

A Survey on Federated Learning and Edge Computing: Applications, Advantages, and Challenges [†]

Theodora Nevratiki ¹, Panagiotis Radoglou-Grammatikis ^{1,2} , Antonios Sarigiannidis ², Panagiotis Sarigiannidis ¹  and George F. Fragulis ^{1,*} 

¹ Department of Electrical and Computer Engineering, University of Western Macedonia, 50100 Kozani, Greece; ece01523@uowm.gr (T.N.); pradoglou@uowm.gr (P.R.-G.); psarigiannidis@uowm.gr (P.S.)

² Department of R&D, K3Y Ltd., 1700 Sofia, Bulgaria; asarigia@k3y.bg

* Correspondence: gfragulis@uowm.gr

[†] Presented at the 8th International Global Conference Series on ICT Integration in Technical Education & Smart Society, Aizuwakamatsu City, Japan, 20–26 January 2026.

Abstract

Federated learning (FL) and edge computing are transformative technologies that enhance privacy, efficiency, and real-time intelligence for distributed machine learning across diverse edge device networks. FL enables devices to collaboratively train models locally, so sensitive data remains on each device, ensuring privacy and regulatory compliance (like GDPR and HIPAA). Edge computing complements FL by bringing data processing closer to sources such as IoT sensors and smartphones, which reduces latency, bandwidth use, and dependence on cloud servers. This architecture is vital for smart cities, healthcare, industry, and autonomous systems, supporting real-time decision-making. The review details challenges such as resource heterogeneity, communication constraints, security risks, and management complexity, while highlighting opportunities for scalable orchestration, decentralized architectures, and blockchain integration. Together, FL and edge computing create a robust paradigm for scalable, privacy-aware distributed intelligence across multiple domains.

Keywords: federated learning; edge computing; edge AI; privacy-preserving machine learning; decentralized learning

1. Introduction

Edge computing is an advanced computing approach that handles and stores data closer to where it is generated, contrasting with the conventional reliance on distant, centralized data centers. This method responds to the increasing need for real-time applications, the explosion of connected devices, and the diverse requirements of sectors like smart cities, healthcare, manufacturing, and autonomous vehicles. Decentralized processing, a key concept in edge computing, means that data is computed locally on edge devices which include anything from IoT sensors and gateways to smartphones and industrial robots. Implementing edge computing reduces latency, lowers bandwidth consumption, boosts security, and facilitates real-time processing responsiveness. The combination of decentralized processing and various types of edge devices is revolutionizing how information is managed, analyzed, and acted upon across countless industries, making technology more immediate, efficient, and adaptive to local needs.

The central philosophy behind federated learning is to make machine learning possible where it used to be impossible: when data cannot leave its source due to privacy, security,



Academic Editors: Debopriyo Roy and Peter Ilic

Published: 20 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

policy, or ownership concerns. Federated learning enables each participant—such as a smartphone, sensor, or hospital—to train the model using its own local data, thereby eliminating the need for extensive and risky transfers of sensitive information to a central server. This setup has a profound impact in sectors like healthcare, finance, and edge computing, where isolated data is widespread. Decentralization also means resilience. Even if some clients drop offline or do not participate in a training round, the process continues with the others. This supports scalability and robust, real-world deployment.

2. Historical Development

2.1. Edge Computing

The development of edge computing dates back to 1998, when Akamai introduced Content Delivery Networks (CDNs). CDNs were created to lower network congestion and improve user access response times speed, mainly by emphasizing data backup and caching. However, edge computing later focused on function caching. The concept of function caching was initially implemented in 2005 to help with bandwidth and latency for personalizing mailbox management services [1,2].

The emergence of cloud computing can be traced back to 2006, when Amazon unveiled its “Elastic Computing Cloud.” The notion of “edge computing” was initially introduced by the Chief Executive Officer of Google in August 2006. While cloud computing was a powerful network service platform, the intense growth of data at the Internet of Things (IoT) led to issues like network bandwidth load, slow response times and security concerns. Therefore, these issues highlighted the need for a new computing paradigm, which lead to the advent of edge computing [3]. Edge computing has evolved as a new computing paradigm that executes calculations near the network’s edge, closer to the user and data source. In the meantime, various new technologies with similar principles but different approaches arose: Content Delivery Networks (CDNs): An internet-connected, resource-rich host at the network edge, acting as a middle layer between mobile services and the cloud [3,4]. Micro Data Centers (MCDs): like cloudlets, they were designed to complement the cloud by packaging computing, storage and networking equipment for applications needing lower latency. Mobile Edge Computing (MEC): An extension of the cloudlet concept, which allows processes to occur in base stations closer to mobile users; it is also known as multi-access edge computing. Fog computing (FC): In 2012, Cisco introduced a highly virtualized and distributed platform that migrates cloud tasks to network edge devices. In 2013, Ryan LaMothe formulated the modern “edge computing” term, describing it as the upstream of IoT services and the downstream of cloud services [1].

The European Telecommunications Standards Institute (ETSI) launched Mobile Edge Computing standardization in 2014, though the IBM and Nokia Siemens network had introduced MEC concepts in 2013. ETSI describes MEC as offering an “IT service environment and cloud computing capabilities at the edge of the mobile network, within the RAN and near mobile subscribers.” The Industry 4.0 era officially began in 2011, and it continues to the present day. During this era, edge computing became essential for enabling smart manufacturing through the Physical–Digital–Physical (PDP) loop, where data flows continuously between physical and digital worlds. This integration allowed for real-time analysis, reduced latency in industrial processes, and improved decision-making capabilities. Finally, the emergence of 5G technology significantly accelerated edge computing adoption. The performance metrics of 5G, which include data rates of up to 10 Gbps, service level latency of less than 1 ms, and an ultra-high reliability of 99.99999%, have made edge computing essential for meeting these stringent requirements. Recent developments focus on edge AI, which involves deploying artificial intelligence and machine learning models directly at the network’s edge. This enables real-time inference and decision-making without needing

cloud connectivity, benefiting applications in healthcare, autonomous vehicles, and smart cities [5].

2.2. Federated Learning

The concept of computing on encrypted data can be traced back to the early 1980s, a period concurrent with the development of cryptographic methods. In 2000, Agrawal and Srikant published work on privacy-preserving data mining, which was an early example of learning from local data with a central server while keeping privacy. Following that, in 2008, Vaidya et al. also contributed to this space with their work on privacy-preserving SVM classification. Over the past decade, leading up to the publication of some of these papers (around 2009–2019), solutions for privacy-preserving data analysis became much more widely deployed at scale. The term “federated learning” was first introduced in 2016. They described it as a way for a “loose federation of participating devices” to solve a learning task collaboratively, managed by a central server.

In 2017, McMahan et al. introduced the Federated Averaging algorithm, which is a common starting point for FL training, and Google AI Blog published an article titled “Federated learning: Collaborative machine learning without centralized training data” in April. The implementation of China’s cybersecurity law and the general principles of civil law began in 2017, emphasizing data protection and posing new challenges for traditional AI data processing. The European Union implemented the General Data Protection Regulations (GDPR) on 25 May 2018, which significantly impacted data privacy requirements, prohibiting model training without user permission and allowing users to delete private data [6]. In 2018, research emerged on applications like “Federated learning for mobile keyboard prediction” and “Applied federated learning: Improving Google keyboard query suggestions.”.

A notable event was the Workshop on Federated Learning and Analytics, which took place on 17–18 June 2019, at Google’s Seattle office, highlighting the growing need to survey open challenges in the field. In 2019, Apple began implementing cross-device FL in iOS 13 for applications such as the QuickType keyboard and voice classifiers for “Hey Siri,” a development that impacted hundreds of millions of users worldwide. Such adoption made FL visible as a real production technology [7]. The move beyond tech giants is shown by companies like doc.ai (health research) and Snips (voice assistants), exploring FL in healthcare, hotword detection, and other commercial challenges, proving FL’s cross-industry potential [8].

Federated learning started being used for more varied, real-world scenarios—hospitals sharing health insights without sharing patient data, banks collaborating on fraud detection, factories optimizing processes jointly, and more. Researchers started focusing on scaling FL to different settings: sometimes millions of end devices, sometimes a few large institutions. Surveys in this period included deep dives into privacy, optimization, and case study analysis in real deployments [9]. The importance of FL for IoT and edge computing was highlighted, showing its role in everything from smart homes to city infrastructure, all of which generate decentralized but valuable information [5].

Federated learning moved from early adoption to deeper integration within edge networks and IoT environments, especially as the number of heterogeneous devices grew. This shift was driven by the need to preserve data privacy in settings like smart cities, healthcare, and transportation, as well as to manage massive amounts of decentralized data [10,11]. Advances included more robust frameworks for edge deployment such as TensorFlow Federated, FedLab, and LEAF, with support for model pruning, compression, customized privacy strategies, and adaptation to device capabilities [7].

Research increasingly addressed the scaling of FL to support not just millions of devices, but also nuanced orchestration across highly diverse hardware (sensors, mobiles, servers, and embedded systems). Real-world deployments considered the energy and carbon footprint of large-scale collaborative ML.

As deployment increased, so did vulnerabilities. Research between 2022 and 2024 introduced advanced techniques for defending FL against model poisoning, backdoor attacks, and data leakage through gradient inversion [12]. Mutually verifiable updates, strong aggregation algorithms, and privacy-preserving techniques like differential privacy and secure aggregation protocols received increased attention and early implementation [13].

From 2022 onward, the focus increased on personalization and local adaptation so models better fit users, devices, or locations, despite non-IID data. Hierarchical, clustered, and federated multi-task learning protocols grew in prominence [9,14]. Federated learning further integrated continual learning, semi-supervised and unsupervised protocols, and “one-shot” approaches for quick adaptation with few labels [15].

FL in 2023–2024 expanded into collaborative domains where data privacy is non-negotiable, like finance, healthcare, wearable tech, autonomous vehicles, and smart manufacturing, often delivering significant practical improvements [16]. Benchmarks such as FedBNIST, MedMNIST v2 (for biomedical images), and collaborative health datasets became more common, and the research community rallied around public codebases and simulation platforms (FedERA, FedML, FedLab, TensorFlow Federated, LEAF) for reproducibility, comparison, and deployment.

3. Applications

3.1. Edge Computing

Edge computing is transforming how and where data is processed by bringing computation closer to the devices and sensors that generate it. This shift enables a wide range of innovative applications across multiple fields like smart cities, IIoT/manufacturing, autonomous vehicles, healthcare, energy/smart grids, agriculture, retail and public safety.

Edge computing is essential for the development of smart cities, as it allows real-time data processing and instant decision-making right at the data source, whether that is traffic sensors, surveillance cameras, or smart device infrastructure. The most frequent applications are the following:

1. **Traffic management:** Edge devices analyze real-time data from road sensors, cameras and connected vehicles to optimize traffic flow, manage congestion, and adjust traffic signals dynamically. This enables rapid response to incidents, reduces travel times and improves road safety [3].
2. **Public safety and surveillance:** Videos from city cameras are processed locally at the edge to detect anomalies such as accidents or suspicious activity, enabling quick alerts and responses while reducing bandwidth and central server workloads. Privacy is also better protected since raw video data does not always leave the local network [3].
3. **Environmental monitoring:** This involves installing sensors throughout the city to gather data on various environmental factors. These sensors track air quality, noise levels, weather conditions, and water quality. Edge computing enables quick analysis and alerts in case of pollution spikes or dangerous conditions, supporting public health efforts.

Edge computing is driving major innovation in industrial IoT (IIoT) and manufacturing by enabling machines, sensors, and production systems to process and respond to data locally without relying on remote cloud servers. This local processing empowers factories and industrial environments to achieve better efficiency, reliability, and intelligence. Key applications include:

1. Predictive maintenance: Edge devices constantly monitor machinery via vibration, temperature, and other sensors, to detect patterns indicating potential failures. They can predict maintenance needs, schedule repairs before breakdowns, and reduce costly unplanned downtime [1].
2. Real-time quality control: Vision systems and smart sensors at the edge inspect products as they move along the production line, instantly identifying defects or irregularities. This allows for immediate correction, reduces waste, and maintains high product quality [11].
3. Process automation and optimization: Edge computing supports autonomous decision-making on the spot. Systems can quickly adapt to changing conditions, optimize resource use, and control robotics or automation equipment with minimal delay, improving speed and flexibility in manufacturing operations.

Edge computing is fundamental to the functioning and advancement of autonomous vehicles because these vehicles must process enormous volumes of sensor data and make critical decisions in real time, often in environments where reliable connectivity to distant cloud servers cannot be guaranteed. The key applications of edge computing in autonomous vehicles are as follows:

1. Real-time perception and decision-making: Edge processors in vehicles combine data from cameras, radars and ultrasonic sensors to accurately interpret the vehicle's immediate surroundings. This allows rapid identification of obstacles, pedestrians, road signs and lane markings, tasks that demand immediate response and cannot tolerate cloud latency [2].
2. Vehicle to everything (V2X) communications: Edge computing powers low-latency communications between vehicles (V2V), infrastructure (V2I), and pedestrians (V2P). This supports applications like cooperative collision avoidance, traffic flow optimization, and hazard warnings.
3. Predictive maintenance: Edge devices monitor the status of vehicle components in real time, detecting anomalies and predicting part failures. Maintenance can then be scheduled proactively, reducing breakdowns while improving safety and reliability.

Edge computing is revolutionizing healthcare by enabling the rapid, local processing of sensitive patient data, improving responsiveness, enhancing privacy, and supporting novel medical applications [17]. Below are the main application areas of edge computing in healthcare:

1. Remote patient monitoring and wearables: Edge devices allow continuous and real-time data analysis. They analyze physiological data from wearables like the heart rate, ECG and blood glucose monitors in real time, detecting anomalies such as arrhythmia or sudden drops in vital signs and sending timely alerts to caregivers or health professionals [18].
2. Medical imaging and diagnostics: Hospitals and clinics can collaboratively train AI diagnostic models (e.g., for pneumonia or tumor detection using X-rays and MRIs) using federated learning at the edge. This approach preserves patient privacy and meets regulatory requirements by keeping sensitive data local [2,19–21].
3. Smart hospitals and clinical workflow automation: Edge platforms track medical equipment, monitor occupancy, and regulate building automation like heating/cooling or lighting for energy efficiency and infection control [13].

Edge computing is driving a new generation of energy and smart grid solutions by enabling real-time, decentralized data processing close to where energy is generated, distributed, and consumed. This revolution supports improved efficiency, reliability, sustainability, and security in modern power systems. The major applications are as follows:

1. Edge computing lets smart grids collect and process data from smart meters, sensors, and distributed energy resources instantly at the edge, not just in the central cloud. This enables immediate detection of power outages, equipment failures, or grid fluctuations, so operators can react in real time, keeping the energy system reliable [3].
2. Edge servers collect and manage inputs from solar panels, wind turbines, electric vehicles, and batteries at the grid's edge. They coordinate resource usage and sync with the central network for optimal performance [3].
3. Edge nodes analyze consumption patterns and can adjust loads dynamically, making it easier to schedule heavy-demand devices or electric vehicle charging when the grid is less busy. This supports demand response and smart pricing models, helping users and utilities save energy and money.

Edge computing has become a transformative technology for agriculture by connecting computing resources nearer to data sources and enabling real-time processing essential for modern farming operations. The primary applications of edge computing in agriculture include the following:

1. Edge computing powers real-time data collection and analytics from field sensors and IoT devices, helping farmers monitor soil conditions, moisture, weather, crop growth, and equipment status right at the field edge. Sensor data is processed locally for instant decisions like when to irrigate, fertilize, or spray pesticides without waiting for cloud responses, which means faster reactions and more efficient resource use. This leads to higher crop yields, water savings, and reduced environmental impact since inputs can be fine-tuned down to specific plants or zones [22].
2. Edge servers collect and analyze temperature, humidity, light, pH, and other environmental data in greenhouses, automatically controlling lighting, heating, ventilation, and watering. Local control at the edge helps maintain ideal growing conditions at all times, even if the internet connection drops, and can issue alerts or fix problems immediately. This improves crop quality, stabilizes yield, and saves resources by reducing unnecessary adjustments [3].
3. Edge-enabled wearable sensors monitor animal health, activity, and environment (like temperatures for calving cows or herd movements), then deliver instant alerts or health assessments to farmers' mobile devices. This ensures timely interventions for animal health issues, reduces risk of disease spread, and improves farming outcomes by monitoring pregnancies or stress in real time [9].

Edge computing has significantly influenced the retail industry, facilitating the relocation of computational resources to the point of sale and customer interaction. This development has enabled real-time processing and intelligent decision-making, thereby transforming the landscape of the retail sector. The main applications are as follows:

1. Edge computing allows shelves fitted with sensors and cameras to monitor product levels, detect when items are misplaced, and update inventory in real time. The analysis happens at the edge, so alerts or restocking requests are instant, even during network slowdowns. This reduces out-of-stock incidents, shrinks inventory loss, and saves staff time by automating manual checks [23].
2. Edge servers analyze data from in-store sensors, loyalty apps, and cameras to provide real-time, personalized promotions like digital displays that greet returning customers or suggest relevant products as they walk by. Moreover, edge devices can process customer preferences and purchase history locally to provide personalized recommendations and offers in real time. Digital signage systems powered by edge computing can display targeted advertisements based on customer demographics detected through computer vision [2].

3. Edge AI systems monitor shopper movement at checkout lanes using cameras and sensors, predicting waiting times and automatically opening or closing registers to minimize lines or redirecting customers to shorter lines. Touchless self-checkout can also use edge processing for fast payment and fraud detection [23].

Public safety

Edge computing has become a transformative technology for public safety, allowing real-time processing and intelligent decision-making directly at the network's edge where vital safety data is collected generated. The key applications in public safety include the following:

1. Edge computing enhances surveillance by allowing real-time video analysis directly at the camera level. Smart cameras with edge processing capabilities can handle object detection, facial recognition, and behavioral analysis locally, eliminating the need to send raw video streams to a central server [3].
2. Edge computing enables intelligent traffic management systems that enhance public safety through real-time traffic flow optimization and incident detection [2].
3. Edge AI at gates, building entrances, or even at airports can perform facial recognition or badge validation, authorize or deny entry within milliseconds, and trigger lockdowns or alerts if an unauthorized person is detected, while keeping sensitive biometric data local [13].

3.2. Federated Learning

Federated learning has become a revolutionary approach in machine learning, allowing multiple devices to collaboratively train models without sharing their data while preserving privacy. This technology is making significant impacts across multiple domains, offering unique advantages that are revolutionizing how machine learning tasks are approached and implemented. The most critical applications are in healthcare, IoT and smart devices, financial and insurance services, telecommunications and networking, natural language processing and autonomous vehicles and transportation [24–26].

Federated learning is having a significant impact in healthcare and medical research, enabling collaboration across institutions and devices while preserving patient privacy and data security. Below are the core application areas:

1. Predictive Modeling for Disease Risk and Outcomes: Federated learning is used by hospitals and research centers to collaboratively train models that predict patient outcomes like the risk of hospital readmission, disease onset (e.g., heart disease), or mortality without pooling sensitive electronic health records into a central database [27].
2. Medical Image Analysis: Federated learning enables MRI, CT, and X-ray scans from different institutions or countries to be used for collaborative training of image recognition models, improving diagnostics for cancer, lung disease, and neurological disorders [28,29].
3. Patient Similarity Search and Representation Learning: Federated learning allows computation of privacy-preserving hash codes or embeddings to find similar patients for research or treatment planning across multiple hospitals without ever exposing individual-level data [9].

IoT and smart devices

Federated learning offers transformative advances for IoT and smart device ecosystems, providing privacy-preserving, scalable, and efficient solutions for a wide array of applications.

1. Smart Home Devices and Personalization: Smart speakers, home assistants, and IoT-enabled appliances learn user preferences for voice commands, schedule automation,

and behavioral routines locally, only sharing model updates, not raw audio or sensor data, with cloud or edge servers [11].

2. **Industrial IoT and Predictive Maintenance:** Federated learning unlocks high-value collaborative analytics across industrial devices and factories, enabling actionable insights while keeping proprietary data secure. Sensors on industrial equipment, remote assets, and operational infrastructure use federated schemes to collaboratively train models for predictive maintenance, ensuring each plant or site retains sensitive operational logs [7].
3. **Anomaly Detection and Intrusion Detection:** Federated learning is used in IoT to collaboratively detect threats like cyberattacks or device tampering based on distributed patterns seen across the network. Each device contributes local alerts or patterns, but does not leak user specifics, improving safety and resilience [30].

Financial and Insurance Services

Federated learning is increasingly adopted in the finance industry to enable collaborative analytics and decision-making models across institutions—improving security, regulatory compliance, and intelligence while maintaining strict privacy.

1. **Fraud Detection and Prevention:** Multiple institutions collaboratively detect and block fraudulent patterns that span several organizations, like money laundering rings and coordinated digital payment attacks, without revealing sensitive client records. When institutions share only model improvements, they are less likely to miss sophisticated fraud that crosses traditional boundaries [9].
2. **Credit Scoring and Lending Risk:** Federated learning enables more accurate and fair credit scoring by combining knowledge from different banks, fintechs, or credit unions, even if their customers never overlap. This approach reduces bias for underbanked groups and creates much richer scoring models [7].
3. **Open Banking, Personalization, and Secure Data Collaboration:** Open banking is all about sharing account access and analytics with third parties. Federated learning makes it safer by letting banks, fintech apps, and aggregators build better recommendation engines, budgeting tools, and credit risk models while never revealing raw transactions [31].
4. **Anti-Money Laundering (AML) and Regulatory Compliance:** Banks and regulators use federated learning to work together, spotting patterns in money movement that indicate money laundering or terrorism financing, all while ensuring regulatory requirements for privacy and data residency are honored [30].
5. **Insurance Underwriting and Claims Analytics:** Federated learning gives insurers a way to train risk models on claims, accident, and customer data from across the industry—making the prediction of claim likelihood or fraud much more robust.

Federated learning is playing a growing role in telecommunications and networking by providing privacy-preserving, distributed intelligence for network optimization, management, security, and new service enablement. Following are the principal applications in this sector:

1. **Network Resource Allocation and Optimization:** Federated learning is used to distribute decision-making for network resource allocation, bandwidth scheduling, and congestion control across nodes or base stations without sharing sensitive or proprietary data [32]. Across 5G/6G, edge, and IoT networks, local nodes can train models on metrics like latency, throughput, energy, or user load and participate in global scheduling via the federated approach [30].
2. **Ultra-Reliable Low-Latency Communication (URLLC) and Vehicular Networks:** In vehicular and industrial IoT, federated learning models are trained to predict queueing

delays or schedule wireless channels, improving reliability and latency without central coordination [22].

3. Smart Caching and Content Delivery: Federated models are used to predict content popularity and manage proactive edge caching in telecoms, allowing for recommendations without network-wide data aggregation [30].

Natural Language Processing

Federated learning is increasingly prominent in natural language processing (NLP), delivering privacy-preserving collaborative models for text, speech, and multilingual applications across institutions and edge devices. Below are the major research use cases:

1. Keyboard Prediction and Query Suggestions: Federated learning enables real-time improvement of keyboard models like autocorrect, next-word prediction, and query suggestions, without sending raw keystrokes or messages to a central server. For example, on Google Gboard, smartphones process local data such as typed text sequences and only share encrypted model updates with a central server [11].
2. Speech Recognition and Keyword Spotting: Federated learning lets smart devices continually refine speech recognition or keyword spotting models on the device itself. Instead of uploading actual voice recordings, which might include names or sensitive info, devices process their data locally, then send only model weights or gradients [9].
3. Sentiment Analysis, Translation, and Multilingual NLP: Federated learning is important in organizations like banks, hospitals and global tech firms that have massive, private chat or text data but want to build robust models for sentiment analysis, translation, text summarization, or intent detection.

Federated learning is rapidly becoming a cornerstone technology for data-driven applications in autonomous vehicles, smart transportation systems, and mobility services. It enables collaborative machine learning across distributed entities (vehicles, infrastructure, edge nodes) without sharing raw data, thus addressing privacy, latency, and bandwidth challenges critical to transportation. Below the core applications are depicted:

1. Privacy-Preserving Data Sharing Across Vehicles: Autonomous vehicles constantly generate huge amounts of sensitive data, like driving behavior, routes, sensor footage, and interactions with other road users. Federated learning allows these vehicles to collaboratively train smarter models for perception, prediction and control, without ever sending raw data off the vehicle. Instead, vehicles only upload encrypted model updates, protecting users' locations and habits while pooling learning benefits across fleets, brands, or even city-wide systems [11].
2. Low-Latency, Real-Time Decision Making: In self-driving applications, latency is critical: vehicles must react instantly to dangers or changes on the road. Federated learning can train edge-deployed models at RSUs (Road Side Units), on-board units, or within platoons to deliver up-to-date, highly adaptive driving policies in real time.

4. Complementarity

Federated learning and edge computing converge naturally: edge brings compute and storage close to where data is generated, while FL brings collaborative model training to those distributed data silos without moving raw data. Together, they enable low-latency, privacy-preserving, scalable intelligence at the network edge. The main takeaway: EC provides the system substrate that makes FL practical at scale; FL provides the learning paradigm that makes EC privacy-preserving and data-efficient. Their integration unlocks robust edge AI across mobile, IoT, and vertical domains such as healthcare, industrial IoT, and vehicular systems [3].

How Edge Computing Enables Federated Learning

Edge computing platforms offer distributed computational resources at the “edge” of the network (base stations, routers, local servers, or even smart devices). This enables federated learning because local model training, updating, and inference can be performed directly on or near the data-generating devices, reducing the need for transmitting large data volumes to the cloud [33]. By processing and aggregating model updates locally, edge computing helps lower the latency, increase the real-time capability of federated learning, and save network bandwidth. This is crucial for applications needing immediate reactions, such as autonomous vehicles or smart healthcare [4].

How Federated Learning Enhances Edge Computing

Federated learning boosts privacy on edge computing platforms. Instead of moving sensitive raw data into the cloud or other centralized storage, data never leaves the edge device, and only model updates, often presented as gradients or weights, are shared. This approach aligns well with the privacy needs of edge-enabled applications, especially in sectors like healthcare, finance, and smart cities [13,34]. Federated learning allows the aggregation of collective intelligence across geographically distributed edge devices, offering improved model accuracy by learning from diverse data, while keeping data private [15]. Through federated learning, edge devices can continue to learn and adapt from local data distributions, enabling real-time adaptation and increased robustness to changing environments or non-IID data scenarios commonly seen in edge networks [35].

Technical Synergy in Architectures

In modern federated learning, cloud–edge–end architectures split the learning workload—edge devices handle local computation and aggregation, while the cloud manages more intensive aggregation, optimization, or storage. This hybrid approach enables real-time responsiveness and the capacity for massive distributed learning [22,33]. Some frameworks leverage hierarchical approaches (device–edge–cloud). Devices send updates to edge servers (first-level aggregation), which then pass to the cloud, balancing communication efficiency and scalability [7]. There is ongoing research on model compression, adaptive communication, and energy management to better suit the constraints of edge devices with federated learning, ensuring efficient, eco-friendly, and reliable collaboration.

Core intersections

1. Proximity and low latency: EC places computation at or near devices, enabling fast local training and inference, FL exploits that proximity to perform local update steps and occasional aggregation rather than frequent round trips to distant clouds [18,36].
2. Bandwidth relief: EC minimizes wide-area transfers while FL transmits model deltas instead of raw data, further reducing backbone load.
3. Privacy and compliance: EC reduces exposure windows; FL keeps data on-device and augments with secure aggregation and differential privacy for additional protections [30].

Both edge and FL face non-IID (non-independent and identically distributed) data as devices/users generate unique data patterns. Collaborative filtering, personalized federated models, and careful aggregation schemes (like FedAvg, FedProx, SCAFFOLD) are used to converge on robust, global models while accommodating local specialization [9,15].

5. Advantages and Challenges

5.1. Advantages

1. Privacy Preservation: Federated learning fundamentally protects user privacy by not transmitting raw data from devices to a central server. Instead, only model updates are shared, safeguarding sensitive information such as health records, personal images, or user behavior [13,34].

2. **Reduced Communication and Efficient Bandwidth Usage:** In FL, only the necessary model parameters or gradients are communicated between devices and the server, reducing data sent over bandwidth-limited or costly networks.
3. **Low Latency and Real-Time Processing:** Model inference and even some training happen locally, providing real-time analytics and decision-making critical for applications like autonomous vehicles, industrial IoT, and smart city infrastructure [5,34,37].

5.2. Challenges

Federated learning at the edge, when implemented in real-world environments, faces numerous deep-rooted and interdisciplinary challenges that span from system and network limitations to algorithmic, privacy, and societal concerns [38].

FL involves many rounds of exchanging model updates between edge devices and a central coordinator (or among devices in decentralized setups). Since models, like neural networks, can be quite large, these updates can quickly saturate bandwidth, especially in wireless or constrained networks [39]. The cost is even higher in environments with intermittent connections, where re-transmissions or delays further increase load [12].

Statistical Heterogeneity: Data at the edge is almost always non-IID. Devices like phones and IoT sensors collect user-specific data leading to updates that may conflict or have varying scales. This slows convergence and increases the risk that the global model will not generalize to all users [9].

System Heterogeneity: Devices greatly differ in processor power, RAM, battery, and network type from powerful laptops to low-power sensors. Some devices may drop out, go offline, or contribute irregularly, causing imbalanced and “straggler” effects in federated rounds.

Privacy and Security Threats: Although FL keeps raw data on the edge, information can still leak through shared model updates. Attacks like model inversion and property inference can reconstruct sensitive personal data from gradients or parameters [12].

6. Discussion—Research Gaps

6.1. Federated Learning

The less explored directions in federated learning (FL) cover several promising but under-researched topics that have the potential to significantly impact the efficiency, applicability, and trustworthiness of FL in practical deployments. The following are the major directions that require further exploration:

Federated learning depends on voluntary participation from clients, but motivating sustained contributions and ensuring high data quality is challenging. Open problems include: Designing fair, robust incentives to encourage participation from diverse clients, especially those contributing valuable or large datasets. Preventing free-riders (participants receiving model updates without meaningfully contributing data or computation). Building mechanisms for accountability, where untrustworthy or malicious contributors can be detected and removed without compromising privacy or fairness. Crafting robust economic models for participation, evaluating reputation systems, and developing verifiable contribution metrics are all areas open to both theoretical and practical research.

Most FL systems focus on learning a single global model, but participant diversity leads to varying data distributions and user needs. Personalized FL aims to deliver tailored models to each participant without sacrificing overall performance or privacy. Research gaps include the following: more scalable and practical multi-task learning algorithms suitable for resource-constrained decentralized environments; new theoretical foundations to understand privacy–personalization trade-offs; and tailored aggregation protocols that

blend global knowledge with individual adaptation. Many proposed solutions are yet to be standardized or widely validated in large-scale, real-world deployments.

Decentralized, Peer-to-Peer, Blockchain, and Hierarchical FL. Most FL relies on a central server, but truly decentralized peer-to-peer and blockchain-based FL aims to strengthen resilience, autonomy, and auditability: gossip-based, ring-based, and hierarchical topologies have been proposed but lack mature communication protocols and robust aggregation methods. Applying FL within cloud–fog–edge or multi-tiered architectures brings additional layers of communication and aggregation challenges, requiring new protocols for synchronization, load-balancing, and fairness.

Standardization and Realistic Benchmarks: There is an urgent need for realistic, widely accepted benchmarks, datasets, and standard evaluation criteria [28]: datasets that capture the true diversity (non-IID-ness, unbalanced quantity, device constraints) seen in real deployments.

6.2. Edge Computing

Less explored directions in edge computing include emerging challenges and opportunities where research remains in early stages or incomplete. These directions often involve integrating with advanced technologies, addressing newly arising limitations, and expanding edge computing paradigms beyond current mainstream applications. Most current edge computing systems rely on at least some central coordination. However, true peer-to-peer or decentralized approaches, where edge nodes manage themselves without a central authority, are still pretty rare. This area needs more real-world system designs and consensus/coordination protocols that can scale and cope with unreliable nodes. There is a lot of talk about energy-efficient algorithms, but studies of holistic energy management like integrating renewable energy, optimizing power distribution, and making edge sites “green” over time are just starting out. This includes work on balancing local generation/storage and dynamic task offloading for lower carbon footprint [2]. Research on how edge nodes can automatically discover each other, recover from failures, and adapt to changes without human intervention is far less mature than in the data center/cloud world. Topics like software-defined self-configuration, distributed monitoring, and trust-less leader election at scale are largely open. Delivering computation/services based on user/context/location is heavily referenced in edge hype, but actual mechanisms for seamless handoff, continuous context sensing, and privacy-aware “service always follows the user” designs remain barely explored at scale. Issues with handover, edge server migration, and network slicing using real mobility traces need attention [22]. Methods for tracking data origin, securing data at rest/in transit, and enforcing differentiated access and usage policies across multiple edge devices are not well explored, especially for dynamic, multi-tenant, and cross-operator environments.

Author Contributions: Conceptualization, T.N. and G.F.F.; methodology, T.N. and P.R.-G.; software, A.S.; validation, T.N., P.R.-G., A.S. and P.S.; data curation, T.N., P.R.-G., A.S. and P.S.; writing—original draft preparation, T.N. and G.F.F.; writing—review and editing, T.N., P.R.-G., A.S., P.S. and G.F.F.; supervision, G.F.F.; project administration, P.S. and G.F.F.; funding acquisition, P.R.-G., A.S. and P.S. All authors have read and agreed to the published version of the manuscript.

Funding: This project has received funding from the European Union’s Horizon Europe research and innovation program under grant agreement No. 101135930 (CoGNETs).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are available upon request from the corresponding author.

Conflicts of Interest: Authors Panagiotis Radoglou-Grammatikis and Antonios Sarigiannidis were employed by K3Y Ltd. The remaining authors declare no conflict of interest.

References

1. Nain, G.; Pattanaik, K.K.; Sharma, G.K. Towards edge computing in intelligent manufacturing: Past, present and future. *J. Manuf. Syst.* **2022**, *62*, 588–611. [\[CrossRef\]](#)
2. Hua, H.; Li, Y.; Wang, T.; Dong, N.; Li, W.; Cao, J. Edge Computing with Artificial Intelligence: A Machine Learning Perspective. *ACM Comput. Surv.* **2023**, *55*, 1–35. [\[CrossRef\]](#)
3. Cao, K.; Liu, Y.; Meng, G.; Sun, Q. An Overview on Edge Computing Research. *IEEE Access* **2020**, *8*, 85714–85728. [\[CrossRef\]](#)
4. Wang, X.; Han, Y.; Leung, V.C.M.; Niyato, D.; Yan, X.; Chen, X. Convergence of Edge Computing and Deep Learning: A Comprehensive Survey. *IEEE Commun. Surv. Tutor.* **2020**, *22*, 869–904. [\[CrossRef\]](#)
5. Yu, W. A Survey on the Edge Computing for the Internet of Things. *IEEE Access* **2018**, *6*, 6900–6919. [\[CrossRef\]](#)
6. Zhang, C.; Xie, Y.; Bai, H.; Yu, B.; Li, W.; Gao, Y. A survey on federated learning. *Knowl.-Based Syst.* **2021**, *216*, 106775. [\[CrossRef\]](#)
7. Khan, L.U.; Saad, W.; Han, Z.; Hossain, E.; Hong, C.S. Federated Learning for Internet of Things: Recent Advances, Taxonomy, and Open Challenges. *IEEE Commun. Surv. Tutor.* **2021**, *23*, 1759–1799. [\[CrossRef\]](#)
8. Li, L.; Fan, Y.; Tse, M.; Lin, K.-Y. A review of applications in federated learning. *Comput. Ind. Eng.* **2020**, *149*, 106854. [\[CrossRef\]](#)
9. Rahman, K.M.J.; Ahmed, F.; Akhter, N.; Hasan, M.; Amin, R.; Aziz, K.E. Challenges, Applications and Design Aspects of Federated Learning: A Survey. *IEEE Access* **2021**, *9*, 124682–124700. [\[CrossRef\]](#)
10. Liu, J. From distributed machine learning to federated learning: A survey. *Knowl. Inf. Syst.* **2022**, *64*, 885–917. [\[CrossRef\]](#)
11. Janaki, G.; Umanandhini, D. Federated Learning Approaches for Decentralized Data Processing in Edge Computing. In *Proceedings of the 2024 5th International Conference on Smart Electronics and Communication (ICOSEC), Trichy, India, 18–20 September 2024*; IEEE: New York, NY, USA, 2024; pp. 513–519. [\[CrossRef\]](#)
12. Neto, H.N.C.; Hribar, J.; Dusparic, I.; Mattos, D.M.F.; Fernandes, N.C. A Survey on Securing Federated Learning: Analysis of Applications, Attacks, Challenges, and Trends. *IEEE Access* **2023**, *11*, 41928–41953. [\[CrossRef\]](#)
13. Syu, J.-H.; Lin, J.C.-W.; Srivastava, G.; Yu, K. A Comprehensive Survey on Artificial Intelligence Empowered Edge Computing on Consumer Electronics. *IEEE Trans. Consum. Electron.* **2023**, *69*, 1023–1034. [\[CrossRef\]](#)
14. Ding, J.; Tramel, E.; Sahu, A.K.; Wu, S.; Avestimehr, S.; Zhang, T. Federated Learning Challenges and Opportunities: An Outlook. In *Proceedings of the ICASSP 2022–2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Singapore, 23–27 May 2022*; IEEE: New York, NY, USA, 2022; pp. 8752–8756. [\[CrossRef\]](#)
15. Yu, B.; Mao, W.; Lv, Y.; Zhang, C.; Xie, Y. A survey on federated learning in data mining. *WIREs Data Min. Knowl.* **2022**, *12*, e1443. [\[CrossRef\]](#)
16. Hu, G.; Teng, Y.; Wang, N.; Han, Z. Faster Convergence on Heterogeneous Federated Edge Learning: An Adaptive Clustered Data Sharing Approach. *IEEE Trans. Mob. Comput.* **2025**, *24*, 5342–5356. [\[CrossRef\]](#)
17. Kaloforidis, N.; Kollias, K.-F.; Radoglou-Grammatikis, P.; Sarigiannidis, P.; Fragulis, G.F. Autism Spectrum Disorder Classification in Children Using Eye-Tracking Data and Machine Learning. *Eng. Proc.* **2025**, *107*, 12. [\[CrossRef\]](#)
18. Sathesh, M.; Ramakrishnan, K.; Raja, M.; Kalaiarasi, K.; Balamurugan, M. Edge Computing Integration in IoT Networks for Real-Time Data Processing. In *Proceedings of the 2024 International Conference on Cybernation and Computation (CYBERCOM), Dehradun, India, 15–16 November 2024*; IEEE: New York, NY, USA, 2024; pp. 585–590. [\[CrossRef\]](#)
19. Choumpaev, A.; Moysis, L.; Lawnik, M.; Fragulis, G. A Generalized Chaotic Neural Network Model for Epilepsy Using Soboleva Hyperbolic Tangent Functions. In *2025 14th International Conference on Modern Circuits and Systems Technologies (MOCAST)*; IEEE: New York, NY, USA, 2025; pp. 1–4.
20. Kafetzis, I.; Hann, A.; Fragulis, G.F. Efficient Selection of Rare Pathology Samples from Unlabeled Medical Data via Deep Active Learning. In *2025 14th International Conference on Modern Circuits and Systems Technologies (MOCAST)*; IEEE: New York, NY, USA, 2025; pp. 1–4.
21. Moysis, L.; Lawnik, M.; Kollias, K.; Baptista, M.; Goudos, S.; Fragulis, G. Dynamic analysis of a generalized attention deficit disorder model with Soboleva activation functions. *Chaos An. Interdiscip. J. Nonlinear Sci.* **2025**, *35*, 083105. [\[CrossRef\]](#)
22. Khan, L.U.; Yaqoob, I.; Tran, N.H.; Kazmi, S.M.A.; Dang, T.N.; Hong, C.S. Edge-Computing-Enabled Smart Cities: A Comprehensive Survey. *IEEE Internet Things J.* **2020**, *7*, 10200–10232. [\[CrossRef\]](#)
23. Kolapo, R.; Kawu, F.M.; Abdulmalik, A.D.; Edem, U.A.; Young, M.A.; Mordi, E.C. Edge computing: Revolutionizing data processing for IoT applications. *Int. J. Sci. Res. Arch.* **2024**, *13*, 023–029. [\[CrossRef\]](#)
24. Papadopoulos, C.; Kollias, K.-F.; Fragulis, G.F. Recent advancements in federated learning: State of the art, fundamentals, principles, IoT applications and future trends. *Future Internet* **2024**, *16*, 415. [\[CrossRef\]](#)
25. Siniosoglou, I.; Argyriou, V.; Lagkas, T.; Moscholios, I.; Fragulis, G.; Sarigiannidis, P. Unsupervised bias evaluation of dnns in non-iid federated learning through latent micro-manifolds. In *IEEE INFOCOM 2022—IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*; IEEE: New York, NY, USA, 2022; pp. 1–6.

26. Siniosoglou, I.; Argyriou, V.; Fragulis, G.; Fouliras, P.; Papadopoulos, G.T.; Lytos, A.; Sarigiannidis, P. Applied federated model personalization in the industrial domain: A comparative study. *IEEE Open J. Commun. Soc.* **2024**, *6*, 3192–3210. [\[CrossRef\]](#)
27. Makris, I.; Lytos, A.; Kyranou, K.; Argyriou, V.; Lagkas, T.; Kollias, K.-F.; Fragoulis, G.F.; Sarigiannidis, P. Detecting personal protective equipment (PPE) utilising YOLOv8 in a federated learning environment. In *AIP Conference Proceedings*; AIP Publishing: Melville, NY, USA, 2024.
28. Shaheen, M.; Farooq, M.S.; Umer, T.; Kim, B.-S. Applications of Federated Learning; Taxonomy, Challenges, and Research Trends. *Electronics* **2022**, *11*, 670. [\[CrossRef\]](#)
29. Nevratiki, T.; Iliadou, A.; Ntolkeras, G.; Sfakianakis, I.; Lazaridis, L.; Maraslidis, G.; Asimopoulos, N.; Fragulis, G.F. A survey on federated learning applications in healthcare, finance, and data privacy / data security. In *AIP Conference Proceedings*; AIP Publishing LLC: Melville, NY, USA, 2023; p. 120015.
30. Lim, W.Y.B.; Luong, N.C.; Hoang, D.T.; Jiao, Y.; Liang, Y.-C.; Yang, Q.; Niyato, D.; Miao, C. Federated Learning in Mobile Edge Networks: A Comprehensive Survey. *IEEE Commun. Surv. Tutor.* **2020**, *22*, 2031–2063. [\[CrossRef\]](#)
31. Yuan, L.; Wang, Z.; Sun, L.; Yu, P.S.; Brinton, C.G. Decentralized Federated Learning: A Survey and Perspective. *IEEE Internet Things J.* **2024**, *11*, 34617–34638. [\[CrossRef\]](#)
32. Qiu, T.; Chi, J.; Zhou, X.; Ning, Z.; Atiquzzaman, M.; Wu, D.O. Edge Computing in Industrial Internet of Things: Architecture, Advances and Challenges. *IEEE Commun. Surv. Tutor.* **2020**, *22*, 2462–2488. [\[CrossRef\]](#)
33. Haibeh, L.A.; Yagoub, M.C.E.; Jarray, A. A Survey on Mobile Edge Computing Infrastructure: Design, Resource Management, and Optimization Approaches. *IEEE Access* **2022**, *10*, 27591–27610. [\[CrossRef\]](#)
34. Cp, V.; Ba, V.; Zulfaquar, S. Edge Computing Assisted Pothole Detection and Lane Identification Augmented by Federated Learning. In *Proceedings of the 2024 International Conference on IoT Based Control Networks and Intelligent Systems (ICICNIS)*, Bengaluru, India, 17–18 December 2024; pp. 1549–1552. [\[CrossRef\]](#)
35. Li, L.; Zhu, L.; Li, W. Cloud–Edge–End Collaborative Federated Learning: Enhancing Model Accuracy and Privacy in Non-IID Environments. *Sensors* **2024**, *24*, 8028. [\[CrossRef\]](#) [\[PubMed\]](#)
36. Boruga, D.; Bolintineanu, D.; Racates, G.I. Federated learning in edge computing: Enhancing data privacy and efficiency in resource-constrained environments. *World J. Adv. Eng. Technol. Sci.* **2024**, *13*, 205–214. [\[CrossRef\]](#)
37. Mughal, F.R. Adaptive federated learning for resource-constrained IoT devices through edge intelligence and multi-edge clustering. *Sci. Rep.* **2024**, *14*, 28746. [\[CrossRef\]](#) [\[PubMed\]](#)
38. Wang, S.; Tuor, T.; Salonidis, T.; Leung, K.K.; Makaya, C.; He, T.; Chan, K. Adaptive Federated Learning in Resource Constrained Edge Computing Systems. *arXiv* **2019**, arXiv:1804.05271. [\[CrossRef\]](#)
39. Zhang, K.; Song, X.; Zhang, C.; Yu, S. Challenges and future directions of secure federated learning: A survey. *Front. Comput. Sci.* **2022**, *16*, 165817. [\[CrossRef\]](#) [\[PubMed\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.