International Journal of Engineering and Information Systems (IJEAIS)
ISSN: 2643-640X
Vol. 9 Issue 8 August - 2025, Pages: 50-62

A Federated Edge-Cloud Architecture for Autonomous Logistics Systems: Enabling Real-Time Coordination and Energy Efficiency

Thar Htet Min1 and Adam Kargbo2

1School of Electronic Information and Communications (EIC), Huazhong University of Science and Technology, Wuhan, China,
tharhtetmin97@gmail.com
2School of Electronic Information and Communications (EIC), Huazhong University of Science and Technology Wuhan, China
adamk@hust.edu.cn

Abstract— *The growing need for autonomous technologies and real-time decision-making is propelling the logistics sector's rapid digital transformation. In order to improve autonomous logistics systems' scalability, security, and energy efficiency, this article suggests a federated edge-cloud architecture. The architecture decentralises data processing by combining edge computing and federated learning, which lowers latency and protects privacy by preventing the transport of raw data to central servers. To improve global learning models, local edge devices execute model updates that are safely aggregated in the cloud. A tri-layered architecture that combines cloud orchestration and edge autonomy, battery-aware dynamic routing for energy optimisation, and federated anomaly detection to provide resilience against disruptions are some of the main achievements. According to experimental evaluation, delivery times can be shortened by up to 35% and energy usage can be decreased by 30%. Future developments like blockchain-enabled trust and 6G-driven smart logistics are covered in the paper's conclusion.*

Keywords— Federated Edge-Cloud, Autonomous Logistics Systems, Edge Computing, Federated Learning, Dynamic Routing, IoT, Security and Privacy.

1. INTRODUCTION

The logistics industry has changed a lot in the past ten years because of the huge growth of e-commerce, globalisation, and customers' growing demands for faster and more environmentally friendly deliveries. The complexity and unpredictability of modern urban supply chains have made traditional centralised logistics systems, which rely on static routing, human oversight, and isolated decision-making, inadequate [6]. These systems often have problems with high latency, scalability, and not being able to adapt to changes in real-time demand or infrastructure outages. Researchers and professionals are looking into Autonomous Logistics Systems (ALS), which are networks of self-driving trucks, drones, and IoT-enabled hubs that can make deliveries without any human help and make the most of real-time coordination. This is being done to solve these problems [1, 13]. Studies have shown that the ALS method lowers costs while also making deliveries much faster and more reliable [5].There is a lot of potential for autonomous logistics, but using them on a large scale is still difficult because they rely on centralised cloud computing infrastructures. In cloud-based models, data processing and decision-making happen on remote servers. This makes communication slower, uses more bandwidth, and makes the system more likely to fail at one point [8]. Also, the huge amounts of raw data that self-driving cars, sensors, and IoT devices collect must be sent to the cloud, which raises serious privacy and security issues [10]. These problems make centralised approaches bad for logistics networks that change quickly, especially in cities where routing decisions need to be made in real time. Edge computing has recently made progress that has led to a new way of doing distributed computing that processes data closer to its source. This cuts down on latency and makes the system more responsive [7]. Federated learning has also become a complementary paradigm that lets machine learning models be trained in a distributed way without having to share raw data. This improves privacy and scalability [1, 11].

A federated edge-cloud architecture built on a strong base of edge computing and federated learning fixes the problems with traditional centralised logistics systems right away. In this hybrid model, the cloud combines information from different places to make models that work best for everyone. Edge nodes, like delivery drones, driverless cars, and local hubs, can do important tasks in real time, like changing routes, avoiding obstacles, and managing energy [2, 12]. This design makes it possible to make quick decisions close to the data source. It also makes use of cloud computing's ability to coordinate large groups of people. According to studies, these distributed methods are especially well-suited for massive fleets of autonomous logistics systems functioning in both urban and rural settings since they not only lower latency but also enhance privacy, robustness, and scalability [18, 24]. These new technologies are driving the development of a new generation of smart, decentralised logistics systems that can better handle operational interruptions [21]. In large-scale autonomous logistics, one of the biggest problems is making sure that fleets of drones and electric vehicles use as little energy as possible and keep running. Traditional logistics planning doesn't take vehicle energy limits into account enough, which makes routing less effective and increases downtime. The federated edge-cloud paradigm [3, 4] makes it possible to use battery-aware routing algorithms that use real-time energy data from each car to change routes on the fly. These distributed methods let self-driving cars change their routes based on changing traffic conditions, battery levels, and the availability of charging stations. They also save time and energy by not running out of power. The system can also find possible risks or failures early and take local corrective action by combining edge-level anomaly detection with federated model updates. This makes the system more resilient without having to rely on the cloud for every decision [16, 23]. This kind of architecture is a big step towards logistics operations that are very independent, environmentally friendly, and safe.

This study suggests a federated edge-cloud architecture for self-driving logistics systems to get around the problems with centralised logistics and make the most of edge intelligence and federated learning. Large-scale, real-time delivery truck coordination is supported by the architecture, which also guarantees operational resilience, energy efficiency, and data privacy. Our main contributions are (i) a tri-layered system that uses federated learning to integrate edge autonomy and cloud orchestration; (ii) a dynamic routing algorithm that is battery-aware and optimised for fleets with limited energy; and (iii) distributed anomaly detection mechanisms that protect system stability without disclosing raw data. Simulation scenarios that mimic real-world logistics problems are used to assess the suggested architecture, and the results demonstrate notable gains in energy consumption and delivery efficiency [5, 18, 27]. This paper's remaining sections are organised as follows: In Section 2, relevant work and foundational technologies are reviewed; in Section 3, the proposed architecture and its components are presented; in Section 4, simulation results and performance metrics are discussed; and in Section 5, future extensions, such as blockchain integration and 6G-enabled mobility, are discussed.

2. RELATED WORK

2.1 Autonomous Logistics Systems

In contemporary supply chain and transportation operations, autonomous logistics systems, or ALS, have become a key breakthrough. Thanks to fleets of self-driving cars, drones, and infrastructure that can connect to the Internet of Things, these systems can handle distribution and coordination with little help from people. According to research done over the past ten years [6], ALS can improve route management, reduce delivery delays, and boost efficiency in crowded urban areas where traditional centralised systems are prone to congestion and delays. The benefits of ALS become even more clear in complicated, fast-moving supply chains where speed and flexibility are important [1].

The growth of e-commerce and the need for just-in-time delivery are two of the main things that have led to the rise of ALS. Unlike most systems, which need to be monitored and changed by people all the time, ALS can run all the time and adjust to changing demand patterns without any human help. Significant gains in cost-effectiveness and system responsiveness have been demonstrated by recent case studies and pilot deployments, with fleets of autonomous cars able to react instantly to emergency situations and erratic traffic circumstances [13, 15]. These systems' scalability enables businesses to satisfy client demands even during busy times, which is now a key component of contemporary logistics plans. Drone integration for last-mile delivery has been a key component of ALS innovation. Drones provide a high-speed alternative for urgent and time-sensitive deliveries by avoiding traffic on the roads. Drone swarming methods, which involve several UAVs coordinating with one another to maximise delivery coverage and minimise redundancy, have been developed by recent research and greatly boost efficiency for short-range logistical jobs [6]. The benefits of coordinated drone fleets have been confirmed by a number of smart city pilot projects, which highlight how their deployment can lower energy use and environmental effect [20]. However, for data analysis and decision-making, early ALS implementations frequently relied heavily on centralised cloud infrastructures. This method limits autonomous agents' capacity to act autonomously during network outages, causes delay, and forms bottlenecks [8]. Concerns regarding data privacy and operational robustness are also brought up by the cloud-based approach, especially in systems that must oversee hundreds or thousands of agents at once. Because of these constraints, researchers and industry are investigating more intelligent and distributed architectures that will enable autonomous systems to learn and coordinate more efficiently at the edge while preserving global visibility [21].

2.2 Edge Computing in Logistics

Because edge computing moves data processing and analytics closer to the point of data generation, it has become a vital enabler for contemporary autonomous logistics systems. Edge computing enables logistics systems to react faster to environmental changes, including demand spikes, weather variations, and route bottlenecks, by decreasing reliance on distant cloud servers [7, 8]. It has been demonstrated that having the ability to make decisions locally at the edge greatly lowers communication latency and eases the burden on network bandwidth. This is especially useful in large-scale urban deployments where thousands of devices may be sending data at once. Edge nodes, which can be mounted on delivery trucks, drones, and local hubs, can carry out real-time calculations like local optimisation, obstacle avoidance, and route recalculation in the context of logistics. These calculations are essential in time-sensitive situations, such avoiding crowded junctions or adjusting to sudden road closures, when milliseconds can impact delivery results [9].

System resilience is further increased by spreading compute duties among edge nodes, which enables businesses to continue functioning even in the event that network communication to the cloud is disrupted. Logistics operations can be revolutionised by edge intelligence, as evidenced by a number of recent studies. For example, multi-agent coordination, in which teams of self-driving delivery vehicles work together in real-time to choose the best delivery routes and assign duties throughout the network, has been managed using edge computing-based systems [22, 23]. In addition to increasing fleet utilisation, this lowers overall energy consumption, which is crucial for electric delivery fleets. Furthermore, warehouse-to-consumer routing has been optimised and demand changes have been predicted using edge-enabled predictive analytics, enabling proactive resource and vehicle allocation. Nevertheless, even with these developments, existing edge implementations still have drawbacks when utilised alone. Fragmented learning can result from localised decision-making, in which nodes optimise according to their immediate surroundings but are unable to access more comprehensive system knowledge. Edge computing by itself might not be able to fully utilise the collective intelligence that exists across thousands of autonomous devices if there is no way for insights to be shared throughout the network [19, 21]. These drawbacks encourage the incorporation of federated learning to enhance edge capabilities and offer a system-wide perspective while maintaining the advantages of local computation's low latency.

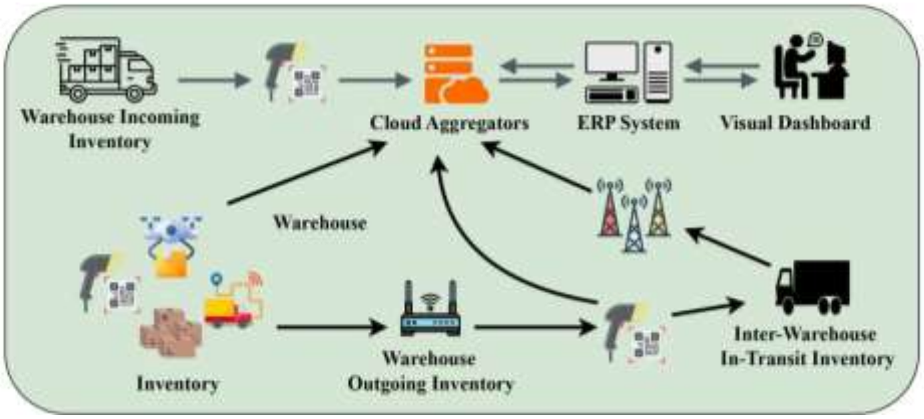

Figure 2.1: IoT-enabled warehouse management: The flow of inventory data between warehouses, cloud aggregators, ERP systems, and visual dashboards. This system ensures real-time synchronization and monitoring of incoming and outgoing inventory.

2.3 Federated Learning for Transportation and Logistics

Federated learning (FL), a potent machine learning paradigm that allows dispersed devices to work together to train models without sharing raw data, has grown in popularity recently. This is especially helpful for logistics and transportation systems, which produce enormous amounts of sensitive data from urban infrastructure, IoT sensors, and driverless cars. FL makes privacy better and communication easier by letting models be trained locally at the edge and then combined every so often to make a global model, instead of relying only on cloud servers [10, 11]. The ability to combine local intelligence from hundreds or thousands of devices opens the door to logistics networks that are more flexible and can grow. In logistics and transportation, FL has been used for a number of things, such as planning routes together, sensing autonomous vehicles, predicting traffic flow, and finding anomalies. By letting autonomous agents train models locally and only sharing model updates, systems can use community knowledge while keeping sensitive raw data, like location traces and operational metrics, decentralised [1, 10, 11]. Many studies have shown that FL is good for driverless and connected cars. They show that federated methods make it easier for cars to adapt to local conditions and reduce the need for large amounts of centralised data collection [12, 18, 19]. This strategy works especially well for logistics fleets that work in a lot of different urban and rural areas, where conditions can be very different from one place to another [21, 23]. Researchers have also begun combining FL with reinforcement learning to make distributed intelligence better in logistics platforms and vehicle networks [22, 24]. Federated learning has shown promise in helping logistics fleets make decisions that take energy use into account, beyond just observation and routing. Battery-powered self-driving cars and drones can benefit from distributed models that can find the best charging schedules, predict how much energy will be used, and change routes based on how much charge they have [3, 4, 5]. Recent research has also shown how FL can be used to coordinate charging infrastructure so that a lot of cars can get power without causing traffic jams at charging stations [16, 23]. These systems can change their strategies to cut down on downtime and make batteries last longer by looking at patterns from many different local situations and putting them together in the cloud [18, 24]. This energy-focused FL application is now very important for creating logistics systems that are both cheap and good for the environment [19, 27].

Federated learning can make logistics systems safer and more private by keeping raw data local. This is another big benefit. The dangers of disclosing sensitive data, including vehicle positions or operational specifics, are greatly decreased because only model parameters are sent throughout the learning phase [10, 21]. To guard against possible model inversion attacks, researchers have also investigated incorporating privacy-preserving strategies like safe aggregation and differential privacy into FL frameworks [22, 24]. These techniques are becoming more and more applicable to big logistics fleets where operations depend on data security and integrity. Furthermore, because no single server houses all operational data, FL's decentralised architecture increases resilience and lessens the impact of possible intrusions [23, 25]. Federated learning's trust-enhancing qualities make it an especially appealing choice for next-generation autonomous transportation networks as logistics systems become more interconnected [12, 26]. Although federated learning has already demonstrated promise in resource optimisation, traffic coordination, and autonomous vehicles, the majority of current applications are still in the experimental stage and have small scales [18, 21]. Combining these distributed models with real-time decision-making frameworks that can operate dependably across thousands of diverse devices is still a challenge [23, 24]. Due to this constraint, there is growing interest in architectures that combine edge computing and federated learning, laying the groundwork for scalable and intelligent autonomous logistics systems [2, 12].

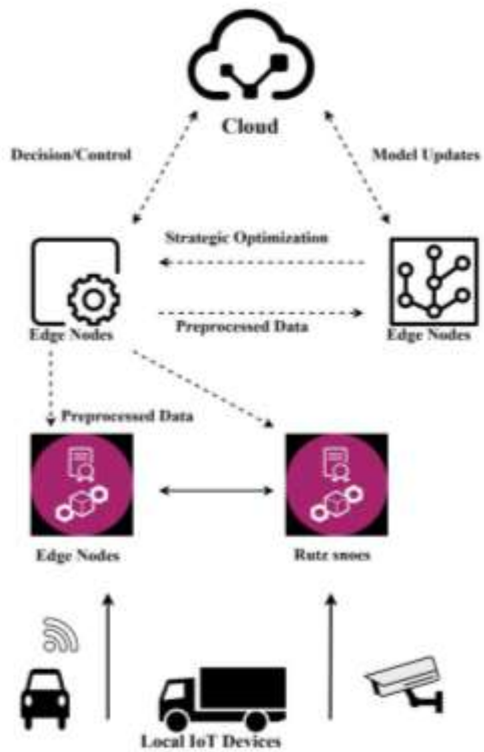

Figure 2.2: Overview of the proposed tri-layered federated edge-cloud architecture

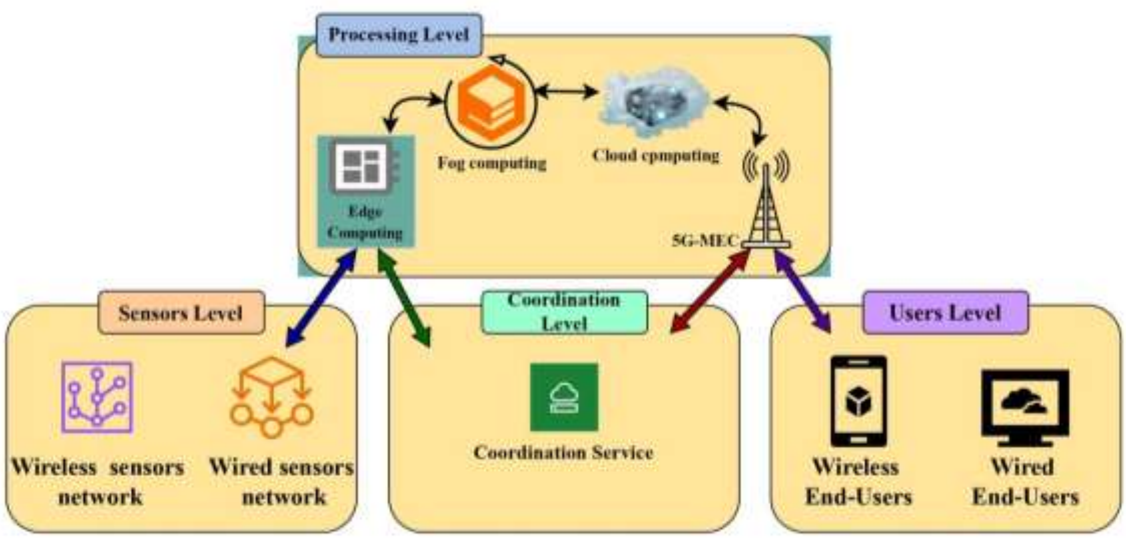

Figure 2.3: Hierarchical data flow architecture: Sensors, Edge, Coordination, and User Levels in the proposed system for autonomous logistics operations.

2.4 Integrated Edge-Cloud Architectures

The difficulties of centralised logistics systems have lately been addressed by hybrid edge-cloud designs. These designs improve robustness, lower latency, and facilitate real-time decision-making by sharing intelligence between cloud servers and local edge nodes. Without waiting for commands from the cloud, edge nodes installed in cars, drones, or local depots can quickly modify routes, identify irregularities, and handle temporary jobs in logistics environments [2, 8, 22]. Long-term optimisations like demand forecasting and fleet-wide coordination are made possible by the cloud's simultaneous aggregation of model updates from these nodes [12, 23]. This job separation makes large-scale deployments more responsive and consistent around the world. Combining edge computing and federated learning in these designs has shown that autonomous logistics has a lot of potential. Studies show that these combined systems protect privacy by keeping sensitive information on the user's device, which also lowers connection costs and latency [10, 18, 24]. Local edge intelligence is great for dynamic situations where cloud access can't be guaranteed all the time because it lets autonomous agents keep working even when the network goes down.

Federated learning also makes sure that insights from many operating locations go into a global model without putting data security at risk [11, 21]. These changes point to a move towards smart and decentralised logistics systems that can work on a large scale and quickly adapt to changing situations.

2.5 Identified Research Gaps

Despite significant advancements in federated learning, edge computing, and autonomous logistics systems, there are still significant gaps in their integration. Many autonomous delivery system solutions are still centralised, which restricts scalability and makes them more susceptible to network outages [6, 8]. Similar to this, edge computing deployments have been shown to work in isolated pilots, but they frequently don't communicate information throughout a fleet, which leads to intelligence that is fragmented [7, 22]. Despite its theoretical strength, federated learning has only been investigated in controlled research settings and has not yet been used in the large-scale, real-time scenarios encountered by contemporary logistics providers [10, 12]. The optimisation of resources and energy is another gap. Current research typically focusses on predictive modelling or autonomous vehicle operational coordination, but it rarely discusses how distributed intelligence can integrate real-time energy constraints, like battery health and charging availability, into scheduling and routing [3, 5, 16]. Moreover, insufficient research incorporates anomaly detection directly into these architectures, making systems more vulnerable to errors or cyberattacks that could cause large-scale disruptions in operations [21, 23]. A number of significant shortcomings in the current autonomous logistics systems are addressed by the suggested federated edge-cloud design. These research gaps and their accompanying solutions are compiled in Table 2.1, which also emphasises the need for a hybrid strategy that combines federated learning with edge computing.

These drawbacks emphasise the necessity of a cohesive architecture that incorporates federated learning and edge computing to accomplish resilience, global coordination, and real-time decision-making while taking energy-conscious logistics into account. The strategy in this work tries to close these gaps and provide scalable, secure, and effective logistics solutions by offering a federated edge-cloud architecture that strikes a balance between local autonomy and centralised optimisation [2, 18, 24].

Table 2.1: Key Research Gaps in Autonomous Logistics Systems and Proposed Solutions

Research Gap	Description	Proposed Solution
High Latency in Centralized Systems	Cloud-centric models cause delays due to remote data processing [6, 8].	Edge computing for local, real-time decision-making (e.g., route adjustments) [7, 18].
Poor Scalability in Edge Computing	Isolated edge nodes lack global coordination, limiting fleet-wide optimization [7, 22].	Federated edge-cloud architecture for distributed learning and global model updates [2, 12].
Limited Privacy in Data Handling	Raw data transfer to cloud servers risks breaches and regulatory issues [10].	Federated learning to keep data local, sharing only model updates [10, 21].
Inefficient Energy Management	Traditional routing ignores battery constraints, leading to downtime [3, 5].	Battery-aware dynamic routing to optimize energy use and reduce consumption [3, 4].

3. Proposed Architecture

3.1 System Overview

The proposed solution is a federated edge-cloud architecture made specifically to make autonomous logistics systems more scalable, reliable, and efficient. The architecture blends the collaborative intelligence of Combining edge computing's low-latency processing power with federated learning [2, 8, 22]. Autonomous vehicles, drones, and local depots all have edge nodes that do time-sensitive calculations and react instantly to shifting operational and road conditions [3, 7, 18]. To enable fleet-wide optimisation without jeopardising data privacy, the cloud acts as an orchestrator, integrating model updates from these scattered nodes to get a system-wide image [10, 12, 23]. Even in situations with sporadic connectivity and diverse settings, the architecture may operate in real time thanks to this dual-layered approach [11, 19, 24]. The necessity to oversee fleets of drones and driverless cars that must function in a variety of uncertain conditions is what inspired this architecture. The suggested structure gives the system the ability to act locally when time is of the essence, in contrast to centralised models where all choices are dependent on a distant data centre. When unexpected road restrictions occur, for instance, a delivery vehicle can use its onboard intelligence to quickly calculate a different route while still receiving periodic updates from the global model pooled in the cloud [7, 18, 21]. Delivery delays are decreased, the system as a whole is more resilient to network outages, and federated learning allows lessons learnt in one area of the fleet to be shared with others [2, 10, 12]. Crucially, the suggested architecture's design also takes privacy into account by eliminating the transfer of raw data; rather, devices exchange model parameters, preserving sensitive operational data inside local nodes [19, 22, 24]. The suggested framework creates a basis for extensive, intelligent logistics operations by combining local autonomy with centralised knowledge sharing. While the cloud continuously improves a common model for future planning, this balancing guarantees that quick decisions may be taken at the network edge [8, 11, 18]. The system's multi-layered structure and the way intelligence and data move across the architecture are explained in the ensuing subsections.

3.2 Architecture Layers

Three levels make up the suggested federated edge-cloud architecture: the edge layer, the cloud layer, and the intermediate coordination layer. In order to provide scalable, flexible, and privacy-conscious logistics operations, each of these layers has a unique function. While the coordination layer connects groups of edge nodes within an area to exchange information and settle disputes, the edge layer handles instantaneous, on-site decision-making, such as modifying delivery routes or identifying nearby risks [2, 3, 22]. The cloud layer at the top is the main intelligence centre. It combines model updates to make global optimisation methods [8, 12, 23]. This layered approach keeps the benefits of global learning while letting information flow up and down. It also makes it possible for quick responses at the edge.

3.2.1 Edge Layer

The edge layer is made up of drones, self-driving cars, and local IoT-enabled devices with processing power. It is the base of the architecture. These edge nodes do things that are time-sensitive, like checking battery life, finding problems in their immediate area, and changing routes when they come across unexpected obstacles [3]. When cars drive in areas with bad connectivity, they can take action without waiting for orders from the cloud because they can process data on their own [7]. Edge intelligence also makes context-aware decision-making possible. For example, a delivery drone can use sensor data to respond to sudden changes in weather or areas that are off-limits without needing outside help [18]. Also, each edge node helps train the model by learning from its own data and then sending only the updated model to higher layers. This is done to protect data privacy [10].

3.2.2 Intermediate Coordination Layer

The intermediate coordination layer connects each edge node to the cloud. Regional hubs in this layer combine local knowledge, deal with problems like delivery routes that cross each other, and get updates from drones and nearby cars [12]. This design makes it easier for nodes in the same area to work together while putting less strain on the cloud's communication [21].

3.2.3 Cloud Layer

The cloud layer is at the top of the architecture and is in charge of orchestration and global learning. It creates optimisation techniques for the entire system by combining model updates from several regional hubs [8]. These tactics include long-term planning techniques that necessitate a comprehensive picture of all operational delivery units, such as fleet scheduling and demand forecasting [18]. Without exchanging raw data, the cloud also makes it easier to transmit knowledge between geographical areas, allowing operations in one city to profit from lessons learnt in another [10]. The cloud keeps an extensive model that is always changing as fresh data comes in from the edge thanks to the combination of these insights. This architecture's tiered design combines the advantages of each tier to create a unified and flexible system. Logistics operations can continue even when connections to the higher levels are unreliable by letting the edge layer make quick, localised decisions. By controlling resources and traffic patterns within a certain area, the intermediate coordination layer—also known as a fog or regional layer—adds a crucial dimension and keeps the cloud from becoming overloaded with requests for constant communication [12]. Lastly, by preserving a global viewpoint that facilitates large-scale optimisation, such as predicted demand models and fleet-wide planning, the cloud layer integrates these dispersed activities [2]. As orders and updates move downhill to enhance field operations, knowledge from vehicles and drones flows upward thanks to the constant interaction between these layers. This collaborative framework serves as the foundation for data flow and model training within the system and is essential for handling the complexity and unpredictability of logistical networks in the real world. The suggested federated edge-cloud architecture uses a number of cutting-edge technologies to make it scalable, responsive, and efficient. These important technologies, together with their main purposes and contributions to system performance, are compiled in Table 3.1.

Table 3.1: Key Technologies Enabling the Federated Edge-Cloud Architecture

Technology	Primary Function	Benefit	Relevance to Architecture
Edge AI	Local data processing and decision-making	Reduces latency, enhances responsivenes	Enables real-time routing and anomaly detection at edge nodes [7, 18].
Federated Learning	Distributed model training without raw data	Preserves privacy, improves scalability	Supports global model updates via cloud while keeping data local [10, 21].
Battery-Aware Routing	Dynamic route optimization based on energy	Reduces energy consumption by up to 30%	Optimizes energy-efficient fleet operations [3, 4].
Anomaly Detection	Real-time monitoring of system irregularities	Enhances resilience, prevents disruptions	Ensures operational stability at edge and regional levels [18, 23].

3.3 Data Flow and Model Training

The three layers of the proposed design make it easy to see how data and models are updated. Drones and cars on the edge are always gathering information from sensors like GPS, cameras, and battery monitors [3]. These raw data streams go through local pre-processing [7] to make sure that only useful features are used when training models for things like planning routes and predicting energy use. Each edge device makes a set of model updates that show what it has learnt instead of sending sensitive data to the cloud. This greatly cuts down on the need for communication and keeps operations private [10]. The intermediate coordination layer puts these changes together and sends them to the cloud, where a global model is updated to include information from the whole network [2].

During the training process, the system goes through cycles that start at the edge. Any edge device, like a car, drone, or local hub, can improve local models with the data it has collected without having to send raw data outside of its borders [11]. After a training session, the device makes a small package of model parameters that shows what it has learnt and sends it to the regional coordination centre [7]. These hubs work as a middle layer that efficiently captures the range of conditions in a given area by combining local updates from multiple devices using federated averaging techniques [21]. After this regional aggregation, only updates for the improved model are uploaded to the cloud for global aggregation. This is to make sure that all of the fleet learns from what they learnt in different operating situations [10]. Our layered federated learning method lowers the chance of bandwidth congestion while still making models that work in a wide range of situations by finding a balance between global accuracy and communication efficiency [2]. When the global model is updated and sent back to regional hubs and then to individual devices, the cycle is complete. This cycle keeps improving decision-making at all levels of architecture [12]. One big advantage is that this organised data flow naturally protects privacy while allowing for flexible, real-time responses. Only model updates, not raw operational data, get to the edge nodes [10]. This means that sensitive data like position records or sensor images never leave the vehicles or drones. Using this method makes it much less likely that there will be big data breaches. Also, the system can handle bad network conditions because it relies on localised model training. This means that devices can keep learning and making decisions even when the connection is bad [18]. Because of the combination of distributed learning and synchronised model sharing [8], the architecture can quickly adapt to unexpected events like traffic jams, extra delivery orders, or environmental threats. The shared global model gets more accurate over time as it takes into account experiences from all parts of the logistics network [2, 12].

This cyclical flow of data and model updates in the suggested design combines the benefits of global coordination with local autonomy. Because of this, the system keeps strict privacy limits while getting more accurate and responsive with each training cycle [7, 11]. Thanks to this foundation, the functional modules described in the next paragraph can work well and reliably in a wide range of logistics situations.

3.4 Key Functional Modules

The suggested architecture will only work if it has a set of specialised modules that work across multiple system tiers. These modules handle things like finding anomalies, scheduling with energy in mind, dynamic routing, and learning that protects privacy [3]. Each module works at the level that works best for it. For instance, long-term optimisation modules use information from the cloud, while routing and hazard detection modules mostly work at the edge, where they can make decisions right away [8]. When put together, these parts create an ecosystem that lets people make quick decisions when they need to and improve slowly through working together with people from other countries [2, 18]. The dynamic routing and scheduling module is a key part of the suggested architecture. It lets cars and drones quickly adapt to changes in the logistical environment. The module at the edge constantly checks sensor and local map data [7] to find delays, traffic jams, or unexpected obstacles. It makes sure that deliveries keep going by recalculating the best routes when there are problems, without having to wait for input from the cloud [3]. The module also uses data from regional hubs to help cars in the same area plan their trips and avoid making unnecessary trips [12]. Over time, the cloud improves a global model of routing techniques that can predict future traffic patterns and better allocate resources with the help of aggregated updates from this module [2]. This multi-level decision process makes logistical operations faster and more reliable. It is much better than the static planning found in centralised systems [18].

The suggested architecture includes an energy-aware management module that takes care of energy efficiency, which is also very important. This module gets information about power use, charging station availability, and battery health directly from each car and drone [3]. It changes delivery routes and assignments based on this information to stay within energy limits, which lowers the chance of service interruptions caused by low battery levels [4]. By using aggregated energy data, coordination hubs can manage the charging schedules for different vehicles at the regional level. This cuts down on traffic at charging stations and spreads the load more evenly [16]. The system uses predictive models that look at past patterns of energy use to plan for large-scale operations and seasonal changes on the cloud side [2]. This all-encompassing approach to energy management supports the goal of creating a sustainable and effective autonomous logistics network [18].

The architecture has a multi-layered anomaly detection and resilience module to make sure that operations are safe and stable. This module at the edge keeps an eye on sensor readings and vehicle performance all the time to find problems like batteries that drain too quickly, communication problems, or mechanical issues [18]. The system can fix a local problem right away by rerouting jobs or isolating the affected node, without having to wait for commands from the cloud [7]. Regional hubs make this feature even better by being able to spot trends among multiple cars in the same area. They also help identify coordinated malfunctions or environmental risks that affect multiple devices at the same time [21]. Cloud-based aggregated anomaly data give fleet managers useful information that can help them make safety rules better and predict possible dangers [10]. This tiered detection method makes autonomous logistics networks much more durable when things go wrong in the real world [12].

When put together, these modules make a coordinated system that helps with independent logistical operations by coordinating energy management, routing, and resilience measures. By splitting these tasks across three layers, the design makes sure that the edge can respond quickly while also getting better all the time through cloud-based aggregated learning [2, 12].

3.5 Advantages of the Proposed Architecture

The proposed federated edge-cloud architecture has many benefits that can help with problems that keep coming up in autonomous logistics. By combining collaborative model training with distributed computing, the approach makes real-time operations more responsive and speeds up decision-making [7]. Also, it protects data privacy by keeping private information, like delivery locations and sensor recordings, on local devices instead of centralised servers [10]. This structure lets fleet operations see things from a global point of view, which makes it easy to adapt to changes in traffic, road conditions, or the environment [2, 12]. One of the best things about the suggested architecture is that it can grow with the logistics network. Adding more cars or drones to a traditional centralised system can make it work much worse by putting more strain on the central servers and slowing down communication [8]. With this federated edge-cloud approach, on the other hand, computing is split up into different levels, so each edge device can process its own data and send model updates instead of raw data [18]. This setup makes sure that the network will stay stable and responsive even when there are thousands of devices connected to it. Also, the layered structure makes fault tolerance even stronger. When communication with higher layers is lost, autonomous cars and drones can still work because they can make decisions on their own at the edge [7]. Intermediary hubs can help fleets of vehicles work together in the same way, so that the operations of an entire region don't stop because of one point of failure [21]. These traits help make a system strong enough to work in a variety of real-world situations that can't be predicted [2].

Another big benefit is how this architecture combines energy efficiency and sustainability into big logistics operations. Using distributed energy-aware modules [3], the system keeps an eye on the power usage and battery health of each drone and self-driving car all the time. Regional hubs can manage charging schedules so that some stations don't get too busy, but the edge uses this information to create paths that use the least amount of energy [4]. Long-term study of these energy trends helps with predictive planning at the cloud layer, which predicts when energy use will be highest and makes sure resources are used as efficiently as possible [16]. This makes the whole fleet run more sustainably, with fewer problems and longer battery life [18]. This combination of optimising local energy use and predicting global trends encourages greener logistics practices in response to the growing demand for environmentally friendly supply chain solutions [5]. When dealing with private information, autonomous logistics systems need to put security and privacy first. Table 3.2 lists the main security risks in the federated edge-cloud architecture and the ways that these risks are reduced to ensure safe and private operations.

To sum up, the proposed architecture brings together sustainability, scalability, resilience, and responsiveness into one system. These strengths work together to make autonomous logistics networks better all the time while still working reliably in the real world [2, 7, 12]. The next section shows how the experimental setup and evaluation were used to test this architecture's performance.

Table 3.2: Security Threats and Mitigation Strategies in Federated Edge-Cloud Systems

Security Threat	Description	Mitigation Strategy	Relevance to Architecture
Data Breaches	Unauthorized access to sensitive data (e.g., vehicle locations) during transmission [10].	Secure communication protocols and encryption to safeguard data while it's in transit [21].	Ensures secure data exchange between edge and cloud [21].
Model Inversion Attacks	Attackers reconstruct raw data from model updates [22].	Data reconstruction is prevented by using differential privacy to introduce noise during model updates [10, 24].	Preserves privacy in federated learning [10].
Adversarial Attacks	Malicious nodes send corrupted model updates, degrading global model accuracy [12].	Secure aggregation to validate and combine updates, filtering malicious inputs [22].	Maintains integrity of global model updates [12, 22].
Communication Interception	Eavesdropping on model updates during edge-cloud communication [21].	Homomorphic encryption to enable computation on encrypted data [24].	Protects model updates in transit [24].

4. Discussion And Future Directions

4.1 Discussion

The proposed federated edge-cloud architecture is a new way to deal with the problems that modern autonomous logistics systems have. Comparing the results of similar research with the structure of our architecture shows that combining hierarchical processing with distributed learning can greatly improve operational efficiency [2]. As shown by many previous studies [7], decentralised processing at the edge lowers latency and makes self-driving cars more responsive to their surroundings. Federated learning, on the other hand, makes sure that knowledge is shared across the network without putting privacy at risk [10]. These results support the idea that a hybrid architecture like this one would work better than traditional centralised methods when it comes to responsiveness, scalability, and adaptability [12].

A federated edge-cloud architecture will have a big effect on how logistics networks are planned and run. Integrating intelligence at the edge lets businesses make decisions locally. This cuts down on the time it takes to respond to things like equipment breakdowns, sudden spikes in demand, or road closures by a lot [7]. This increased responsiveness makes service more reliable, which is important in fields where customers want faster and more predictable delivery. The architecture actually makes sure that operations continue at a high level by allowing fleets of drones and self-driving cars to work even when there isn't much or any cloud connectivity [3]. In addition to these immediate benefits, logistics operators can find inefficiencies, plan their resources more strategically over time, and look at performance trends across thousands of trucks with the help of global coordination that federated learning makes possible [2]. These skills are especially important for businesses that manage complex networks across multiple cities or around the world because they help make decision-making processes more consistent while still taking into account the unique needs of each local area [18]. Another important effect is cost-effectiveness: decentralising processing workloads means that less cloud-based computation is needed all the time, which lowers bandwidth use and encourages fairer use of computer resources [12]. Also, this structure naturally helps with legal compliance when it comes to data protection because sensitive data like location traces and sensor images are stored locally instead of being sent to central servers [10]. This system is a good candidate for widespread use in modern smart logistics ecosystems because companies that use this kind of architecture can improve service quality, lower costs, and strengthen compliance all at the same time.

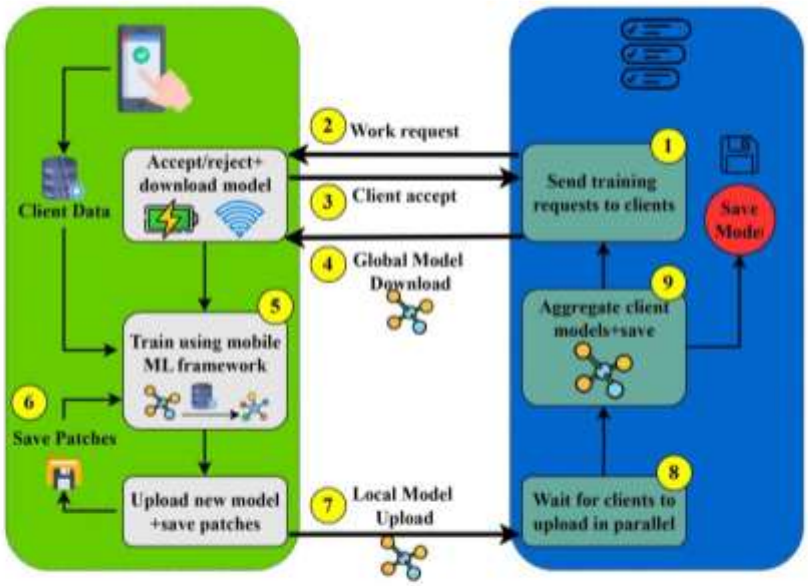

Figure 4.1: Federated learning workflow between mobile/edge devices and cloud aggregator. Devices accept work requests, train local models, upload results, and participate in global model updates in a privacy-preserving manner.

The advantages of a federated edge-cloud architecture are obvious, but before these systems can be used extensively, there are several problems that must be fixed. The initial infrastructure expense, which involves establishing regional coordination hubs and mounting computing units on cars, is one of the main obstacles [21]. Interoperability is another challenge because logistics networks sometimes combine hardware and software from several manufacturers, This may make communication between devices and layers more challenging [18]. There are dangers associated with model update security and governance as well; if a hostile node delivers tainted updates, the accuracy of the global model may suffer [12]. Furthermore, synchronising model updates across numerous heterogeneous devices becomes more difficult with federated learning, particularly when the network is unreliable [2]. Last but not least, before these technologies can be widely implemented, operational norms including adherence to privacy laws and the moral use of data need to be standardised [10].

All things considered, the conversation shows that the suggested architecture may significantly increase the responsiveness, resilience, and efficiency of autonomous logistics networks. However, overcoming obstacles pertaining to cost, interoperability, and governance is necessary to actually realise these benefits. To fully utilise federated edge-cloud systems for upcoming logistics operations, these obstacles must be removed [2, 7, 10].

4.2 Future Directions

Opportunities to expand and improve the suggested federated edge-cloud architecture are presented by a number of new technologies. Integrating blockchain technology with federated learning to produce unchangeable records of model updates and guarantee that only reliable nodes add to the global model is one exciting topic [21]. Simultaneously, the introduction of 6G networks is anticipated to offer extremely dependable low-latency communication, greatly enhancing synchronisation between edge devices and cloud servers, especially for logistics activities that require quick response times [12]. The creation of digital twins, which can replicate real-world logistics systems in virtual settings to model routes, forecast failures, and assess tactics before implementing them in actual situations, is another intriguing field [27]. Future logistics networks may be able to achieve previously unheard-of levels of coordination and flexibility when these technologies are paired with federated edge-cloud systems [18].

Applying this architecture to multi-modal logistics systems—where trucks, drones, railroads, and even marine vehicles collaborate—is a logical next step. For commodities to move effectively across regional and worldwide networks, future autonomous supply chains will rely on the smooth integration of these many modes of transportation [5]. Since each mode can make decisions locally and still benefit from global optimisation through federated learning, the layered architecture suggested in this study is especially well-suited to accommodate such complex systems [2]. Furthermore, the realisation of these multi-modal ecosystems would necessitate cooperation between sectors like manufacturing, retail, and urban planning, necessitating platforms that can facilitate data sharing and interoperability without sacrificing privacy [12]. Logistics firms will be able to move from single-mode planning to a comprehensive strategy where items go through an adaptive, networked system thanks to these developments. Increasing the federated edge-cloud systems' scalability to accommodate much bigger and more intricate networks is another avenue for future research. Model synchronisation becomes increasingly difficult as the number of autonomous devices increases, necessitating adaptive communication techniques to maintain timely and accurate updates [18]. Mechanisms for asynchronous model aggregation, which allow nodes with varying connectivity circumstances to contribute to learning without waiting for all devices to be online simultaneously, are probably going to be incorporated into future systems [21]. This will be supplemented by continuous learning, which does not rely on predefined training cycles but instead allows the system to improve its models as new data comes in [23]. These features will let global logistics platforms adapt to long-term changes in things like how people shop, how cities are built, and how demand changes with the seasons. Ethics and sustainability will also have a big impact on the directions of future research. Combining federated learning with green computing methods can help lower the energy footprint of large-scale distributed systems [4]. Adding explainable AI methods to the architecture can also make sure that the actions of autonomous agents are clear, allowing human operators to check results and follow the law [22]. People are starting to realise that their trust in AI-powered logistics systems will depend on more than just how well they work. They will also need to be able to handle privacy, fairness, and the environment [10]. Therefore, in order to make sure that these systems are in line with societal and environmental objectives, future work must incorporate frameworks for monitoring and auditing them.

Developing cross-industry data ecosystems that enable manufacturers, city authorities, and other logistics service providers to work together via federated edge-cloud infrastructures is one exciting avenue for future study. By sharing model updates via a secure, federated network, several stakeholders could obtain insights from a much bigger dataset without revealing sensitive information, as opposed to each organisation maintaining separate models [18]. City-level optimisations, in which delivery fleets from various suppliers plan their routes to avoid traffic jams and more effectively share charging infrastructure, may be made possible by such ecosystems. Additionally, this strategy would cut down on unnecessary travel and energy use, which would help businesses and communities alike. New guidelines for compatible federated learning protocols and reward systems would have to be created in order to make such ecosystems feasible [12]. Future systems must be more resilient to disturbances brought on by malfunctions, natural catastrophes, or cyberattacks as logistics networks get bigger and more sophisticated. In order to share local knowledge and assist one another in completing tasks during connectivity outages, autonomous vehicles or drones might dynamically establish ad hoc clusters. This is one area of research that focusses on building collaboration mechanisms among edge devices [21]. Peer-to-peer federated learning of this kind has the potential to build self-healing networks that carry on even in the event that certain nodes or communication channels malfunction [7]. In addition to resilience, autonomous collaboration would allow for specialised roles in the fleet: certain devices may be responsible for gathering environmental data, while others serve as coordinators, temporarily taking over decision-making duties in the event that central coordination is not accessible [2]. Logistics operations would be much more resilient and flexible in unstable, real-world settings with this capabilities. The close integration of autonomous logistics systems with intelligent urban infrastructure is another avenue with a lot of promise. In order to supply logistics networks with useful data, smart cities are progressively implementing sensors, intelligent road systems, and connected traffic lights [5]. Autonomous fleets can improve their route planning, prevent traffic jams before they arise, and handle deliveries more precisely by incorporating these infrastructure-generated datasets into federated edge-cloud systems. In order to reduce conflicts and increase efficiency, such integration would also enable anticipatory coordination with other city services, such emergency vehicles or public transportation [18]. For city services and private logistics providers to work together seamlessly, future designs must prioritise secure communication protocols and real-time APIs [23]. In the future, logistics networks will not be standalone systems functioning independently, but rather an integral component of the smart city ecosystem.

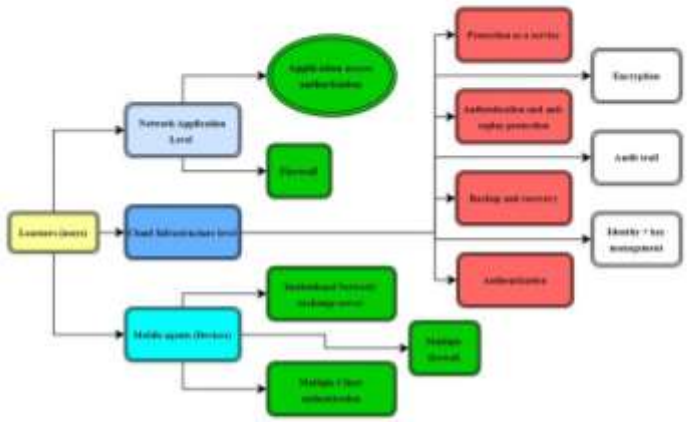

Figure 4.2 : Proposed data protection and security architecture integrating mobile agents, cloud infrastructure, and network applications to support trusted, decentralized logistics environments.

In conclusion, federated edge-cloud logistics architectures will progress from controlled pilot deployments to large-scale, multi-modal, and ecologically responsible systems over the course of the next ten years. These technologies can turn the logistics sector into a safe, flexible, and sustainable ecosystem by fusing the distributed intelligence described in this work with developments in 6G, blockchain, digital twins, and explainable AI [2, 12, 27]. This concept will lay the groundwork for the upcoming generation of autonomous logistics platforms and spur continuous innovation. The discussion's collective observations and the suggested research avenues demonstrate how federated edge-cloud systems have the potential to develop into a key technology for the logistics sector. The next step is to test these concepts through cross-industry collaboration and real-world deployments, even though the current work provides the architectural principles. The study's primary contributions are outlined in the part that follows, along with how these conclusions can direct future advancements in autonomous logistics.

5. Conclusion

This article introduced a federated edge-cloud architecture to help fix the problems of scalability, reactivity, and privacy that modern autonomous logistics systems have. The design improves global coordination and lets people make decisions in real time by combining collaborative learning across a multi-layered network with local intelligence at the edge. The proposed framework has specialised modules for dynamic routing, energy management, and robustness, which means that fleets of self-driving cars and drones can work well even when their surroundings change unexpectedly [2, 7]. The conversation made it clear that decentralisation greatly lowers latency, improves data security, and increases fault tolerance compared to traditional centralised logistics methods [3]. Federated learning lets fleets share knowledge without putting important operational data at risk. This means that systems can adapt to changing traffic patterns, energy needs, and other logistical challenges [10, 18]. When used together, these features could make regional and urban distribution networks smarter, more adaptable, and more environmentally friendly. Even though there are still technical problems to be solved, the proposed architecture sets the stage for the next generation of smart logistics networks. These systems will advance from conceptual ideas to real-world implementations as enabling technologies like digital twins, blockchain, and 6G connectivity develop, improving sustainability, resilience, and efficiency [12, 21]. Large-scale validation, industry-wide interoperability, and the incorporation of cutting-edge AI models for local and global decision-making will be the main areas of future research.

6. References

[1] Chellapandi, V. P., Yuan, L., & Brinton, C. G. (2023). Federated learning for connected and automated vehicles: A survey of existing approaches and challenges. IEEE Transactions on Intelligent Vehicles.

[2] Zhao, L., Cai, L., & Lu, W.-S. (2023). Transform-domain federated learning for edge-enabled IoT intelligence. IEEE Internet of Things Journal, 9(23), 6205–6216.

[3] Chen, Y., Wang, W., Yang, C., Liang, B., & Liu, W. (2025). An efficient energy management strategy of a hybrid electric UAV considering turboshaft engine speed regulation: A deep reinforcement learning approach. Applied Energy, 350, 125837.

[4] Tang, X., Chen, J., Qin, Y., Liu, T., & Yang, K. (2024). Reinforcement learning-based energy management for hybrid power systems: State-of-the-art survey, review, and perspectives. Chinese Journal of Mechanical Engineering, 37, 43.

[5] Ghorbani, B., Zendehboudi, S., & Saady, N. M. C. (2025). Advancing hybrid energy systems: A comprehensive review of processes and performance optimization. Energies, 18(6), 1443.

[6] Kuru, K. (2025). Swarms of autonomous drones in logistics within smart city: Opportunities, challenges and future directions. In Proc. 3rd Int. Conf. Logistics Engineering, Supply Chain and Digital Transformation.

[7] Liu, D., Liang, H., Zeng, X., Zhang, Q., & Zhang, Z. (2022). Edge computing application, architecture, and challenges in ubiquitous power internet of things. Frontiers in Energy Research, 10, 850252.

[8] Zhou, Z., Chen, X., Li, E., Zeng, L., Luo, K., & Zhang, J. (2019). Edge intelligence: Paving the last mile of artificial intelligence with edge computing. Proceedings of the IEEE, 107(8), 1738–1762.

[9] Li, J., Gu, C., Xiang, Y., & Li, F. (2022). Edge-cloud computing systems for smart grid: State-of-the-art, architecture, and applications. Journal of Modern Power Systems and Clean Energy, 10(4), 813–822.

[10] Quan, M. K., Pathirana, P. N., et al. (2025). Federated learning for cyber-physical systems: A comprehensive survey. IEEE Communications Surveys & Tutorials.

[11] Zhang, X., Wang, Y., Lin, J., & Chen, T. (2021). Federated learning for Internet of Things: Applications, challenges, and opportunities. arXiv preprint arXiv:2111.07494.

[12] Al-Quraan, M., Mohjazi, L., Bariah, L., Centeno, A., Zoha, A., & Imran, M. A. (2023). Edge-native intelligence for 6G communications driven by federated learning: Trends and challenges. IEEE Transactions on Emerging Topics in Computational Intelligence.

[13] Mohsen, B. M. (2024). AI-driven optimization of urban logistics in smart cities: Integrating autonomous vehicles and IoT for efficient delivery systems. Sustainability, 16(24), 11265.

[14] Borah, A., Paranjothi, A., & Thomas, J. P. (2025). A survey on distributed approaches for security enhancement in vehicular ad hoc networks. Computer Networks, 250, 110238.

[15] Priyanta, I. F., Freytag, J., & Schulte, P. (2023). Towards 6G-driven digital continuum in logistics. Logistics Journal: Proceedings, 2023.

[16] Min, W., Muthanna, M. S. A., Ibrahim, M., & Alkanhel, R. (2025). Privacy-preserving federated UAV data collection framework for autonomous path optimization in maritime operations. Applied Soft Computing.

[17] Shivaanivarsha, N., & Swetha, J. (2025). Securing intelligent vehicular networks using AI-driven federated learning. In Proc. IEEE Int. Conf. Computing and Communication.

[18] Paul, A., Singh, K., & Li, C. P. (2025). A multi-agent federated DRL model for vehicular task offloading in WPT-aided eROAD environment. IEEE Transactions on Intelligent Transportation Systems.

[19] Bouzoubia, R., Zaidi, S., Khamer, L., & Bourouis, A. (2024). Deep reinforcement learning-based resource allocation for the Internet of Vehicles: A systematic mapping study. Computer Science and Information Systems, Springer.

[20] Long, Q. (2025). Collaborative distributed machine learning: from knowledge reuse to sparsification in federated learning. PhD Thesis, University of Glasgow.

[21] Song, M., Lin, Y., Wang, J., Sun, G., & Dong, C. (2025). Trustworthy intelligent networks for low-altitude economy. IEEE Network.

[22] GK, S. K., Muniyal, B., & Rajarajan, M. (2025). Explainable federated framework for enhanced security and privacy in connected vehicles against advanced persistent threats. IEEE Open Journal of the Communications Society.

[23] Wu, Z., Zhang, H., Li, L., Lu, Y., & Sun, C. (2025). Mobile edge intelligence and computing with Star-RIS assisted intelligent autonomous transport system. IEEE Transactions on Vehicular Technology.

[24] Selvaraj, A. K., & Govindarajan, Y. (2025). Federated learning and digital twin-enabled distributed intelligence framework for 6G autonomous transport systems. IEEE Transactions on Intelligent Transportation Systems.

[25] Singh, G. (2025). Improving distributed intelligence solutions in 6G systems. PhD Thesis, University of Bologna.

[26] Mathieu, P., & De la Prieta, F. (2024). Advances in Practical Applications of Agents, Multi-Agent Systems, and Digital Twins. PAAMS 2024, Springer.

[27] Muthanna, M. S. A., et al. (2025). Autonomous maritime logistics with federated learning for ship traffic coordination. Applied Soft Computing.