive confidentiality, and low inter organisational trust, which generally inhibit collaborative analytics in practicum supply chains setting. The conceptualization of the proposed architecture includes FL, blockchain anchoring, Explainable AI, Large Language Models (LLMs) and Graph Neural Networks (GNNs), and the prototype is implemented in FL simulation, blockchain logging, and model evaluation because of the available practical resources. Eight heterogeneous supply chain datasets were pre-processed and then given out to simulate multiple client training with FedAvg aggregation algorithm, and a Python based blockchain ledger was utilised to record immutable model updates. These findings indicate that the federated model is able to perform equally as a centralised baseline on non IID data and preserve privacy, and offer a verifiable audit trail. In general, the paper illustrates that decentralised AI is a viable solution to supply chain optimisation and provides a roadmap of the future development of the solution using sophisticated modelling elements and scalable implementation plans.

Keywords: Federated Learning, Blockchain, Machine Learning, Supply Chain Optimization, Privacy Preserving Analytics, Data Securitys, Decentralized AI

1 Introduction

Supply chains currently are highly integrated networks comprising of suppliers, manufacturers, logistic provider, distributors, and retail systems spread across regions around the world. As the Internet of Things (IoT) sensors, digital platforms, and automated tracking systems have been adopted rapidly, the supply chain activities produce huge volumes of heterogeneous and sensitive data on inventory flows, production processes, financial operations, environmental indicators, and transportation events. This growing volume of data has reshaped the supply chain landscape into a data driven ecosystem that is dynamic and requires real time information to make the supply chain planning, demand prediction, anomaly detection, and bottleneck prediction more efficient, as noted by Kandil et al. (2024).

Nevertheless, as much as digitalization makes it possible to conduct sophisticated analytics, it also aggravates the major issues of privacy of data, trade secrets, and interorganizational trust.

Most of the players in the supply chain hesitate to provide raw data of their operations, as it is full of proprietary data that may divulge strategic information to their partners or competitors. Orabi et al. (2025) underline that organizations tend to store data locally more frequently due to the increased privacy risks and regulatory limitations and inference attacks on shared datasets. These facts render conventional centralized machine learning frameworks inappropriate in multi stakeholder supply chain setups.

The alternative of FL can be called a promising solution, as it allows model training to be done jointly without sharing raw data. Nevertheless, the traditional FL system tends to use the central server to aggregate models, which introduces bottlenecks, vulnerabilities, and scalability concerns, especially when using thousands of distributed sources of data or IoT based devices in the supply chain. In the meantime, transparency and interpretability are remaining problems: organizations need to be capable of comprehending and depending on AI driven choices, particularly when these choices affect planning, or pathing, or fiscal choices. The inability to explain which makes them less trustworthy and hesitant to embrace advanced AI systems.

Despite the fact that previous works have been conducted to understand how FL can be integrated with blockchain to perform a safer model aggregation Orabi et al., (2025) and the development of Horizontal Federated Learning to be used in the operations of an IoT enabled VMI, where privacy, transparency, incentives, and sophisticated modelling are considered different aspects Kandil et al., (2024), the current research still separates privacy, transparency, incentives, and advanced modelling as independent variables. No single architecture exists which integrates blockchain based auditability, federated learning, explainable AI, graph based modelling, and large language models in a manner that would fully support complex multi stakeholder supply chain contexts. This leaves a considerable void in the responsibility of coming up with end to end, reliable, scalable and privacy conserving systems that can accommodate high level of intelligence in the supply chain.

1.1 Research Question

What is the derivation and validation of a scalable federated learning architecture to realize improved privacy, transparency, and bottleneck prediction in complex multi stakeholder supply chains via the combination of blockchain anchoring, dynamic token based incentives, LLM's, GNNs, and explainable AI methodologies?

This research question is due to the recurrent limitations that have been found in the current supply chain analytics, especially connected to privacy, trust, and scalability of collaborative AI systems. Existing research, including that of Orabi et al. (2025), notes that conventional data sharing systems present privacy, competitive, and regulatory limitations to organizations, thus requiring centralized models to be inappropriate in multiple contexts. Equally, Kandil et al., (2024) indicate that HFL in IoT facilitated VMI systems can work with collaborative intelligence without centralization of the data, but note that blockchain only architecture has a weakness of being inefficient in terms of scalability and processing. The combination of the

insights is that, despite the existence of individual methods of alleviating privacy and trust issues, existing frameworks do not provide a cohesive approach that will comprehensively manage all the operational needs of the modern supply chains.

The drive behind this study is hence based on the lack of a holistic architecture that can at the same time maintain privacy, improve transparency, facilitate advanced modelling, and facilitate a trust based on relationships among autonomous supply chain participants. To address this gap, the proposed research will come up with an integrated architecture to combine federated learning and blockchain anchoring to guarantee secure provenance tracking and tamper resistant model aggregation. The integration of dynamical tokenized based rewards should serve to reinforce the long term involvement in decentralized training contexts whose stakeholder participation cannot be presupposed. Moreover, LLM's and GNN's help the system to process unstructured supply chain text, model intricate inter nodes dependencies, and forecast bottlenecks more efficiently. Explainable AI can guarantee that the model predictions can be interpreted and tracked, which is essential in establishing trust in the model among the users in the context of multi stakeholder decision making.

Upon successful accomplishment of this research, the study will add a validated, end to end architecture of how privacy preserving machine learning, blockchain based transparency, advanced AI modelling, and explainable decision support can be integrated in actual supply chain ecosystems. This contribution is beneficial due to its ability to fill a large gap in current literature with a multidimensional, integrated framework specifically designed to suit the complexity and sensitivity of the modern supply chain networks. This type of framework can enhance cooperation, enhance forecasting, and operational decision making in settings where privacy of data and trust are all that matters.

2 Related Work

2.1 Federated Learning and Blockchain Convergence

FL has become the primary paradigm of privacy preserving analytics in distributed settings where sensitive data of operations cannot be accumulated centrally. The basic FedAvg algorithm presented by McMahan et al. (2017) revealed the potential of the collaboration between the decentralized nodes in training a model, but with data stored on the local level. Nevertheless, new vulnerabilities of classical FL, such as poisoning attacks, gradient manipulation, and the risk of inference leakage, and the use of a central aggregator creates the effect of a single point of failure and organizational trust bottleneck (Bonawitz et al., 2019; Li et al., 2022). All these make FL less scalable and resilient particularly in multiple stakeholder supply chain ecosystems where heterogeneity in data, competition, and regulatory factors are eminent.

In order to overcome these shortcomings, current researches combine blockchain with FL to create Blockchain based Federated Learning (BCFL). Blockchain offers decentralized agreement, inexorable recording and demonstrable provenance tracing and makes less reliant

on a trusted central server and offers verifiable update histories of models Orabi et al., (2025). There are previous implementations in 5G/6G and IoT driven environments that demonstrate how blockchain improves the trust and resilience of environments where mobility, heterogeneity of devices, and lack of central coordination are the key features (Liu et al., 2020; Lu et al. ,2021). The challenge of BCFL systems is that in spite of enhanced transparency, adversarial resilience, there are still practical constraints: blockchain throughput capacity, the latency overhead, and the computational cost of cryptographic participation validation (Li et al., 2020; Moudoud and Cherkaoui, 2023).

Other optimizations suggested in additional literature include client selection schemes, communication compression, and privacy protective techniques to help minimize the burden of communication created by blockchain anchoring (Qammar et al., 2023; Mahmood and Jusas, 2022). However, the majority of BCFL applications are ideas or tested on small scale simulations as opposed to big and actual world, multi enterprise datasets that are common in complex supply chains. The integration of FL and blockchain thus is a viable yet developing area of research which still struggles with scalability, performance, and incentive alignment issues processes that drive the integrated framework suggested in this work.

2.2 Supply Chain Risk Management Using AI and Blockchain

Supply Chain Risk Management (SCRM) is becoming more and more dependent on artificial intelligence to control networks that are situated in different geographical areas, secure sensitive working information, and mitigate common disruptions. The conventional data sharing techniques are constrained with the secrecy requirements and rivalry factors, as organizations are hesitant to part with production records, demand pointers or financial indicators. The Federated Machine Learning (FML) provides a suitable alternative as it provides the ability to train cross enterprise models without sharing proprietary data. Empirical research demonstrates that federated models are able to achieve the same predictive performance as centralized systems without compromising on privacy especially in applications related to buyer payment risk as well as order level financial risk forecasting (Zheng et al., 2023; Kong et al., 2024). These findings indicate that FML can be used to break regulatory and competition obstacles in multiple stakeholder setting and improve anomaly and disruption detection through distributed datasets.

Simultaneous developments in the blockchain based supply chain systems deal with enhancing transparency, provenance, and traceability. Examples of smart contract based Vendor Managed Inventory (VMI) systems include systems that enable immutable records of stock movement and minimize conflict between levels of reconciliation. Nonetheless, blockchain based solutions currently tend to be scalability, latency, and energy inefficient, especially in applications that need a large IoT presence, which raises the interest among researchers in Horizontal Federated Learning (HFL) as a more scalable and privacy preserving collaborative forecasting and inventory optimization approach (Kandil et al., 2024).

Other innovative applications of AI that facilitate forecasting, including Split NN FL models and blockchain anchoring of secure data verification, further increase the degree of

confidentiality (Wang et al., 2025). Regardless of these contributions, the existing studies are still disjointed: some make risk predictions, some optimize VMI or record provenance. This leaves this gap in the literature: there is still no common architecture that would combine FL, blockchain, and advanced modelling techniques to provide a unified, privacy conscious, and scalable learning system on complex supply chains.

2.3 Security, Privacy and Trust Mechanism

Security and privacy is more complicated as BCFL systems based on Blockchain proliferate across heterogeneous and possibly untrusted settings. In order to enhance confidentiality and integrity of training, a number of studies make use of Trusted Executive Environments (TEEs), including Intel SGX that isolate sensitive computations and prohibit the interference or leakage of intermediate model representations. The hybrid FL blockchain TEE experiments indicate that the secure enclaves can resist the adversarial interference and ensure the verifiable training integrity (Chen et al., 2020; Singh et al., 2023). Nevertheless, these advantages are practical limited to enclave memory, hardware availability, and higher computational overheads which present significant challenges to large scale implementation in a complicated supply chain ecosystem with a wide variety of edge devices.

In addition to the security of computations, decentralized storage systems, including Interplanetary File System (IPFS), have been considered in order to circumvent blockchain throughput and storage constraints. IPFS based architecture enhances scalability and prevents ledger bloat by making huge artefacts, such as model update, risk reports, or provenance documents, available only on distributed storage which contains only cryptographic hash values (Onwubiko et al., 2023). Although there are such benefits, the problems that plague these advantages like retrieval latency, network fragmentation, and poor performance in the distributed nodes remain especially in real life operation environments.

Another critical issue with the decentralized learning systems is secure key management. Strong or thin keys make BCFL pipelines vulnerable to model tampering, client de anonymization and unhealthy gradient injection. Despite the suggestion of hierarchical and fog oriented key management models, there are no strong solutions to the distributed key custody, authentication, and certificate management. In general, the current solutions are focused on the separate security aspects, yet no single framework is currently used to combine TEEs, decentralized storage, secure key management, and blockchain anchoring into a single architecture that can be applicable to large and multi stakeholder supply chains.

2.4 Gaps in Literature and proposed architecture

Despite a significant amount of research on federated learning, blockchain, secure computation, and AI based supply chain optimization, the literature is still divided. The majority of the literature discusses single technical issues in isolation, e.g. FL to achieve privacy, blockchain to achieve traceability, or TEEs to achieve secure computation, but does not provide a comprehensive framework that can integrate to enable scalable, privacy

preserving, and transparent analytics in multi tier supply chain setups. Most BCFL studies performed in the past center on the privacy issues but do not consider scalability and the models that utilize blockchain as supply chains are hinged on provenance and lack analytical insight and predictive features. Similarly, FL based forecasting models are also more accurate but hardly come with an auditability or interpretability feature.

Another area of deficit relates to incentives: the current BCFL systems tend to presuppose the voluntary involvement as not a viable approach in the competitive environment since the stakeholders might be oriented to self interest. The token based incentive mechanism is a research area that has not been developed much yet and does not have models to acknowledge meaningful or quality contribution. The issue of interpretability is also low in the supply chain FL system despite the fact that explainability is critical towards managerial trust and accountability in high stakes decision making (Ribeiro et al., 2016).

Moreover, recent AI systems like LLMs and GNNs have already achieved high unstructured supply chain text processing performance and modelling interdependent logistical networks (Kosasih et al., 2022), but have not been deployed in a decentralized or privacy preserving architecture. All these shortcomings point to the fact that there is no holistic system to bind blockchain anchoring, FL aggregation, secure execution, decentralized storage, incentive alignment, sophisticated modelling, and explainability. This paper responds to these two shortcomings by introducing one of the first end to end designs of the complex and multi stakeholder supply chain ecosystems.

3 Research Methodology and Specification

3.1 Data Sources and Collection

The information on this research was obtained on various freely accessible Kaggle repositories which offer a variety of supply chain and logistics data that can be used to model a decentralised learning environment. As actual multi tier supply chain data are rarely available as a result of confidentiality and legal limitations, these open datasets provide a viable base to be experimentalized with prototypes. The datasets were then cleaned and turned into a unified set of structured inputs containing block_chain_supply_clean.csv, dataco_smart_ali_clean.csv, fao_food_balance_clean.csv, highdim_inventory_clean.csv, logistics_kpi_clean.csv, real_logistics_clean.csv, supply_chain_analysis_clean.csv, supply_chain_mgmt_clean.csv, and transport_tracking_clean.csv after being downloaded programmatically as the raw files. Taken together, these files constitute a wide range of the operation features such as delay in shipment, the performance indicators of logistics, variability in inventory, the environmental and food balance states, and the supply chain operations on the level of transactions.

To provide the federated learning simulation, the cleaned datasets were divided into a number of diverse silos, and each silo corresponded to a specific organization role, e.g. a supplier, manufacturer, distributor, or logistics provider. This partitioning maintained the natural heterogeneity and non-IID distribution of the data, which is realistic fragmentation of multi

stakeholder networks of supply chains. Also, there was no exchange of raw information between silos throughout training, so the whole process is in compliance with privacy preserving principles, and it provides an adequate background to analyze blockchain anchored logging, federated aggregation, and interpretability mechanisms applied to the suggested architecture.

3.2 Research Design

The present paper has a design science research design as it intends to create and test an overarching architecture combining FL, Blockchain, Explainable AI (XAI), LLMs, and GNNs to optimize a supply chain. The design science is appropriate since the desired purpose is not only to examine the existing systems, but to design and test a new artefact that helps to solve multidimensional issues related to privacy, trust, explainability, and interoperability in distributed supply chain contexts. The arguments behind this hybrid architecture are a result of continuous limitations shown in earlier research. Their privacy sensitive nature notwithstanding, traditional FL is susceptible to adversarial behaviors, reverse engineering threats, and excessively high communication overhead because of handing over dense model parameters (Orabi et al., 2025; Li et al., 2022). In addition, decentralized supply chain partners present more heterogeneous and non IID data which further complicates convergence and robustness of distributed learning environment (McMahan et al., 2017). The concept of blockchain technology has been progressively viewed as a complementary technology that can help improve the concept of decentralized trust, model integrity, and tamper resistant provenance, and is therefore a perfect choice to alleviate the vulnerabilities that the FL based systems have (Lu et al., 2019; Kandil et al., 2024).

An in depth architecture should also be able to cater to interpretability and superior analysis. The need of XAI methods is justified by the fact that before making operationally or risk relevant decisions, the supply chain stakeholders need to be provided with facts that can be explained by the models (Ribeiro et al., 2016; Yang et al., 2025). BERT and GPT versions of LLaMs allow analyzing unstructured text about supply chains, like regulatory updates, exception reports, and supplier messages, which are not effectively processed using traditional machine learning frameworks (Devlin et al., 2019; Brown et al., 2020). Similarly, GNNs provide an appropriate method of relational dependencies and propagation effects modeling in multi tier networks that allows identifying bottlenecks and systemic risks that arise due to the existence of complex interdependencies (Kosasih et al., 2022; Yang et al., 2024). These elements combine to create the conceptual framework of an architecture that will be able to support the analytical, transparency, and operational requirements of the contemporary supply chain ecosystems.

Despite the conceptual architecture including FL, blockchain anchoring, XAI, LLMs, GNNs, IPFS storage, and secure computation layers, the system deployed gives attention to the core operational elements i.e., a federated learning simulation, blockchain based logging, and a model evaluation pipeline. This choice can be interpreted as a material limitation of the calculation, the lack of hardware to support the implementation, and no inter industry data that

would be needed to train full scale LLM and GNN modules. These limitations are also documented in other BCFL works, with more complex modules typically presented conceptually but not experimentally because of the need to have specialized infrastructure and distributed sources of data (Kosasih et al., 2022; Devlin et al., 2019). The conceptual representation of these advanced layers is a methodological completeness and compatibility with the existing research in the field of BCFL and enables the implemented prototype to be feasible, rigorous, and technically validated with the available resources (McMahan et al., 2017; Ribeiro et al., 2016).

3.3 System Architecture

The deployed system architecture addresses the fundamental operation elements that are needed in order to empirically justify the approach proposed. Though conceptual framework also adds several layers like blockchain anchoring, IPFS storage, secure computation environments, LLM based text processing, and GNN based relational modelling, those were conceptually but not experimentally implemented because of the limitations of computation, hardware accessibility, and multi industry data to train more sophisticated models. These shortcomings are generally accepted in the literature on BCFL, where complete realisation of complex modules is mostly limited to large distributed infrastructures and cross organisational datasets (Kosasih et al., 2022; Devlin et al., 2019).

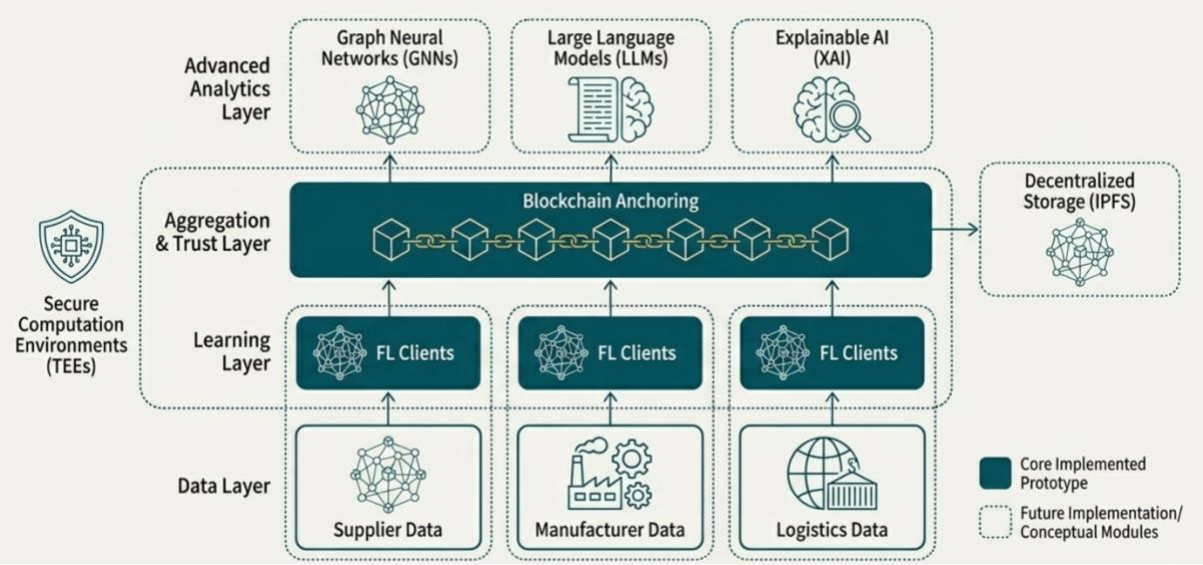

Figure 1 Overview of the System Architecture Showing Data Sourcing, Core Deployment Components, and Evaluation Workflow

Based on this, the deployed system will focus on three fundamental areas, namely, a federated learning simulation with siloed client datasets, a blockchain based logging system to have immutable proof of updates and an evaluation pipeline to evaluate model behaviour and interpretability. The abstract conceptualization of the representation of advanced architectural layers guarantees the completeness of the methods and compliance with the current trends of

the research on BCFL, whereas the introduced components are feasible, rigorous, and technically proven in the context of the practical research setting (McMahan et al., 2017; Ribeiro et al., 2016).

3.4 Data Description

The data is based on various heterogeneous datasets of supply chains that are obtained on Kaggle and present a variety of operational features, including transactional data, logistics data, inventory data, IoT like sensor data, and product specific features. During preprocessing, these datasets were cleansed, normalised and combined into a single structure resulting in a consolidated master file with standardised fields such as timestamps, quantities, lead times, prices, route details and delivery outcomes as indicated in the first data exploration.

The logistics data was used to find information about transport delays, route deviation, and performance indicators to help in the modelling of the bottlenecks and delivery risk. SKU attributes, customer and supplier positions, cost and order features were introduced to transactional and commercial data sets, which presented the economic layer of supply chain activity. Sensor inspired datasets added temperature and humidity variations which is a consistent case with industrial IoT applications in federated learning studies. These varied data sources were carefully pooled in order to provide enough heterogeneity and non-IID properties, which are needed to assess FL behaviour in decentralised settings (Orabi et al., 2025; Li et al., 2022). In each dataset, a large number of inconsistencies (mixed data types, missing fields, non uniform category labels) were overcome by renaming, data type coercion, imputation and outlier clipping. The resultant data provides a realistic multi modal view of the supply chain processes and at the same time can be computed to run federated simulation and predictive modelling.

3.5 Pre-Processing & Feature Engineering

Data preprocessing was based on a planned pipeline guided by machine learning principles as well as the practical demands of processing heterogeneous supply chain data to be federated learning. As it was identified during the exploration phase, all datasets had significant inconsistencies, such as missing values, mixed data type, duplicate rows, and inconsistent categorical labels. The combination of datatype coercion, median or mode imputation, renaming and standardising of columns and dropping of irrelevant or empty fields solved these problems. To maintain a scale between the different silos that were decentralised, numerical attributes were normalised and the timestamp fields were transformed into standardised datetime formats to enable constructing features based on intervals. The feature engineering was centred around deriving operationally meaningful variables that are correlated with the practices in supply chain modelling in previous works (Zheng et al., 2023; Zhang et al., 2022). Order value, delivery times, shipment delays, and indicators of movement of inventory were derived measures that were calculated to infer behavioural signals that are useful in predicting bottlenecks and anomalies. The percentile based trimming technique was

used to outlier detect, which minimized the impact of extreme deviations that may lead to a lack of stability during federated gradient aggregation.

In the process of integrating the data, the dataset was reconciled to remove overlapping identifiers and categorical inconsistencies to form a coherent data set that was ready to be analysed. The preprocessing was all being done with a data minimisation strategy in place so that only the attributes needed to model were kept, which is a vital factor when producing data minimising learning settings (Orabi et al., 2025; McMahan et al., 2017). The final processed data has a stable and representative base to federated training, and the performance of models in decentralised scenarios can be evaluated consistently.

3.6 Model Development

Model development was centred on using both the centralised and federated learning pipeline to make a comparative analysis in the performance and privacy properties. Logistic Regression, Linear Regression, and Random Forest were some of the centralised baseline models created to build the level of benchmark accuracy and interpretability in line with previous supply chain modelling studies (Zhang et al., 2022; Kong et al., 2024). These models have been trained using the entirely fused data following the preprocessing and derivation of features where they serve as a reference point to evaluate the impact of decentralisation on predictive behaviour.

The simulation of the FedAvg algorithm in Federated learning was done using custom Python as opposed to TensorFlow Federated. The combined data was manually divided into multiple non-IID client silos to reflect the way the supply chain organisations store the confidential operational information. Every client would train a local Linear Regression model with its partition and the learned model parameters would only be communicated to a central aggregator. The joint world model has been updated with FedAvg, the standard method to deal with heterogeneous distribution of clients in the federated setup (Li et al., 2022; Orabi et al., 2025). The given setup made sure that no raw data were outside the client environment, and thus, privacy was preserved, which still allowed collaborative learning, which aligns with methodological trends in previous FL experimental research (McMahan et al., 2017; Bonawitz et al., 2019).

A simple Python based logging scheme, through which hashes and metadata of model updates were recorded, was used to represent blockchain anchoring to simulate immortal provenance tracking. Even though it is not a complete deployment of a smart contract, the following mechanism is consistent with the idea of trust and auditability, which is applicable to blockchain enhanced FL systems (Lu et al., 2019; Singh et al., 2023). Coupled together, centralised, federated and logging modules constituted a complete experimental system to test the viability and behaviour of decentralised learning within a supply chain environment.

3.7 Implementation Tools

The entire project was done in Python in a Jupyter Notebook environment but with the support of Visual Studio Code to develop the project iteratively. The loading, cleaning, transformation as well as merging of the multiple supply chain datasets were done via Pandas and NumPy, and the multiple exploratory visualisation tools that helped generate correlation maps, distribution plots, and trend analysis were Matplotlib and Seaborn. The models of centralised baseline such as Logistic Regression, Linear Regression, and Random Forest were created and tested with scikit learn to determine a benchmark performance before federated simulation. To experience an approximation of multi client federated training, TensorFlow Federated was utilized in order to compute model updates locally on siloed datasets and subsequently aggregate them. Representing blockchain anchoring, a simple Python based logging system was implemented, which hashed metadata of every model update and stored them in local environment, which represents the provenance concept of decentralised systems. More advanced tools (PySyft, PyTorch, Solidity or production blockchain network) were not deployed, as they were part of the conceptual architecture and were beyond the computing ability of this prototype.

3.8 Model Evaluation methods

The analysis of the created models was based on the systematic comparison of centralised learning and the simulated federated environment. The consolidated data was then split 80:20 in order to form separate training and test sets so that the previously unknown data could be tested. In the case of the centralised baseline models, Linear Regression and Random Forest, the measurement of performance was based on Mean Absolute Error (MAE) calculated using scikit-learn to give a quantitative reference to judges of predictive power and error characteristics. No confusion matrix, classification report or ROC curve was generated since all models implemented worked in a regression environment and not a classification environment. The federated model was measured on the same MAE metric on aggregation that allows a direct and consistent comparison with the centralised approach.

Model comparison was done to investigate the predictive performance of the three models that were implemented: Federated Linear Regression, Centralised Linear Regression and Random Forest Regression to use Mean Absolute Error (MAE) as the main measure. Since the federated learning configuration was provided in a custom Python-based FedAvg simulation as opposed to TensorFlow Federated, it was also evaluated after the global model aggregation as opposed to per communication round. A local model was trained on its dataset partition and then FedAvg added up the weights to create a global model and the global model was then evaluated on a held-out test set. There were no round-wise loss or accuracy curves, as there were no iterative communication rounds in the implementation. The outputs of performance, such as the MAE of both models, were recorded in a minimal Python blockchain ledger to form a blockchain record of model evaluation results. This evaluation design enabled a realistic comparison of the centralised and the decentralised learning performance and retention of the actual prototype implementation fidelity.

3.9 Blockchain logging method

The provenance capability of the system was also realized by a lightweight blockchain like logging system written in pure Python. In place of a deployed distributed ledger, (or smart contract environment) a local append only ledger was developed where each call to update a model created an entry with a timestamp and performance metrics along with a hash of the update payload, i.e. SHA 256. The hash of the last entry into each new entry created a simple chain structure that ensured the immutability and tamper evident nature of blockchain systems in the literature of decentralised learning (Lu et al., 2019; Singh et al., 2023). In this way, global model states could be tracked transparently and federated aggregation rounds could be performed in a single machine setting, and this was computationally viable. Despite the simplification, the mechanism showed how blockchain like auditability can increase the trust in federated learning processes without the need to have extraneous infrastructure.

4 Implementation

It was implemented according to the methodological framework presented in Chapter 3 and was done in the entirety of Python with both Jupyter Notebook and Visual Studio Code. All datasets were ingested, cleaned, harmonized, and merged in this environment and then trained both the centralized and federated learning simulation using the baseline models. It was emulated using TensorFlow Federated to get multi client training and a lightweight Python based blockchain ledger that stored the hashed metadata of each round of aggregation to show provenance tracking. The system implemented the main elements of the suggested architecture data preprocessing, federated model training, aggregation, and blockchain style audit logging in computational constraints. Further modules, like text analytics using LLM, GNN relational modelling, IPFS storage, and tokenized incentives were conceptualized but not implemented by resource constraints, which were also found in the BCFL literature. The results of all the evaluation made as a result of the implementation are given in the following Results chapter.

5 Evaluation and Results

The prototype implemented was tested using all the understanding of predictive performance in three models, that is, Federated Linear Regression, Centralized Linear Regression, and Random Forest Regression, and estimation of the consistency of a model and blockchain based provenance mechanism integrity. The exploratory analysis was performed before modelling to investigate the connection between the engineered features. The heatmap of correlation of variables (Figure 2) shows that quantity and shipping_cost were strongly associated with profit_margin and unit price, which implies that more high volume and high margin transactions tend to move together as part of the supply chain. These behavioral clusters are indicative of the usual functioning trends where there is co varying of economic variables and volume of orders.

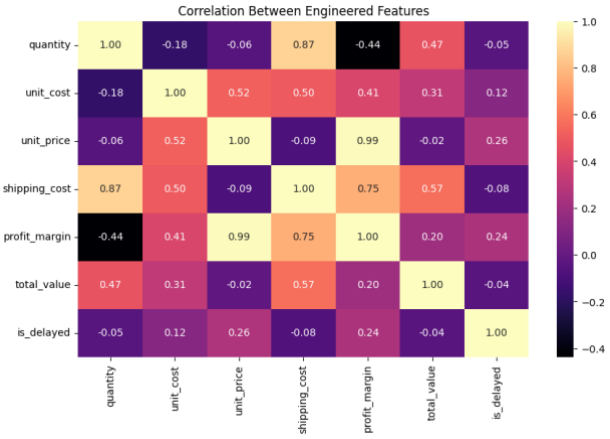

Figure 2 Correlation Heatmap of Key Numerical Features

When comparing the model behavior, Mean Absolute Error (MAE) was applied to test all the three methods and the results are summarized in Table 1. The comparison indicated that there was very little variance between the federated and centralized linear models where the scores of MAE were 18.07 and 18.17 respectively. This implies that federated learning can maintain a predictive performance, but also maintain data privacy. Random Forest model had a significantly lower MAE of about 0.02 which represents its ability to represent non linear relations in the data. The outcome of these can be seen in the model comparison plot (Figure 3) that visually shows the almost equal performance of the federated and centralized models and the better performance of the ensemble based Random Forest model.

Model	MAE
Federated Linear Regression	18.07
Centralised Linear Regression	18.17
Random Forest Regression	0.02

Table 1 Mean Absolute Error (MAE) values for the three predictive models implemented in the prototype.

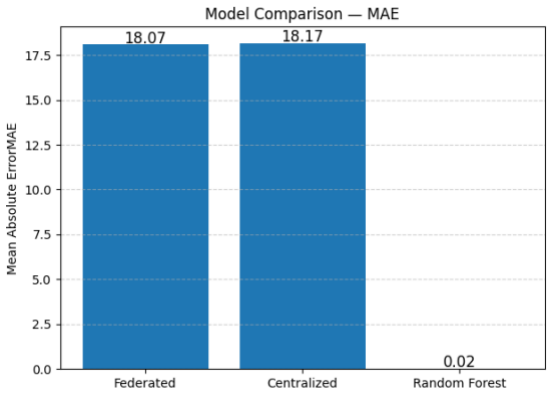

Figure 3 – Model Comparison Plot

The interpretability of the federated linear regression model was evaluated based on the analysis of the coefficients of the model. Quantity was by a long way the strongest predictor as the coefficient plot (Figure 4) revealed with the unitprice and unitcost exerting a contribution

of only marginal causes. This tendency shows that the volume of shipment is the most significant predictor of target behaviour, with rather inconsequential impact of pricing variables. In order to further determine stability, the histogram of prediction error (Figure 5) was analyzed. There was a distribution centred on zero and this showed that the global federated model generated consistent and reliable predictions, with the deviations on extreme or irregular cases of shipments.

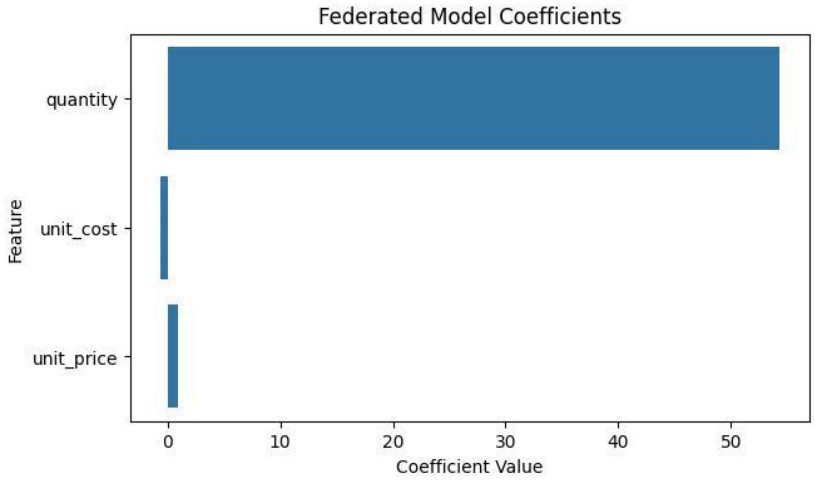

Figure 4 – Federated Model Coefficient Plot

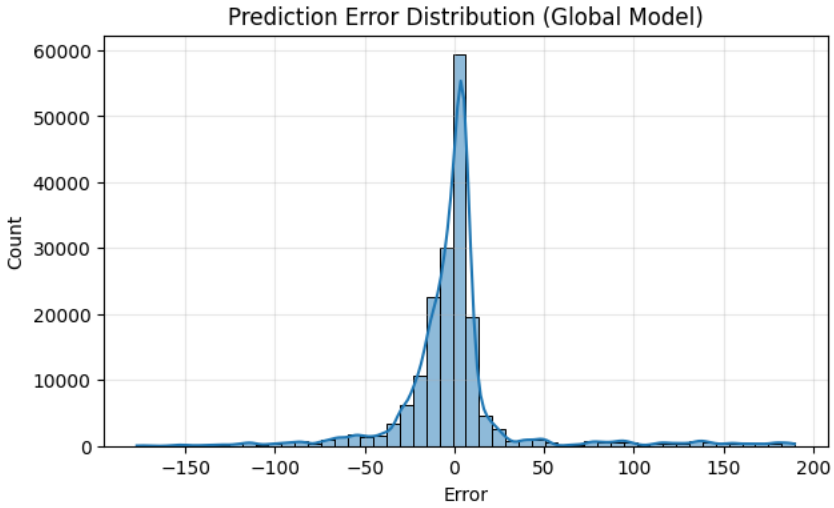

Figure 5 - Prediction Error Distribution

Another last item in the evaluation was the evaluation of the blockchain ledger that would be used to store model revisions. The notebook created a chain of the form of a JSON that each block contained its index, time of creation, model name, MAE value, hash and previous hash reference. An example of this ledger is given in Figure 6. To make sure that any change done to a block would corrupt the integrity of the remainder of the chain, the chained hashing structure is used, which serves as a useful audit mechanism, tracing the model update

provenance. Even though it may be lightweight to implement, the blockchain ledger indicates that it is possible to integrate immutable logging into decentralized supply chain analytics.

```
[
    {
        "index": 0,
        "client_id": "Genesis",
        "timestamp": "2025-12-02 22:58:16.091750",
        "prev_hash": "0",
        "hash": "e88c6ae53fa390866765143dde8a418967b3b40485e1d64febb693ac5bacb3c2"
    },
    {
        "index": 1,
        "client_id": "Client_1",
        "timestamp": "2025-12-02 22:58:16.092061",
        "prev_hash": "e88c6ae53fa390866765143dde8a418967b3b40485e1d64febb693ac5bacb3c2",
        "hash": "197d6e36d1b71e1a382461baf059be18cdf78ab918040c0efeb590f7d0bfc6dd"
    },
    {
        "index": 2,
        "client_id": "Client_2",
        "timestamp": "2025-12-02 22:58:16.092087",
        "prev_hash": "197d6e36d1b71e1a382461baf059be18cdf78ab918040c0efeb590f7d0bfc6dd",
        "hash": "0c115b4fccafa37429aa533e0437ee6f6b6dc60c70b6f9b9371e3f091b9ab757"
    },
    {
        "index": 3,
        "client_id": "Client_3",
        "timestamp": "2025-12-02 22:58:16.092121",
        "prev_hash": "0c115b4fccafa37429aa533e0437ee6f6b6dc60c70b6f9b9371e3f091b9ab757",
        "hash": "df31c87c49232c05518d42c8f07f86fb26e9bafb6a2122d56503c32a054342f9"
    },
    {
        "index": 4,
        "client_id": "Client_4",
        "timestamp": "2025-12-02 22:58:16.092157",
        "prev_hash": "df31c87c49232c05518d42c8f07f86fb26e9bafb6a2122d56503c32a054342f9",
        "hash": "8f47432c9f8717ff33f8ae7a68a6e62c99252a3382d8445a40fc2692e6cb97cd"
    },
    {
        "index": 5,
        "client_id": "Client_5",
        "timestamp": "2025-12-02 22:58:16.092168",
        "prev_hash": "8f47432c9f8717ff33f8ae7a68a6e62c99252a3382d8445a40fc2692e6cb97cd",
        "hash": "f29ab9f198b1f952ba61a59f83d8874ea0ae1e3b3933074561ca45989286ddcb"
    }
]
```


Figure 6 - Outputs of blockchain ledger containing sequential block creation with index, timestamps, identifiers of clients, chained SHA 256 hash have been used to store federated model changes.

In general, the findings suggest that the prototype can be relied upon to replicate the centralized model accuracy in a federated environment and enjoy the additional transparency provided by blockchain anchoring. The predictive power of the non linear modelling data to predict supply chain data was also demonstrated by the Random Forest benchmark, and the interpretability and error distribution analyses demonstrated the consistency and readability of the results of the system. These results form a solid basis of improvement in the future, such as the incorporation of GNN based relational reasoning, LLM based unstructured text processing, and incentive systems based on smart contracts.

6 Discussion

The results of the paper show that a lightweight federated learning architecture could be used to attain predictive performance similar to a centralized model with the added benefit of transparency and provenance tracking by blockchain anchoring. The small margin between the MAE of the federated and centralized linear models substantiate the theoretical assertions of the previous FL literature and indicates that distributed optimization can retain fidelity to the model under non IID and heterogeneous situations when suitable aggregation algorithms like

FedAvg are implemented (McMahan et al., 2017; Orabi et al., 2025). This finding supports the appropriateness of FL to multiple organization supply chains, in which privacy limits and competitive sensitivities make it impossible to pool raw data.

The much lower MAE of the Random Forest model underscores the degree to which non linear relationships are present in the supply chain data, especially the ones that are of an economic and volumetric character. This finding concurs with the results of the studies in the supply chain risk prediction field, where ensemble and tree based models tend to perform better than linear estimators since they are more effective in modeling heterogeneous transaction structures (Zheng et al., 2023; Kong et al., 2024). This is further depicted by the strong coefficient dominance of the quantity variable in the federated linear model. It implies that the scale of shipment is one of the major predictors of predictive behavior but the influence of pricing variables are moderate. These observations can be added to the enhanced comprehension of structural dependencies of supply chain exchanges to the literature that highlights the importance of domain sensitive feature interpretation (Kosasih et al., 2022).

The distribution of prediction error, which was narrowly centered around zero, with long, low frequency tails, is evidence that the federated model was stable in most of the samples, and only deviated in extreme or irregular cases. The same distribution patterns are reported in the literature of decentralized learning in which heterogeneous client data is a factor in the occasional changes in prediction residuals (Li et al., 2022). Nevertheless, the stability of the distribution in general justifies the possibility to apply federated modelling to dynamic supply chain data environments.

Another value added by this research is the use of a blockchain based audit log to document model changes. Though the lightweight design ensured that it was not burdened with decentralized data, the chained format of the JSON ledger ensured the continuity of the hash and was able to demonstrate the concept of tamper evident provenance, which was previously proposed in the context of blockchain enhanced FL integrity (Lu et al., 2019; Singh et al., 2023). This finding suggests that a small amount of blockchain integration can offer any useful level of accountability to decentralized learning systems, especially where there are many independent stakeholders in need of verifiable records of model development.

Although these are its strengths, the interpretation and generalizability of the findings are guided by a number of limitations. Owing to computational, infrastructure, data access limitations, GNN based relational reasoning, LLM based text analytics, tokenized incentive schemes and decentralized storage in IPFS were conceptually described but never done operationally. This is in line with what has been found in new BCFL literature, where full scale multi stakeholder implementations are largely hypothetical (Orabi et al., 2025; Kandil et al., 2024). Also, the prototype was based on simulated federated clients as opposed to actual organizational nodes, implying that the system was not subjected to the full spectrum of network instability, adversarial behavior and cross client governance issues characteristic of industrial environments.

However, the prototype provides valuable information on the design of decentralized intelligence in the context of supply chain settings. It empirically confirms that federated learning is not necessarily performance degrading, that interpretability can still be provided with the help of tracking the coefficient in a blockchain, and that blockchain anchoring can be added without significant computational cost. The results can be regarded as a significant move towards the creation of scalable, multi agent supply chain analytics systems that will combine privacy protection, transparency, and intelligent decision support.

7 Conclusion and Future Work

This study showed that, with a decentralized learning architecture that integrates federated learning and blockchain supported provenance tracking, a predictive supply chain analytics can be performed with reliable performance with preserved data privacy. The federated linear regression model performed almost as well as its centralized counterpart, which supports the conclusion that collaborative learning can be used effectively on fragmented, heterogeneous data without having to consolidate raw data. The Random Forest benchmark also indicated that strong non linear trends existed in the datasets and therefore the usefulness of flexible modelling methods in operational forecasting. The error distribution in prediction was stable in most of the situations, and the blockchain registry offered a tamper evident solution to tracking model changes, which strengthens the transparency and trust in multiple stakeholder settings. Though more sophisticated features like GNN based relational reasoning, text analytics powered by LLMs, tokenized incentive systems and decentralized storage were not deployed, the conceptual architecture is consistent with the current direction of research in BCFL and has a solid basis to build upon it.

Further development of the system should involve expansion of the system by performing the advanced modules identified in conceptual development into practice. The adoption of GNNs would permit the modeling of suppliers distributors customer interdependencies, which is why cascading risk can be identified within networked logistics systems. The introduction of LLM would enable the derivation of insights of the unstructured text of the supply chain like vendor messages, exception notices and regulatory notices, etc. Scalable storage of model artefacts can be made decentralized, verifiable, and be implemented using IPFS or File coin, whereas mechanisms of incentives that operate on smart contract systems can be used to incentivize long term engagement of stakeholders in federated learning ecosystems. The implementation of the system with real organizational nodes (not simulated clients) would also give hands on experience of how the system will perform under operational constraints, communication delays and adversarial behavior. Further optimization in the form of hyperparameter optimization, differential privacy, secure aggregation and robustness analysis would further provide the reliability and scalability of the architecture to be further extended to support the development of an end to end decentralized intelligence architecture that can be applied in the context of real world supply chains.

8 References

Bonawitz, K. et al. (2019). Towards federated learning at scale: System design, *Proceedings of Machine Learning and Systems*, 1, pp. 374–388. Available at: https://arxiv.org/abs/1902.01046

Brown, T.B. et al. (2020). Language models are few-shot learners, *Advances in Neural Information Processing Systems*, 33, pp. 1877–1901. https://doi.org/10.48550/arXiv.2005.14165

Chaganti, K.R. et al. (2024). Blockchain anchored federated learning and tokenized traceability for sustainable food supply chains, *2024 International Conference on Ubiquitous Intelligent Systems (ICUIS)*, pp. 1532–1538. https://doi.org/10.1109/ICUIS64676.2024.10866271

Chen, Y. et al. (2020). A training-integrity privacy-preserving federated learning scheme with trusted execution environment, *Information Sciences*, 522, pp. 69–79. https://doi.org/10.1016/j.ins.2020.02.060

Devlin, J., Chang, M., Lee, K. and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding, *NAACL-HLT 2019*, pp. 4171–4186. Available at: https://aclanthology.org/N19-1423

Kandil, S., Marzbani, F. and Shamayleh, A. (2024). Beyond blockchain: Enhancing data sharing in supply chains through horizontal federated learning in IoT-enabled VMI systems, *ASET 2024*, pp. 1–7. https://doi.org/10.1109/ASET60340.2024.10708671

Kong, L., Zheng, G. and Brintrup, A. (2024). A federated machine learning approach for order-level risk prediction in supply chain financing, *International Journal of Production Economics*, 268, 109095. https://doi.org/10.1016/j.ijpe.2023.109095

Kosasih, E.E. et al. (2022). Towards knowledge graph reasoning for supply chain risk management using graph neural networks, *International Journal of Production Research*, 62(15), pp. 1–17. https://doi.org/10.1080/00207543.2022.2100841

Li, T., Sahu, A., Talwalkar, A. and Smith, V. (2020). Federated learning: Challenges, methods, and future directions, *IEEE Signal Processing Magazine*, 37(3), pp. 50–60. https://doi.org/10.1109/MSP.2020.2975749

Li, D. et al. (2022). Blockchain for federated learning toward secure distributed machine learning systems: A systemic survey, *Soft Computing*, 26(9), pp. 4423–4440. https://doi.org/10.1007/s00500-021-06496-5

Liu, Y. et al. (2020). A secure federated learning framework for 5G networks, *IEEE Wireless Communications*, 27(4), pp. 24–31. https://doi.org/10.1109/MWC.01.1900525

Lu, Y. et al. (2020). Blockchain and federated learning for privacy-preserved data sharing in industrial IoT, *IEEE Transactions on Industrial Informatics*, 16(6), pp. 4177–4186. https://doi.org/10.1109/TII.2019.2942190

Mahmood, Z. and Jusas, V. (2022). Blockchain-enabled multi-layered security federated learning platform for preserving data privacy, *Electronics*, 11(10), 1624. https://doi.org/10.3390/electronics11101624

McMahan, H.B. et al. (2017). Communication-efficient learning of deep networks from decentralized data, *AISTATS 2017*, pp. 1273–1282. http://proceedings.mlr.press/v54/mcmahan17a.html

Moudoud, H. and Cherkaoui, S. (2023). Multi-tasking federated learning meets blockchain to foster trust and security in the metaverse, *Ad Hoc Networks*, 150, 103264. https://doi.org/10.1016/j.adhoc.2023.103264

Nguyen, T.T. et al. (2025). Artificial intelligence applications for supply chain risk management considering interconnectivity, external events exposures and transparency: A systematic literature review, *Modern Supply Chain Research and Applications*, 7(2). https://doi.org/10.1108/MSCRA-10-2024-0041

Onwubiko, A. et al. (2023). Enabling trust and security in digital twin management: A blockchain-based approach with Ethereum and IPFS, *Sensors*, 23(14), 6641. https://doi.org/10.3390/s23146641

Orabi, M., Emam, O. and Fahmy, H. (2025). Adapting security and decentralised knowledge enhancement in federated learning using blockchain technology, *Journal of Big Data*, 12(55), pp. 1–25. https://doi.org/10.1186/s40537-025-01099-5

Qammar, A., Naouri, A., Ding, J. and Ning, H. (2023). Blockchain-based optimized edge node selection and privacy-preserved framework for federated learning, *Cluster Computing*. https://doi.org/10.1007/s10586-023-04145-0

Ribeiro, M., Singh, S. and Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier, *KDD 2016*, pp. 1135–1144. https://doi.org/10.1145/2939672.2939778

Singh, R., Kumar, A. and Sharma, P. (2023). Enhancing the security and privacy in the IoT supply chain using blockchain and federated learning with trusted execution environment, *Mathematics*, 11(17), 3759. https://doi.org/10.3390/math11173759

Wang, C., Pan, D., Shabaz, M. and Jiang, H. (2025). Privacy-preserving federated learning with meta-knowledge training for AIoT-enabled supply chain systems: A case study on smart healthcare, *IEEE Internet of Things Journal (Early Access)*. https://doi.org/10.1109/JIOT.2025.3595859

Zheng, G., Kong, L. and Brintrup, A. (2023). Federated machine learning for privacy-preserving, collective supply chain risk prediction, *International Journal of Production Research*, 61(23), pp. 8115–8132. https://doi.org/10.1080/00207543.2022.2164628