information

Review

Exploring Federated Learning Tendencies Using a Semantic Keyword Clustering Approach

Francisco Enguix [1,*], Carlos Carrascosa [1] and Jaime Rincon [2]

[1] Valencian Research Institute for Artificial Intelligence (VRAIN), Universitat Politècnica de València (UPV), 46022 Valencia, Spain; carrasco@dsic.upv.es
[2] Departamento de Digitalización, Escuela Politécnica Superior, Universidad de Burgos, 09006 Miranda de Ebro, Spain; jarincon@ubu.es
* Correspondence: fraenan@upv.es

Abstract: This paper presents a novel approach to analyzing trends in federated learning (FL) using automatic semantic keyword clustering. The authors collected a dataset of FL research papers from the Scopus database and extracted keywords to form a collection representing the FL research landscape. They employed natural language processing (NLP) techniques, specifically a pre-trained transformer model, to convert keywords into vector embeddings. Agglomerative clustering was then used to identify major thematic trends and sub-areas within FL. The study provides a granular view of the thematic landscape and captures the broader dynamics of research activity in FL. The key focus areas are divided into theoretical areas and practical applications of FL. The authors make their FL paper dataset and keyword clustering results publicly available. This data-driven approach moves beyond manual literature reviews and offers a comprehensive overview of the current evolution of FL.

Keywords: federated learning; analysis; review; multi-agent system (MAS)

Citation: Enguix, F.; Carrascosa, C.; Rincon, J. Exploring Federated Learning Tendencies Using a Semantic Keyword Clustering Approach. *Information* **2024**, *15*, 379. https://doi.org/10.3390/info15070379

Academic Editor: Peter Z. Revesz

Received: 7 May 2024
Revised: 18 June 2024
Accepted: 26 June 2024
Published: 28 June 2024

1. Introduction

Federated learning (FL) has emerged as a revolutionary paradigm in collaborative machine learning [1]. It empowers multiple devices or institutions to train a model while collectively safeguarding data privacy. This decentralized approach contrasts traditional methods where data are centralized for model training, potentially compromising user privacy and data ownership. FL accomplishes this collaborative learning by keeping raw data distributed on individual devices, and instead of sharing the raw data, participants exchange the model updates.

The field of FL is experiencing explosive growth, leading to a vast and ever-expanding body of research literature. This presents a significant challenge to researchers attempting to identify current trends and emerging sub-areas within FL. Traditional manual literature reviews with a global approach, while valuable, become increasingly impractical for analyzing field trends as the number of publications and the intricate interplay of FL concepts continue to grow exponentially, as depicted in Figure 1. To address this challenge, this paper proposes the use of an automated semantic keyword clustering technique as a critical tool for analyzing FL research trends.

Automated semantic keyword clustering leverages advanced natural language processing (NLP) techniques to extract meaningful data from the vast amount of interconnected areas in FL. Using pre-trained transformer models [2], the research article keywords can be transformed into dense vector spaces that capture their semantic relationships. This empowers the creation of clusters based on thematic relevance, revealing the underlying thematic structure of the FL research landscape.

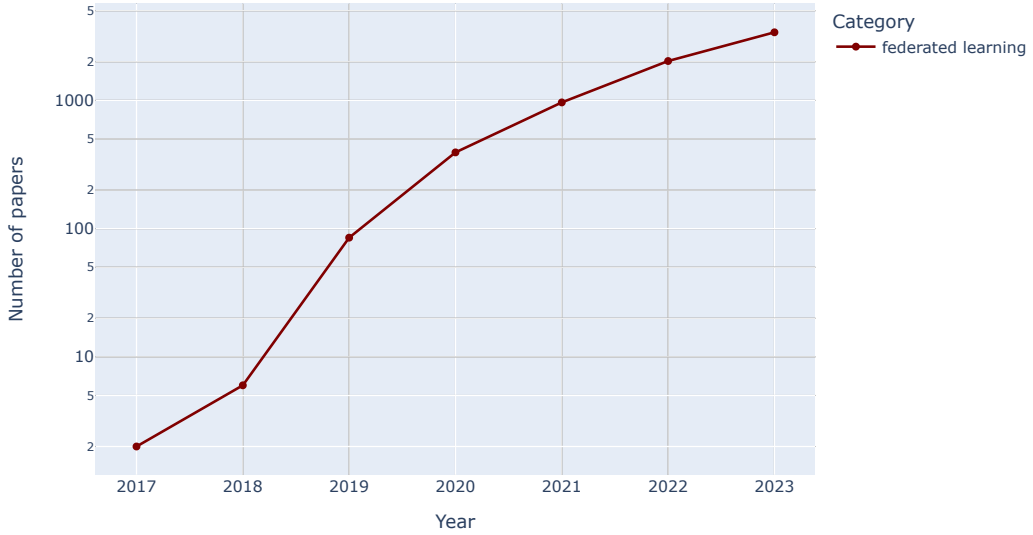

Figure 1. Number of papers containing the keyword "federated learning" across the years, on a logarithmic scale on the y-axis, based on the public dataset presented in Section 2.

This paper presents a semantically-based literature analysis of the 7 953 papers about FL. The primary objective is to uncover and explore the major theoretical categories and practical application areas of FL and examine the current trends of the field, along with the emerging sub-areas that have received less research attention. First, we formulate a series of research questions (RQs) that guide the investigation. These RQs delve into the current trends in FL (RQ1), the tendencies of these trends (RQ2), and the application domains where FL techniques are finding utility (RQ3 and RQ4). Recognizing the potential in under-explored areas, we propose additional research questions (RQ5 and RQ6) that focus on identifying emerging sub-areas within FL that have received limited research focus, and investigating how existing FL techniques can be adapted to address these application domains. The final question (RQ7) looks ahead to predict potential future directions and areas of growth. Formally, we formulated the following research questions:

RQ1: What are the current trends in FL?
RQ2: What are the tendencies of the current trends in FL?
RQ3: What are the application domains where FL techniques are applied?
RQ4: What are the tendencies of the application domains?
RQ5: What are the emerging sub-areas within FL?
RQ6: What are the tendencies of the emerging sub-areas?
RQ7: What are the potential future trends of FL?

A data-mining technique and a transformer-based semantic analysis of the literature's keywords will be employed to address these RQs and uncover the trends and tendencies within this extensive collection. This approach permits automatically grouping keywords into clusters, revealing the thematic relationships and dominant topics within the current body of FL research.

The Structure of the Survey

The structure of this survey is designed to address the research questions and present the findings. Figure 2 provides a classification scheme outlining these categories. Then, we will delve deeper into each category and explore the relevant advancements from the existing literature. Following the introduction, this paper unfolds across several key sections:

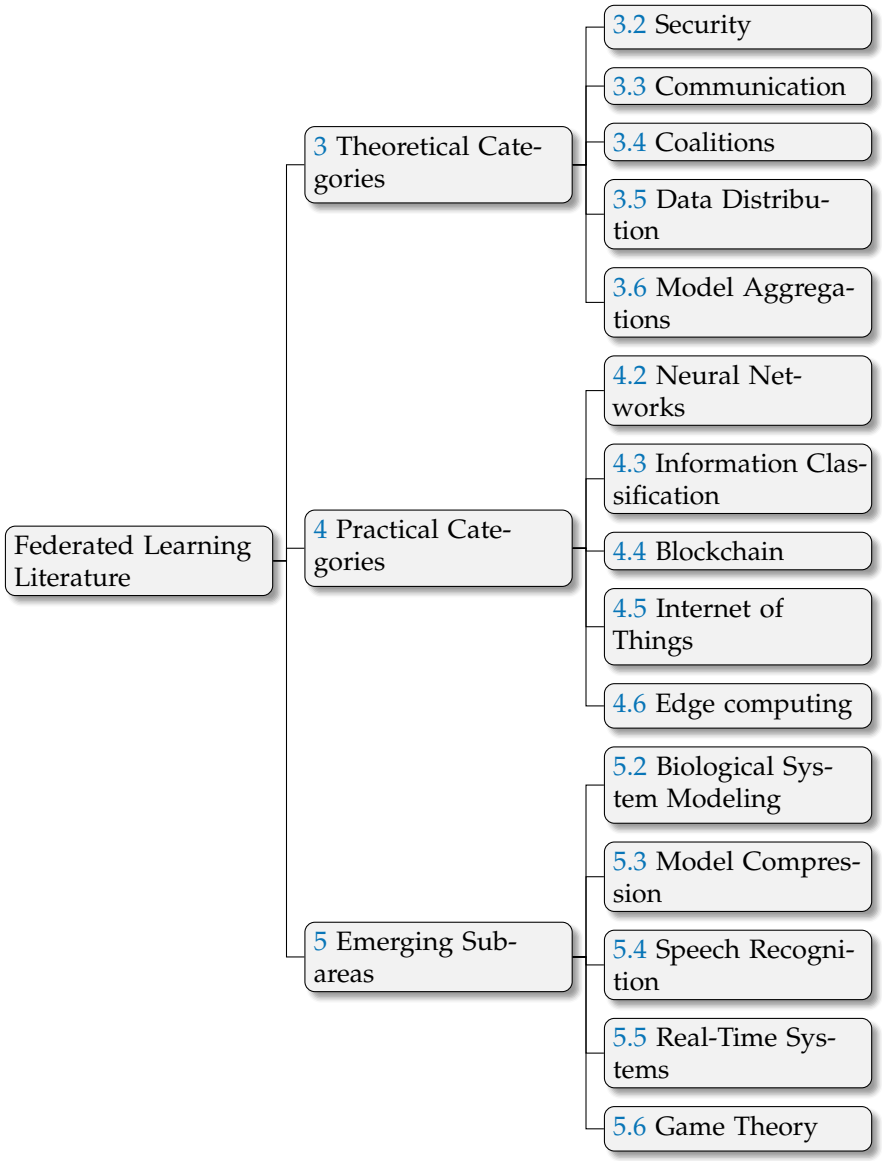

Figure 2. Taxonomy of this paper.

- Research Method (Section 2). This section delves into the approach employed to analyze trends and sub-areas within FL. It details the utilization of keyword extraction and automated clustering techniques to gain insights from the vast FL research landscape.
- Theoretical Categories (Section 3). Here, we present a detailed analysis of the key theoretical areas of FL. This section explores crucial aspects such as security mechanisms, communication protocols, coalition formation, data distribution strategies, and model aggregation techniques.
- Practical Categories (Section 4). Shifting the focus to the practical applications of FL, this section examines its implementation in various domains. We will explore how FL empowers neural networks, facilitates information classification tasks, integrates with blockchain technology, and finds applications in the Internet of Things (IoT) and edge computing environments.
- Emerging Sub-Areas (Section 5). This section explores the sub-areas of FL research that have emerged as a result of previous research directions. Here, we will identify and analyze these emerging trends that hold significant promise for the future development of the field, including biological system modeling, model compression techniques, advancements in speech recognition, the application of FL to real-time systems, and the utilization of game theory for improved performance.

- Conclusion (Section 6). Building upon the foundation in the preceding sections, this section will synthesize the key findings. It will address the research questions and explore potential future research directions of FL.

2. Research Method

This study aims to analyze the current trends and sub-areas within the field of FL while examining the tendencies over the years. We leverage a data-driven approach that utilizes keyword extraction and automated clustering techniques to achieve this.

Our analysis begins by collecting a comprehensive dataset of research papers from the Scopus database. We employ a query to identify relevant publications of FL in computer science that are written in English. The exact query we used in advance Scopus searcher is: TITLE-ABS-KEY ("federated learning") AND (LIMIT-TO (EXACTKEYWORD, "Federated Learning")) AND (LIMIT-TO (SUBAREA, "COMP")) AND (LIMIT-TO (LANGUAGE, "English")).

The result of this query, on 15 April 2024, reveals 7953 results without counting the 11 duplicated papers. Subsequently, we extract the keywords from each paper, forming a collection of 22,841 unique keywords that represent the research landscape in FL. Then, to uncover the underlying thematic structure within this keyword collection, we turn to Natural Language Processing (NLP) techniques. We employ a pre-trained transformer model, specifically the `all-mpnet-base-v2` model, to convert each keyword into a 768-dimensional dense vector space. We used the `all-mpnet-base-v2` transformer because it is trained for a total number of sentence pairs above 1 billion sentences (https://huggingface.co/sentence-transformers/all-mpnet-base-v2, accessed on 26 April 2024) and this corpus includes the Semantic Scholar Open Research Corpus (S2ORC), which is a general-purpose corpus for NLP and text mining research over scientific papers [3]. In addition, the `all-mpnet-base-v2` model has the best average performance between the performance of sentence embeddings and the performance of semantic search, over all the Hugging Face pre-trained sentence transformers models (https://www.sbert.net/docs/pretrained_models.html, accessed on 26 April 2024).

These embeddings capture the semantic relationships between keywords, allowing us to group them based on their semantic meaning. We perform agglomerative clustering on the vector embeddings to identify the major thematic trends and sub-areas. This clustering algorithm starts with each keyword as an individual cluster and iteratively merges the most similar clusters based on a distance metric. In this case, we utilize the Euclidean metric to measure the distance between cluster centroids and Ward's linkage to determine the optimal merging strategy. We used the Euclidean distance because effectively captures the inherent semantic relationships among the keywords, ensuring that the clustering process reflects true semantic groupings [4]. Moreover, the Euclidean distance is computationally efficient, facilitating the iterative process of agglomerative clustering, which involves repeated distance calculations between clusters. The final number of clusters, set at 100, provides a granular view of the thematic landscape while maintaining a manageable number of groups for analysis.

By examining the keywords within each cluster, we can identify the key thematic trends and sub-areas that are currently shaping the field of FL. In Table 1 are shown the number of papers of five keyword groups, over the years, of each category presented on this paper. The number of papers, over all the years, of the keywords groups can be found in Tables A1 and A2. This novel data-driven approach allows us to move beyond manual literature reviews and capture the broader dynamics of research activity within the domain. We can then delve deeper into specific clusters to understand the research questions, methodologies, and potential applications that are driving the current evolution of FL.

We made the FL paper dataset public and the keyword clustering results. You can find those files under the following public GitHub repository: https://github.com/FranEnguix/datasets/tree/main/2024%20FL%20Tendencies (accessed on 26 April 2024).

Table 1. The number of papers over the years of the selected keyword groups.

Category	Total	2017	2018	2019	2020	2021	2022	2023	2024
communication	1110	1	1	20	84	182	300	382	140
security	1076	1	1	7	43	110	255	498	161
coalition	942	1	1	8	77	104	297	355	99
data distribution	671	0	0	6	27	72	170	297	99
model aggregations	574	0	0	5	34	92	139	232	72
neural networks	2592	2	3	28	137	327	657	1097	341
classification (of information)	1292	0	1	10	65	172	321	536	187
blockchain	1281	1	0	21	68	147	340	515	189
Internet of Things	1262	0	1	12	53	116	328	541	211
edge computing	1142	0	0	16	73	158	325	417	153
biological system modeling	288	0	0	0	5	26	59	140	58
model compression	277	0	0	3	18	43	58	109	46
speech recognition	273	0	0	1	26	30	84	99	33
real-time systems	241	0	1	5	18	35	53	94	35
game theory	232	0	0	6	16	23	57	90	40

3. Main Theoretical Categories

This section dissects the research landscape by analyzing the publication trends within the following core theoretical areas: security, communication, coalitions, data distribution, and model aggregation. Our analysis, presented in the following subsections, leverages a data-driven trend analysis approach examining the yearly publication volume across these categories. Subsequently, we will present each category, highlighting the novel advances in each sub-area.

3.1. Data Analysis

While the current main sub-areas of FL started with just a handful of publications in 2017 and 2018, there has been a steady rise across all categories, with a sharp increase from 2019 onward, as depicted in Figure 3. This growth highlights the growing interest in FL as a method to collaboratively train ML models without compromising data privacy. Notably, as Figure 4 exposed, the category of "security" shows the most significant rise, reflecting a growing focus on addressing potential vulnerabilities in FL systems. Interestingly, "communication" research, though increasing, has not grown at the same exponential rate as other categories. This suggests that researchers might be prioritizing core security and privacy challenges over delving deeper into optimizing communication efficiency in FL. Overall, the data indicate a maturing field of FL research with a focus on building robust and secure systems for collaborative ML.

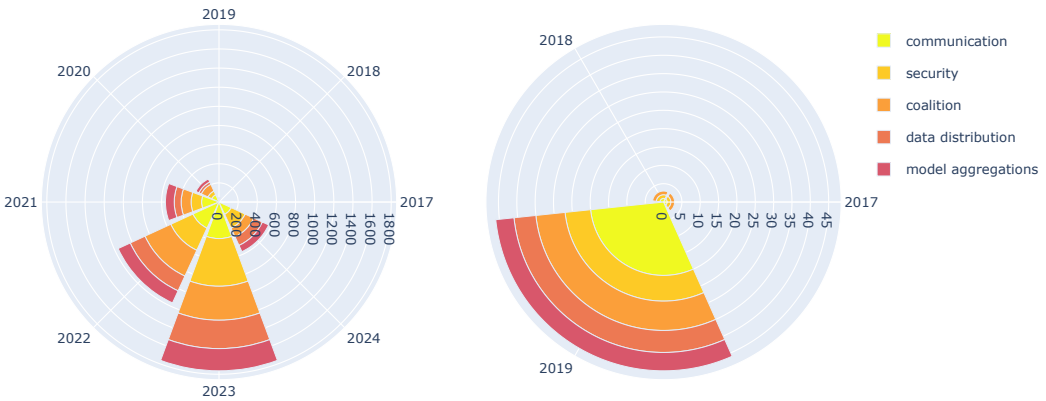

Figure 3. All theoretical keyword category groups over the years.

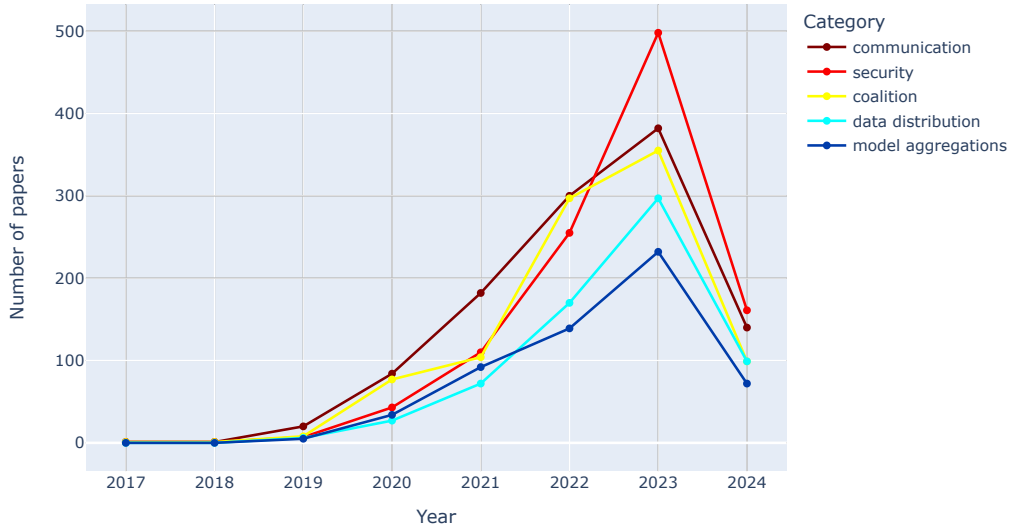

Figure 4. Tendencies of the selected keyword groups over the years.

3.2. Security

FL offers a compelling solution for collaborative machine learning while safeguarding data privacy. However, its core strength—keeping data distributed across devices—also presents a significant security challenge. The FL security research addresses these challenges through a multi-pronged approach, focusing on protecting both model parameters and the underlying data.

3.2.1. Model Inversion Attacks

One major concern is model inversion attacks. In these attacks, malicious participants attempt to reconstruct the training data used to build the model by analyzing the model updates exchanged during the FL process [5–8]. Researchers are developing differential privacy techniques to address this [9,10]. Differential privacy injects controlled noise into model updates, making it statistically impossible to infer any information about individual data points used for training. This technique provides strong data privacy for participants [11].

3.2.2. Poisoning Attacks

Another security threat involves poisoning attacks. Here, malicious actors attempt to manipulate the training process by injecting poisoned data or updates. This can lead to a degraded or biased model [12].

Data Poisoning Attacks

In these attacks, malicious actors inject tampered data points into the training process, aiming to manipulate the FL model to their advantage. These points are designed to mislead the FL model during training, forcing it to learn incorrect patterns or biased outputs that benefit the attacker.

Since FL relies on local participant updates, it can be challenging to detect poisoned data points, especially if they are disguised. Additionally, the distributed nature of FL makes it difficult to pinpoint the source of the attack. Also, there is a novel approach that directly inverts the loss function, generating strong malicious gradients at each training iteration to push the model away from the optimal solution [13].

Techniques like outlier detection algorithms can identify suspicious data points during aggregation [14]. Additionally, robust aggregation methods that down-weight or eliminate extreme updates can further reduce the impact of poisoned data [15].

Model Poisoning Attacks

Unlike data poisoning attacks that focus on corrupting data points, model poisoning attacks target the model updates exchanged during FL. Malicious participants can contribute strategically modified model updates that steer the global model in the desired direction.

Successful model poisoning attacks can cause the global model to learn faulty patterns or biased outputs. This can lead to inaccurate predictions, hindering the functionality of the FL system and potentially causing harm depending on the application. A novel technique named the model shuffle attack (MSA) introduces a unique method of shuffling and scaling model parameters. While the attacker's model appears accurate during testing, it secretly disrupts the training of the global model [16]. This sabotage can significantly slow convergence or even prevent the global model from learning effectively.

Several approaches can help mitigate model poisoning attacks. Cryptographic techniques like SMPC combined with blockchain [17] can be employed to prevent participants from directly observing the model updates, making it harder to inject malicious modifications. Additionally, federated Byzantine fault tolerance (Byzantine-FL) protocols can identify and exclude unreliable or malicious participants from the training process [15], safeguarding the integrity of the federated model.

3.2.3. Membership Inference Attacks

These attacks attempt to determine whether a specific data point belongs to a particular participant's dataset that contributes to the FL training process.

Attackers can potentially infer membership by analyzing the model's predictions on strategically crafted data points. If the model's behavior deviates significantly for a certain input compared to the general prediction pattern, it might indicate the presence of that data point in a participant's training dataset. The PAPI attack is a novel poisoning-assisted property inference attack that targets properties of the training data that are not directly relevant to the model's purpose [18]. By strategically manipulating data labels, a malicious participant can leverage updates to the central model to infer these sensitive properties, even from benign participants.

Existing works have proposed homomorphic encryption and secure multiparty computation (SMC) to address this issue, but these approaches do not apply to large-scale systems with limited computation resources. Differential privacy methods inject noise into the model updates during training, making it statistically harder to link specific data points to participants. Still, it brings a substantial trade-off between privacy budget and model performance. A novel FL framework, based on the computational Diffie–Hellman (CDH) problem to encrypt local models, safeguards against inference attacks [19]. The framework achieves this with minimal impact on model accuracy and computational/communication costs and eliminates the need for secure pairwise communication channels.

3.2.4. Backdoor Attacks

A backdoor attack is a malicious attempt to manipulate a model during training. This is achieved by introducing triggers embedded in the training data. When a sample containing this trigger is fed to the model, it will be misled into producing a specific, attacker-defined output, while functioning normally for all other data. These attacks can be untargeted, aiming to simply degrade the model's overall performance, or targeted, aiming to force the model to misclassify specific triggered samples into a particular category [20].

The attacker achieves this by poisoning the training data. Pairs of data points are created: one being the original training sample and its correct label, and another being the same sample altered with the backdoor trigger and a desired, potentially incorrect, label. The attacker can manipulate the model's learning process by strategically including these poisoned pairs in a small portion of the training data without raising major red flags. This way, the backdoor becomes embedded in the final model, causing it to malfunction when encountering the specific trigger.

The attacks occur during the training phase and rely on a universal trigger that can be added to any sample to activate the backdoor functionality. Backdoor attacks can be particularly concerning as they can bypass standard privacy-preserving techniques in FL. An attacker might steer the model's predictions toward a specific outcome, as they are designed to be subtle and difficult to detect.

An example of a backdoor attack is Cerberus Poisoning (CerP), a new distributed backdoor attack against FL systems [21]. CerP works by having multiple attackers collaborate to fine-tune a backdoor trigger for each of their devices. This makes the poisoned models from the attackers appear more similar to the unpoisoned models from honest users, allowing CerP to bypass existing defenses and successfully embed a backdoor in the final FL model.

While some defenses against label-flipping attacks exist, backdoor attacks are a significant threat. The defense mechanism defending poisoning attacks in FL (DPA-FL) tackles this issue in two phases [22]. First, it compares model weights from participants to identify significant differences, potentially indicating a malicious actor. Second, it tests the aggregated model's accuracy on a dataset, potentially revealing attackers through low performance.

3.3. Communication

Initially, a common depiction featured a central server orchestrating model aggregation, while clients performed local training. This configuration, known as centralized FL, typically employs a star topology. In contrast, decentralized FL, adopting a mesh topology, has gained prominence. In decentralized FL, no central server exists. Instead, clients use peer-to-peer (P2P) communication, exchanging local models directly. This decentralized approach enhances privacy and mitigates reliance on potentially untrusted central servers.

3.3.1. Centralized FL (CFL)

CFL takes a coordinated approach to training a model while keeping data private. Unlike traditional centralized learning—where all data go to one place—CFL leverages a central server to manage the process without ever directly accessing the raw data residing on participants' devices or institutions. This server acts as a conductor, first distributing a starting global model to all participants.

Participants train this model locally on their own datasets, tailoring it to their specific data. Afterward, only the updated model weights, representing the learning from the local training, are uploaded back to the central server. This server then plays a crucial role by aggregating these updates from multiple participants. Combining the knowledge embedded in each update, the central server refines the global model, effectively incorporating the insights from all the distributed datasets. This iterative process of distributing, training locally, and aggregating updates continues until the desired level of model performance is achieved.

3.3.2. Decentralized FL (DFL)

DFL presents an alternative approach that tackles limitations inherent to the central server in CFL. Unlike CFL, DFL dismantles the single point of control, fostering a collaborative learning environment that is both more distributed and potentially more privacy-preserving. This paradigm thrives on direct communication between participating devices or institutions, eliminating the need for a central server altogether. This P2P approach offers potential benefits in reducing communication overhead compared to CFL, as updates can be exchanged directly between participants.

However, removing the central server also complicates the training process. DFL relies on techniques like consensus algorithms [23] to ensure all participants agree on the current state of the global model, a task that becomes more intricate without a central authority. Additionally, ensuring robust security measures remains an active area of research in DFL [24]. DFL offers advantages in privacy and potentially reduces the communication burden compared with the CFL architecture.

3.4. Coalitions

The traditional FL framework treats all participants as equals, raising challenges in efficiency and communication overhead. This section explores the concept of coalitions in FL, a method for grouping agents based on specific criteria. These groupings, known as coalitions, can be formed based on the semantic similarity of the data participants manage or can be formed based on the geographic location and communication radius of participants. Here, we explore these two key approaches to coalition formation:

3.4.1. Semantic-Based Formation

In semantic-based formation, agents are grouped based on the similarity of their data. This ensures that participants within a coalition contribute data that share similar meanings and underlying patterns. This approach can be further classified into:

Static Formation

Here, coalitions are formed based on pre-defined semantic criteria. This could involve analyzing the metadata associated with the data held by each agent and initially classifying the agents into clusters. With static coalitions, once agents are grouped together, these coalitions remain fixed throughout the training process.

Dynamic Formation

Coalitions are formed or reformed continuously based on the semantic similarity of the data itself. ML techniques like automatic semantic clustering, topic modeling, or content analysis can be employed to dynamically assess data similarity and adjust coalition membership accordingly.

3.4.2. Positional-Based Formation

Positional-based formation relies on the geographical proximity of agents and their communication range. This approach is particularly relevant for scenarios where the agents are in different locations and when agents are moving.

Static Formation

Agents within a specific geographical region with a fixed communication range or that are neighbors in the communication graph are grouped into a coalition. In static coalitions, after the initial formation of groups, the group memberships do not change over time.

Dynamic Formation

In dynamic formation, agents can form or leave coalitions based on real-time location updates or changes to their communication range. This could be beneficial in scenarios where data collection is ongoing and the spatial distribution of the agents is constantly changing. Wireless ad hoc networks (WANETs) are examples of this scenario, where agents join or leave groups based on their availability within the wireless range [25].

3.5. Data Distribution

FL deals with training a model collaboratively across multiple participants, each holding their own private data. However, the data distribution across participants can be imbalanced, leading to challenges.

One of FL's major challenges lies in handling statistical heterogeneity within the data. In this context, statistical heterogeneity refers to the non-IID nature of FL data, which deviates from the assumption of identical data distributions across clients. Unlike traditional centralized machine learning, where data are typically drawn from a single source, FL data originates from diverse clients, each with its own unique data distribution. These variations can impact the quality of local models and subsequently affect the performance of the aggregated global model.

3.5.1. Label Distribution Skew

Label distribution skew refers to the unequal distribution of class labels within the training data held by different clients. Some clients may possess a surplus of data belonging to specific classes, while others may have a scarcity for the same classes. This imbalance can significantly impact the performance of the model. Imagine that participant A primarily has data for the class "cat" and very little for "dog", while participant B has the opposite distribution.

When the global model aggregates updates from clients with skewed label distributions, it can become biased toward the over-represented classes. This phenomenon occurs because local models trained on data-rich in certain classes heavily influence global updates. Consequently, the federated model prioritizes learning these dominant classes and neglects the underrepresented ones, leading to decreased accuracy for minority classes and potentially even failing to recognize them altogether.

To address this challenge, exists a novel FL method called FedMGD [26]. FedMGD aims to mitigate the performance degradation caused by label distribution skew. The key innovation lies in introducing a global generative adversarial network (GAN). This GAN operates without access to the raw local datasets, preserving data privacy. However, it can still model the global data distribution by learning from the aggregated model updates received from participants. This allows the global model to be trained using information about the overall data distribution without compromising privacy.

3.5.2. Feature Distribution Skew

While label distribution skew focuses on class imbalance, this phenomenon arises when the distribution of feature values for the same class differs significantly across client datasets. Imagine client A possesses data primarily representing cats with long, white fur, while client B's cat data depicts mostly short-haired black cats. Even if the overall number of cat images (labels) is balanced, the underlying feature distributions (fur length, color) diverge. This disparity affects the model during the training phase.

The model struggles to learn a unified representation of the "cat" class due to the conflicting feature portrayals across clients. This can lead to increased training difficulty and ultimately result in a model with poorer generalization capabilities. The model might perform well on data that resembles the specific feature distributions it encountered during training, but it could struggle with unseen data that deviates from those distributions.

3.5.3. Quantity Skew

Quantity skew refers to the unequal distribution of data samples across participating clients. In this scenario, some clients possess significantly more data points compared to others.

Clients with abundant data exert a greater influence on the global model updates due to the sheer volume of local updates they contribute. This can lead to the model becoming biased toward the data distribution of clients holding more samples. Even if the label and feature distributions are balanced globally, the model might prioritize learning patterns specific to the dominant data source, potentially neglecting valuable information present in smaller datasets from other clients.

This results in a model that performs well on data resembling the dominant client's distribution but exhibits decreased performance on data from clients with less representation.

As presented in this section, a key obstacle in FL is training an effective model when devices possess heterogeneous data, which cannot be directly exchanged. This includes imbalances in label distribution (label skew), feature distribution (feature skew), and data quantity (quantity skew) across devices. To address this issue, a method with a hierarchical FL approach utilizing a hypernetwork (HN) aims to mitigate the negative influence of non-IID data [27]. This method is presented in a landscape of Digital Twin in Industrial IoT. The lower layer of this method leverages hypernetworks to generate local model parameters for each device. The upper layer then refines these hypernetworks by aggregating the

model parameters from all devices. This approach decouples the number of parameters transmitted between the upper and lower layers, leading to improved communication efficiency, reduced computation costs, and ultimately, better model accuracy.

3.6. Model Aggregation

As highlighted, FL thrives in scenarios with heterogeneous data distributions across devices. While this protects data privacy, it also presents the challenge of effectively combining these diverse local models into a single, robust global model. This is where model aggregation techniques come into play. These techniques aim to intelligently combine the knowledge learned from individual devices, mitigating the negative effects of non-IID data and leading to a well-performing global model.

3.6.1. Synchronous Aggregation

Synchronous aggregation offers advantages in terms of convergence guarantees and ease of implementation. However, it can be susceptible to stragglers (devices that take significantly longer to train the model locally), delaying the entire update process and potentially hindering training efficiency. Additionally, communication overhead can be high due to the waiting periods before updates are uploaded.

3.6.2. Asynchronous Aggregation

Asynchronous aggregation techniques offer an alternative approach to synchronous aggregation, aiming to address limitations in scalability and efficiency. Unlike the coordinated update scheme of synchronous aggregation, asynchronous aggregation allows devices or institutions participating in FL training to upload their local model updates to the central server as soon as they become available, without waiting for others to finish. This eliminates delays caused by stragglers.

It avoids the communication bottlenecks associated with waiting periods in synchronous methods but introduces complexities in ensuring convergence of the global model, as participants contribute updates at varying times based on their local training speeds. FedTAR is an example of an FL model that uses asynchronous aggregation to minimize the sum energy consumption of all edge computing nodes of a wireless computing power network (WCPN) [28]. There is also the AMA-FES (adaptive-mixing aggregation, feature-extractor sharing) framework, which aims to mitigate the impact of the non-IID data and reduce computation load in a practical scenario where mobile UAVs act as FL training clients to conduct image classification tasks [29].

3.6.3. Hierarchical Aggregation

Hierarchical aggregation emerges as an optimization technique that addresses potential communication bottlenecks in scenarios with large numbers of participants or geographically distributed devices [30]. It also addresses privacy concerns by introducing a layered approach to update aggregation between user devices and the central server.

Hierarchical aggregation mitigates the privacy risk by having devices send their updates to intermediate servers first. These intermediate servers can then aggregate local updates before forwarding them to the central server, reducing the amount of individual data exposed. This approach is particularly valuable for the Industrial Internet of Things (IIoT) where sensitive data from various devices are involved [31].

Participants are organized into groups, forming a hierarchical structure. Local updates within a group are first aggregated, resulting in intermediate updates. These intermediate updates are then sent upwards in the hierarchy for further aggregation until they reach the central server.

Compared to directly sending individual updates to the central server, hierarchical aggregation significantly reduces communication costs. Only a condensed version of the updates travels through the network, alleviating bandwidth limitations and potentially accelerating the training process.

The specific structure of the hierarchy (number of layers, group sizes) can significantly impact efficiency. Additionally, techniques like selective aggregation, where only significant updates propagate through the hierarchy, can further optimize communication costs.

While hierarchical aggregation reduces communication overhead, it introduces an additional layer of information compression during the intermediate aggregation steps. This compression might lead to a certain loss of accuracy in the final global model.

A novel hierarchical FL framework is proposed for cloud–edge–robot collaborative training of deep learning models [32]. This framework allows robots to train the model for quality defect inspection of civil infrastructures without sharing sensitive data among themselves. The system is designed for resource-constrained robots, employing a lightweight model for efficient training and communication.

3.6.4. Robust Aggregation

As presented in Sections 3.2 and 3.5, FL models are susceptible to outliers within participant datasets and even malicious actors injecting poisoned data to manipulate the training process. Robust aggregation methods aim to detect and mitigate the influence of such anomalies on global model updates.

Various approaches can be employed for robust aggregation. These include clipping techniques that limit the magnitude of updates, outlier detection algorithms to identify and down-weight suspicious contributions, and median filtering to prioritize central tendencies within the updates [33,34].

A novel framework is secure and robust FL (SRFL), which is introduced to address security vulnerabilities in existing methods [35]. SRFL tackles the issue of model parameter leakage during aggregation using trusted execution environments (TEEs). This approach safeguards sensitive model components on resource-constrained IoT devices, even in situations with non-IID data. Evaluations demonstrate SRFL's effectiveness in improving accuracy and reducing backdoor attack success rates compared to traditional FL methods.

4. Main Practical Categories

Having explored the main theoretical trends across FL categories, we now delve into the applications driving this field forward. This section focuses on areas where FL is solving real-world problems. We will examine the distribution of research within these categories, including neural networks, classification, blockchain, Internet of Things, and edge computing. Through this analysis, we aim to identify the most promising and actively researched practical applications of FL technology.

4.1. Data Analysis

FL research shows a clear interest in leveraging powerful ML models for practical applications. The category of neural networks dominates the field, as Figures 5 and 6 depicted, with publications experiencing a staggering growth from 2019 to 2023. This highlights the focus on utilizing complex models to achieve superior performance in FL tasks. There is also a significant rise in classification, indicating a strong interest in using FL for tasks like image categorization. The emergence of blockchain and IoT (2019 onward) as prominent categories reflects the growing importance of integrating FL with secure and distributed data architectures. Similarly, edge computing has gained traction as researchers explore enabling FL on resource-constrained devices at the network edge.

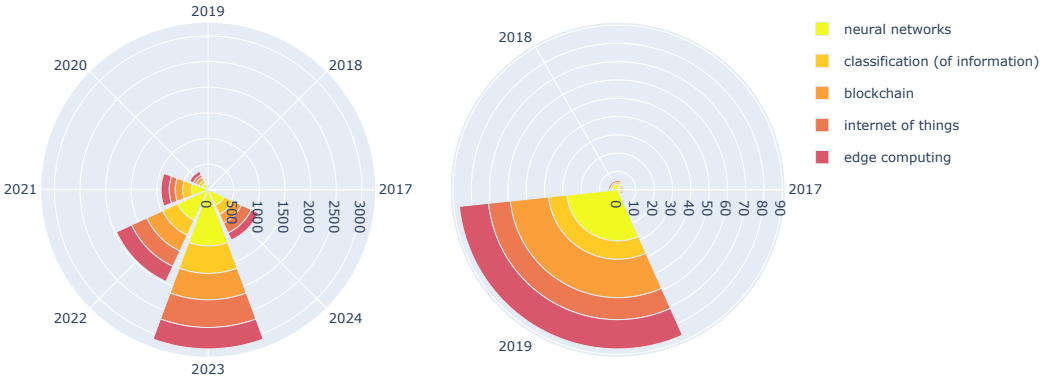

Figure 5. All practical keyword category groups.

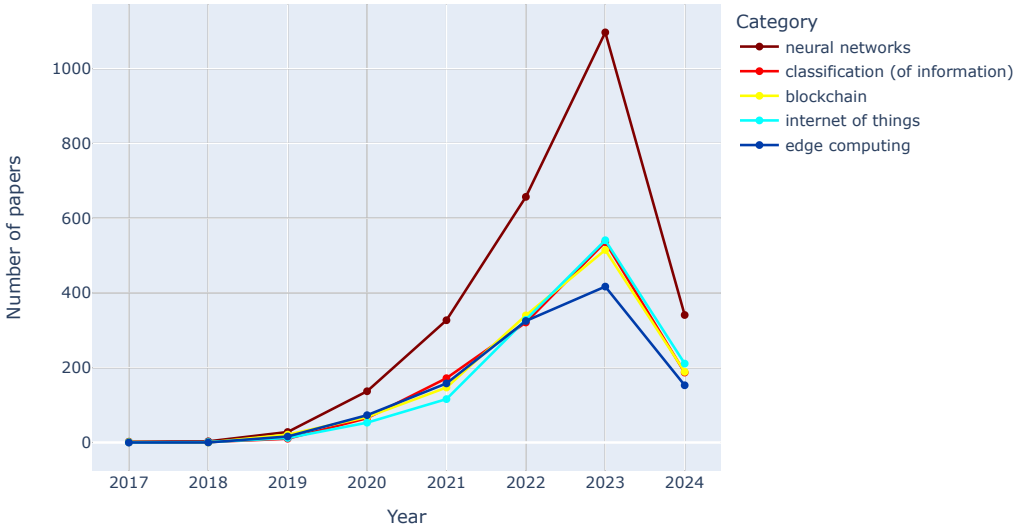

Figure 6. Tendencies of the practical keyword groups over the years.

4.2. Neural Networks

As FL continues its ascent as a privacy-preserving approach to training ML models, the role of neural networks (NNs) within this framework has become a focal point of research. With a growing number of papers dedicated to this topic, it is important to mention the relevant advancements in this field.

This section delineates NN architectures used under the FL framework, where data remain distributed across decentralized nodes while facilitating collaborative model training. Deep neural network (DNN) models tailored for FL encompass convolutional neural networks (CNNs), adept at feature extraction crucial for image processing tasks, and recurrent neural networks (RNNs), specialized in decoding sequential data and temporal dependencies. Furthermore, generative adversarial networks (GANs) demonstrate promise in generating realistic magnetic resonance imaging (MRI) images from undersampled data, while Transformers, initially developed for natural language processing (NLP) tasks, are repurposed to address image capture, information matching, and reconstruction challenges within the FL framework [36].

4.2.1. Traditional DNNs in FL

This section explores the application of CNNs and RNNs for collaborative training while preserving data privacy. We will explore specific use cases in domains like healthcare and cybersecurity, showcasing how FL empowers distributed learning on sensitive data.

CNN (Convolutional Neural Network)

CNNs excel in image processing tasks. Their core strength lies in capturing low-level to high-level features through convolutional operations. This makes them ideal for FL scenarios involving image data, such as medical imaging analysis [37] or object recognition in sensor networks [38]. CNNs can be trained on image data distributed across various devices without compromising privacy. For instance, FL with CNNs can be used to train models for disease detection in medical images without requiring hospitals to share the raw patient data [39].

Stacked CNNs (SCNNs) also excel in the cybersecurity field. A novel intrusion detection system (IDS) for wireless sensor networks (WSNs) based on the FL SCNN-Bi-LSTM model exists, which addresses limitations of traditional methods allowing sensor nodes to collaboratively train a central model without revealing their private data [40]. The SCNN-Bi-LSTM architecture analyzes both local and temporal network patterns to effectively identify even sophisticated and unknown cyber threats.

RNN (Recurrent Neural Network)

RNNs are adept at handling sequential data and capturing temporal dependencies. While not the primary choice for typical image processing tasks, RNNs can be valuable in FL settings where dynamic adjustments are needed. RNNs in healthcare are used for breast cancer detection, which allows hospitals to train an RNN on their mammogram data without sharing the raw images. This study proposes a hybrid approach combining FL with meta-heuristic optimization [41]. Another paper focuses on FL for pancreas segmentation, where data heterogeneity across institutions can hinder performance. To address this, their authors introduce FedRNN, a method that uses an RNN to adjust the aggregation weights based on the past performance of each participating site [42].

4.2.2. Emerging Applications of NN in FL

This section explores how emerging applications of NN in FL offer a revolutionary approach to training ML models while keeping data distributed across devices or servers. We will explore how GANs and Transformers are being leveraged to unlock new potential in FL applications.

GAN (Generative Adversarial Network)

A new approach called "federated synthesis" is emerging within FL. This technique aims to create synthetic data with the same properties as real data but without any privacy risks [43,44]. Researchers are exploring this method using GANs, to combine data from multiple sources while keeping it private. GANs consist of two competing NNs: a generator that creates new data, and a discriminator that tries to distinguish real data from generated data. This adversarial training allows GANs to generate highly realistic synthetic data.

Traditional GAN training requires sending large amounts of data to a central server. CAP-GAN is a novel framework that allows for collaborative training between cloud servers, edge servers, and even individual devices [44]. To address challenges caused by non-IID data, CAP-GAN incorporates a mix generator module that separates general and personalized features, improving performance on highly personalized datasets.

Transformers

Originally developed for NLP tasks, Transformers are powerful architectures based on the attention mechanism. This mechanism allows the model to focus on relevant parts of the input data, making it well-suited for tasks requiring long-range dependencies.

A recent research tackles challenges in medical image analysis with a Transformer-based FL framework. The method uses self-supervised pre-training with Transformers directly on individual institutions' data [45]. This approach overcomes limitations of data sharing and limited labeled data. The study shows significant improvements in accuracy on

medical image classification tasks compared to traditional methods, even with variations in data across institutions.

4.3. Classification (of Information)

The field of FL is actively exploring its potential for various classification tasks, including image classification, object detection, and emotion recognition. This is particularly appealing due to the vast amount of labeled data often residing on private devices, which FL can leverage while preserving privacy.

A recent study [46] investigated a privacy-preserving approach to diagnosing skin lesions using FL. While the FL model achieved comparable performance to a traditional centralized model on data from a new hospital, it fell short when tested on data from a different source. Overall, the findings suggest that FL shows promise for melanoma classification while protecting patient privacy.

Another research proposes a new FL framework called FedCAE for fault diagnosis in industrial applications [47]. Traditional approaches require sharing large amounts of data, which can be impractical due to privacy concerns. FedCAE tackles this by using convolutional autoencoders (CAEs) on local devices to extract features from the data. These features are then uploaded to a central server for training a global fault diagnosis classifier, without revealing the raw data itself. The trained classifier is then downloaded to all devices for performing local diagnoses.

4.4. Blockchain

Blockchains enable secure, verifiable interactions between devices without a central authority [48]. The field of FL with blockchain integration, also known as blockchain-based FL (BCFL), is a rapidly evolving area [49]. Researchers are looking to leverage the strengths of both technologies to address limitations in traditional FL.

Recent research proposes a new FL method for blockchain named loosely coupled local differential privacy blockchain federated learning (LL-BCFL) that addresses data privacy and efficiency concerns on federated sharing methods for massive data in blockchain [50]. Traditional blockchain storage can be slow and unsuitable for private data. LL-BCFL tackles this by combining FL on user devices with blockchain storage. The system uses a client selection mechanism to ensure data integrity and participant honesty. Additionally, a local differential privacy mechanism protects against inference attacks during training.

To protect the FL process against poisoning attacks, two models have been developed under BCFL, namely, centralized aggregated BCFL (CA-BCFL) and fully decentralized BCFL (FD-BCFL) [24]. Both leverage secure off-chain computations to mitigate attacks without compromising performance. The study demonstrates that BCFL effectively defends against poisoning attacks while keeping operational costs low.

4.5. Internet of Things

FL has emerged as a powerful approach for the Internet of Things (IoT) domain. It tackles the challenge of training ML models on data generated by vast numbers of resource-constrained devices while preserving user privacy. The FL literature reflects this synergy, highlighting several key areas of advancement.

A major focus is on addressing the limitations of resource-constrained IoT devices. Traditional FL algorithms may not be suitable for devices with limited battery power, storage, and processing capabilities. Researchers are developing techniques like model compression (Section 5.3) and efficient communication (Section 3.3) mechanisms to reduce the computational burden on these devices. This ensures participation from a wider range of IoT devices in the FL algorithm process.

Another area of exploration is heterogeneity. IoT devices often generate data with varying formats and qualities [51]. This heterogeneity can negatively impact the performance of the model. Researchers are proposing data distribution (Section 3.5) techniques to improve the performance and model aggregation methods (Section 3.6) that can handle

such inconsistencies. These techniques aim to improve the accuracy and robustness of the collaboratively learned model.

4.6. Edge Computing

While both IoT and edge computing are related to FL, they represent distinct concepts. IoT devices generate the data, while edge computing represents the layer of processing power located at the network's periphery, closer to the data source, and performs local computations [28].

One of the primary research areas is optimizing model performance and resource utilization in resource-constrained edge environments. In Section 5.3 techniques such as quantization and knowledge transfer are exposed, which are tailored to minimize the computational and memory requirements of FL models, making them suitable for deployment on low-power edge devices with limited processing capabilities. Furthermore, edge computing platforms with accelerators like GPUs and TPUs accelerate model inference and training, enhancing the efficiency and scalability of the systems.

FL is well-suited for edge devices, where data processing occurs locally [30]. It enables collaborative model training across devices at the network edge. CAP-GAN is a novel framework using GANs (presented in Section 4.2) in network edge [52]. This research tackles training GANs on devices at the network edge due to privacy and bandwidth limitations. However, traditional GAN training methods struggle with data that is not uniformly distributed across devices. To address this, CAP-GAN allows for parallel training of data and models across devices, cloud servers, and the network edge, overcoming isolated training issues. CAP-GAN introduced a mix generator module to handle highly personalized datasets that are common at the edge. Experiments show that this framework outperforms existing methods in handling non-uniformly distributed data.

5. Emerging Sub-Areas

Having explored the core theoretical categories and the practical application areas of FL research, we now turn our attention to emerging sub-areas. These sub-areas represent new lines of inquiry that have gained significant traction in recent years. Unlike the previously established categories, these sub-areas are distinguished by their later emergence and they are rapidly growing interest within the FL research community. This section delves into five such sub-areas: biological system modeling, model compression, speech recognition, real-time systems, and game theory.

5.1. Data Analysis

While all sub-areas show a clear rise in publications since 2019 and 2020, as Figure 7 depict, some demonstrate a more explosive growth trajectory. Figure 8 shows that biological system modeling exhibits the most dramatic increase, with publications nearly tripling from 2022 to 2023. This suggests a rapidly growing focus on applying FL to model complex biological systems like brain–computer interfaces (BCIs). Model compression also shows a steady and significant rise, highlighting the importance of reducing model size for deployment on resource-constrained devices in FL applications, like IoT or edge devices. Speech recognition and real-time systems show a more moderate but consistent growth, indicating a growing interest in integrating FL with these domains. Game theory, while experiencing a steady rise, has a slightly lower overall number of publications, suggesting it is a relatively new but promising sub-area exploring strategic interactions within FL systems.

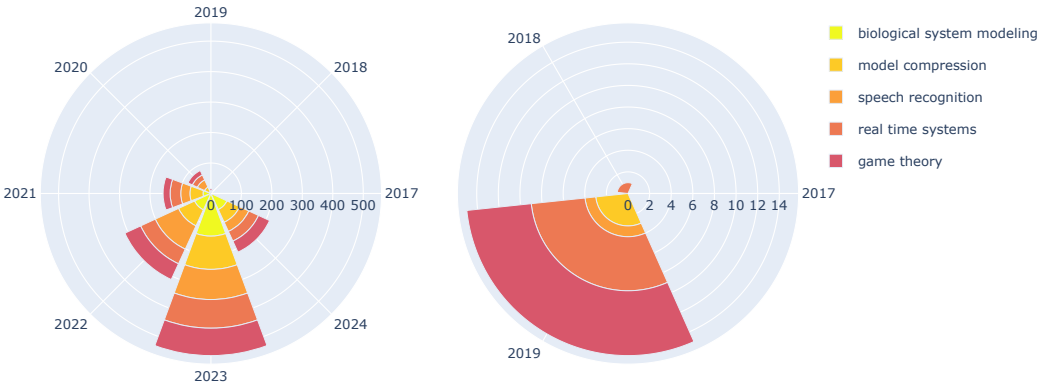

Figure 7. All emerging category groups over the years

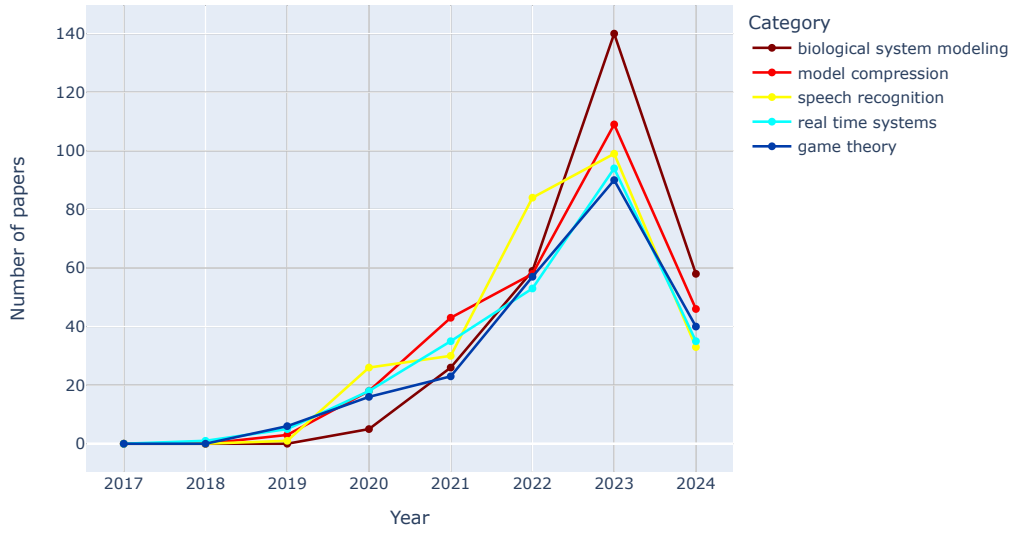

Figure 8. Tendencies of the emerging keyword groups over the years

5.2. Biological System Modeling

Brain–computer interfaces (BCIs) create a bridge between the brain and external devices by translating brain activity into commands [53]. These systems translate brain activity, captured through electroencephalogram (EEG) signals, into commands for external devices. However, a major hurdle in BCI development is the scarcity of data needed to train high-performance models. This is where FL steps in.

FL offers a privacy-preserving approach to training models on distributed datasets residing on individual devices. This eliminates the need for centralized data storage, addressing security and privacy concerns that plague biological datasets.

One recent paper proposes a novel framework called hierarchical personalized FL for EEG decoding (FLEEG) [54]. FLEEG tackles the challenge of device heterogeneity, where BCIs collect data from various sources with potentially different formats. This framework facilitates collaboration in model training across these diverse datasets, enabling knowledge sharing and boosting BCI performance. The studies presented by its researchers have shown that FLEEG can significantly improve classification accuracy, particularly for smaller datasets.

Another paper investigates the application of FL in classifying motor imagery (MI) EEG signals [37]. This approach utilizes a CNN on the PhysioNet dataset and compares two aggregation methods (FedAvg and FedProx) within the FL framework to traditional centralized ML approaches. The results demonstrate that FL can achieve classification accuracy comparable to centralized methods, while significantly reducing the risk of data

leakage. This suggests that FL holds significant promise for MI-EEG signal classification in BCI systems.

5.3. Model Compression

FL allows collaborative machine learning without compromising data privacy. However, training these models across distributed devices or servers presents a challenge: model size. Large models can lead to slow communication and hinder the scalability of FL systems. To address this, researchers are actively exploring various model compression techniques.

5.3.1. Quantization

The distributed nature of FL can lead to communication bottlenecks due to the large size of model parameters. Here is where quantization emerges as a powerful technique to address this challenge. Quantization reduces the number of bits required to represent model parameters, significantly shrinking the model size. This translates to faster communication during the FL training process, making it more efficient and scalable. However, accuracy degradation can occur during the quantization process and researchers are actively developing methods to minimize this accuracy loss.

A recent study addresses communication efficiency in hierarchical FL, where model training is distributed across devices, edge servers, and a cloud server [55]. While existing approaches leverage hierarchical aggregation and model quantization to reduce communication costs, this study proposes an accurate convergence bound that considers model quantization. This bound informs practical strategies for client-edge and edge-cloud communication, such as dynamically adjusting aggregation intervals based on network delays. The effectiveness of these strategies is validated through simulations.

Another prominent and recent area of study is the 1-bit quantization. A study proposes a new scheme that uses 1-bit compressive sensing to significantly reduce the amount of data transmitted during model updates [56]. To optimize this method, they analyze the trade-off between communication efficiency and accuracy caused by data compression. The researchers then formulate a solution to minimize these errors through scheduling devices and adjusting transmission power. While an optimal solution exists, it is computationally expensive for large networks. To address this, they develop a more scalable method suitable for real-world applications with many devices. Simulations show this approach achieves comparable performance to traditional FL with significantly less communication, making it a promising technique for large-scale FL.

5.3.2. Knowledge Distillation

Another key approach is knowledge distillation. This technique involves training a smaller, student model to mimic the behavior of a larger, pre-trained teacher model. The student model learns from the teacher's predictions, resulting in a compressed model with comparable accuracy. Knowledge distillation is particularly useful in FL as it allows for transferring knowledge from a powerful trained model to smaller models deployed on user devices.

Recent research develops an intrusion detection method based on a semi-supervised FL scheme via knowledge distillation [57]. The study proposes an intrusion detection method in IoT devices. Existing FL methods for intrusion detection raise privacy concerns and struggle with non-private data distributions. To address this, the authors developed a method that leverages unlabeled data to improve detection accuracy while protecting privacy. Their approach uses a special NN model to both classify traffic data and assess the quality of labels generated by individual devices. This, combined with a hard-label strategy and voting mechanism, reduces communication overhead.

5.3.3. Pruning

Another promising direction is pruning. Pruning techniques identify and remove redundant or unimportant weights within a model. This process reduces the model's

overall size without significantly impacting its performance. Advanced pruning algorithms can identify weights with minimal influence on the final output, allowing for compression while maintaining accuracy.

PruneFL is a new framework for FL that improves training efficiency on resource-constrained devices [58]. FL trains models on distributed data while protecting privacy, but edge devices often lack processing power and bandwidth. PruneFL tackles this by dynamically reducing model size during training through a distributed pruning approach. This reduces communication and computation requirements while maintaining accuracy. The method involves an initial pruning step and further pruning throughout the FL process, optimizing the model size for efficiency. Experiments on real-world datasets running on devices like the Raspberry Pi demonstrate that PruneFL significantly reduces training time compared to traditional FL and achieves comparable accuracy to the original model with a smaller size.

5.3.4. Sparsification

Researchers are also exploring sparsification techniques. Here, the focus is on converting model weights from dense matrices to sparse ones, containing mostly zeros. Sparse models require less memory and communication bandwidth, making them ideal for FL applications. Recent advancements involve combining sparsification with other compression methods like pruning to achieve even more compact models.

GossipFL is a novel framework that utilizes sparsification and gossiping to optimize bandwidth usage while ensuring training convergence. The authors designed a novel sparsification algorithm that enables each client to communicate with only one peer using a highly sparsified model [59]. Theoretical analysis and experiments using GossipFL demonstrate that this framework significantly reduces communication traffic and time compared to existing solutions while maintaining similar model accuracy.

5.4. Speech Recognition (SR)

SR is a technology that allows computers to translate spoken words into written text. This is achieved by analyzing speech's sound waves and identifying patterns corresponding to specific words or phonemes, which are the basic units of sound in a language.

Traditional SR models require vast datasets centralized in one location for training. This raises privacy concerns, especially for applications like forensic analysis, where data sensitivity is paramount.

The fight against online child exploitation is an example, where European law enforcement agencies (LEAs) require advanced tools to analyze the growing volume of audio data. Recent research explores FL as a solution for training SR models in this domain [60]. While the study compares the effectiveness of WAV2VEC2.0 and WHISPER models, the main focus lies in leveraging FL to overcome data privacy concerns.

The results show that FL models achieve word error rates (WERs) comparable to those trained in a traditional, centralized manner. This is particularly significant considering the challenges of non-IID data distribution, where the data used have unique characteristics due to languages, accents, or recording environments.

5.5. Real-Time Systems

The traditional approach of FL involves a central server aggregating updates from participants periodically. This raises limitations for applications demanding real-time performance. The research in FL delves into techniques for enabling real-time FL that ensure low latency.

Traditional periodic updates can introduce delays that hinder real-time responsiveness. Synchronous FL can lead to slow learning due to stragglers, which are devices that take longer to process information. A novel approach that breaks away from the limitations of synchronous FL uses scalable asynchronous FL for real-time surveillance systems [61]. Asynchronous FL allows devices to participate in the training process at their own pace,

eliminating the bottleneck created by stragglers. This makes asynchronous FL a more suitable solution for large-scale, real-time applications where fast response is critical.

As it is presented in Section 3.2, security and privacy are paramount concerns in any FL system, and real-time settings pose additional challenges. Researchers are actively developing privacy-preserving communication protocols for real-time FL. Techniques like differential privacy [11] are being explored to achieve this balance between real-time performance and data security [62].

5.6. Game Theory

Game theory is a powerful mathematical field used to analyze situations where multiple parties (agents or players) interact and make decisions that can impact each other's outcomes [63]. Imagine a game of chess, where each agent considers not only their own possible moves but also how their opponent might respond. Game theory extends this concept to any situation where competing actors make strategic decisions in a setting with defined rules.

The core concept in game theory is the game itself, which acts as a model for the interactive situation. Each agent is a rational entity with well-defined preferences and a set of possible strategies they can employ. The key element is that an agent's success depends not only on their own choices but also on the strategies chosen by other players. A game will define the players, their available strategies, and how these strategies influence the final outcome for everyone involved.

Existing solutions based on game theory often assume perfect rationality in participants, so motivating participants to contribute to FL systems is an open field of research, which is used for collaborative training. A new model based on evolutionary game theory acknowledges participants' non-perfect decision-making in the long run [64]. By analyzing various scenarios, they identify strategies for parameter servers (coordinating the training) to maintain a sustainable FL system where participants are incentivized to contribute.

As commented, a limitation in existing game theory-related FL frameworks is the assumption of voluntary participation, but also it is the lack of defense against malicious actors. To address this, researchers propose a new scheme based on privacy-preserving techniques and game theory [62]. This scheme incentivizes participation through truthful mechanisms and limits the influence of malicious clients, all while achieving privacy guarantees.

6. Conclusions and Future Work

This study leverages advanced automated semantic keyword clustering techniques to analyze trends, tendencies, and emerging areas within the growing field of FL. By employing a transformer-based model, particularly the `all-mpnet-base-v2` model, the research identifies and groups 22,841 unique keywords of 7953 research articles based on their semantic meaning, providing a comprehensive view of the current state and future directions of the FL research landscape.

We present key research questions (RQs), revealing significant trends in security and communication as dominant areas of interest. The surge in publications related to these categories highlights the importance of addressing vulnerabilities and optimizing communication efficiency in FL systems. Furthermore, the analysis identifies the rising significance of coalitions, data distribution strategies, and model aggregation techniques, which are crucial for tackling challenges related to non-IID data and improving the performance of global models.

Emerging sub-areas such as biological system modeling, model compression, speech recognition, real-time systems, and the application of game theory present promising avenues for future research. These sub-areas show the field's dynamic nature and its potential for interdisciplinary applications.

To conclude, the answers to the RQs provide a structured understanding of the FL landscape: identifying the current trends (RQ1), examining their tendencies (RQ2),

exploring practical application domains (RQ3), analyzing the tendencies within these domains (RQ4), uncovering emerging sub-areas (RQ5), investigating their tendencies (RQ6), and predicting potential future trends (RQ7).

6.1. RQ1: What Are the Current Trends in FL?

This question focuses on identifying the theoretical dominant areas of interest in FL research. FL is experiencing a period of significant research growth, as evidenced by the substantial increase in publications across all categories analyzed in Section 3.1. The data reveal several key trends that are shaping the current landscape of FL research: security, communication, coalitions, data distribution, and model aggregations.

The most prominent trend is the surge of interest in security, with 498 publications in 2023. This highlights a growing concern for addressing potential vulnerabilities in FL systems, as data privacy is paramount when training models collaboratively. Similarly, the rise in communication (382 publications in 2023) reflects the importance of optimizing communication efficiency, especially as the number of participating devices and the complexity of models increase.

6.2. RQ2: What Are the Tendencies of the Current Trends in FL?

This section delves deeper into RQ1 by analyzing the direction of the identified trends in FL. The analysis of publication trends across the theoretical FL categories reveals not only a surge in interest but also the trajectory of these trends.

The most notable observation is the explosive growth in both security and communication research since 2019. Coalitions show a consistent upward trend with a peak in 2023. This indicates a sustained interest in exploring how devices or institutions can group together to optimize FL speed convergence and accuracy of the trained models. For data distribution and model aggregation categories, the substantial rise suggests a growing interest in tackling challenges related to non-IID data and improving model aggregation techniques.

6.3. RQ3: What Are the Application Domains Where FL Techniques Are Applied?

This question explores the most relevant practical applications where FL techniques are being utilized. The data presented in Section 4.1 reveals a diverse range of application domains where FL techniques are finding utility. The most important key trends are the dominance of neural networks (NNs) and the emergence of secure and distributed architecture.

6.4. RQ4: What Are the Tendencies of the Application Domains?

As RQ2 delves deeper into RQ1, this RQ investigates the trends within the RQ3 identified domains.

Firstly, NNs stand out as the most prevalent category. This signifies a strong focus on leveraging powerful ML models to achieve superior performance in FL tasks. The significant and steady rise in publications suggests that researchers are actively exploring how to adapt and optimize complex NNs for collaborative learning in FL systems.

Beyond NNs, the data highlight a growing interest in integrating FL with secure and distributed data architectures. The rise of categories like blockchain and the Internet of Things (IoT) reflects this trend.

6.5. RQ5: What Are the Emerging Sub-Areas within FL?

Recognizing the potential for further exploration, we propose additional research questions that focus on under-researched areas of FL. This question aims to identify new or niche areas that have received less attention but hold promise for future development.

The analysis in Section 5.1 reveals several promising sub-areas that have garnered increasing attention in recent years. Among the most prominent emerging sub-areas are biological system modeling and model compression. Speech recognition and real-time

systems are other emerging sub-areas with significant potential, with a close number of publications to model compression in 2023.

6.6. RQ6: What Are the Tendencies of the Emerging Sub-Areas?

We analyze the identified sub-areas in RQ5 to understand their growth trajectory and potential impact on the broader FL landscape.

Biological system modeling is the most rapidly growing sub-area with topics like bioinformatics and brain–computer interfaces. Game theory, while not exhibiting the most dramatic number of publications, also appears as an emerging sub-area with initial exploration beginning around 2019. This sub-area investigates strategic interactions within FL systems, which could be beneficial for areas like resource allocation or ensuring fairness among participants.

6.7. RQ7: What Are the Potential Future Trends of FL?

Finally, to provide a more comprehensive picture, we introduce this additional question, which looks ahead to predict potential future directions and areas of growth in FL research.

The consistently increasing number of publications in NNs suggests a continued focus on leveraging powerful models for FL tasks. Also, we can expect sustained research efforts in core areas like security and communication efficiency, as the significant rise in publications until 2023 highlights their importance. Researchers might focus on developing more robust security mechanisms to address evolving threats and optimizing communication protocols for specific federated learning applications.

Another core area is data distribution, which is likely to see continued growth. With the increasing interest in applying FL to real-world scenarios involving non-IID data, researchers will likely explore more sophisticated techniques to handle data heterogeneity and improve model performance.

6.8. Future Work

Future work will focus on expanding the software developed to include other database sources and utilizing the software to experiment with the linkage method and the distance metric of the agglomerative clustering algorithm and explore different clustering algorithms.

The results using Euclidean distance and Ward's linkage are used in this research article to group the keywords by their semantic meaning, offering significant insights into FL research trends. In future work, experimenting with different parameter values will enable us to assess the impact of different distance metrics, such as cosine similarity and Manhattan distance, on the clustering results. Additionally, experimenting with various linkage methods, including single linkage, complete linkage, and average linkage, will allow us to compare strategies for forming thematic clusters.

Funding: This research was funded by MCIN/AEI/10.13039/501100011033 and "ERDF A way of making Europe" grant number PID2021-123673OB-C31 and funded by VAE-VADEN UPV grant number TED2021-131295B-C32 and funded by GUARDIA grant number PROMETEO CIPROM/2021/077 and funded by Ayudas del Vicerrectorado de Investigacion de la UPV grant number PAID-PD-22.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Appendix A

Table A1. The first 50 keyword groups ordered by the overall number of papers.

Rank	Category	Total	2017	2018	2019	2020	2021	2022	2023	2024
0	federated learning	7953	2	6	85	393	964	2027	3394	1082
1	learning systems	5028	2	3	42	238	571	1266	2184	722
2	privacy	4175	2	3	47	245	521	1035	1754	568
3	machine learning	3458	1	2	46	183	425	902	1446	453
4	neural networks	2592	2	3	28	137	327	657	1097	341
5	global models	1568	0	2	16	63	211	402	679	195
6	data models	1551	0	3	29	101	202	382	581	253
7	computational modeling	1460	0	1	9	74	145	336	645	250
8	classification (of information)	1292	0	1	10	65	172	321	536	187
9	blockchain	1281	1	0	21	68	147	340	515	189
10	modeling accuracy	1269	0	1	16	63	154	294	539	202
11	Internet of Things	1262	0	1	12	53	116	328	541	211
12	artificial intelligence	1262	1	2	18	80	154	318	509	180
13	decentralized	1233	1	2	19	78	145	308	523	157
14	performance	1183	0	0	8	57	155	302	495	166
15	state of the art	1177	0	1	19	80	183	296	473	125
16	learning frameworks	1163	0	2	24	73	188	302	442	132
17	edge computing	1142	0	0	16	73	158	325	417	153
18	personalizations	1115	0	2	18	59	148	307	461	120
19	communication	1110	1	1	20	84	182	300	382	140
20	poisoning attacks	1095	0	1	6	41	115	270	485	177
21	security	1076	1	1	7	43	110	255	498	161
22	job analysis	1065	0	2	18	46	126	271	416	186
23	large amounts	1055	0	1	15	55	139	276	417	152
24	computational efficiency	1014	1	1	18	68	145	257	395	129
25	distributed machine learning	1013	0	3	13	91	171	264	364	107
26	over the airs	967	0	1	7	35	81	255	448	140
27	coalition	942	1	1	8	77	104	297	355	99
28	centralized	934	0	1	7	27	108	263	393	135
29	servers	929	0	0	2	39	95	227	390	176
30	commerce	908	0	2	19	53	133	243	359	99
31	wireless networks	899	0	2	11	48	106	234	378	120
32	information management	894	0	1	12	46	99	217	397	122
33	optimizations	735	0	1	7	48	85	190	293	111
34	network architecture	704	0	1	10	51	93	182	299	68
35	numerical methods	687	0	1	5	47	88	177	267	102
36	budget control	673	0	0	10	31	80	185	267	100
37	data distribution	671	0	0	6	27	72	170	297	99
38	iterative methods	664	0	1	7	38	67	149	296	106
39	smart city	656	0	1	11	42	100	168	257	77
40	benchmarking	637	1	1	9	53	99	142	255	77
41	energy utilization	627	0	1	6	26	86	173	252	83
42	human	613	0	0	4	20	48	156	293	92
43	forecasting	611	0	0	7	26	82	159	239	98
44	cloud computing	605	0	1	10	35	68	165	244	82
45	transfer learning	596	0	1	6	21	56	153	253	106
46	distillation	596	0	1	5	25	69	135	263	98
47	health care	594	0	0	4	21	77	139	280	73
48	diseases	588	0	0	0	14	60	123	296	95
49	model aggregations	574	0	0	5	34	92	139	232	72

Table A2. The last 50 keyword groups ordered by the overall number of papers.

Rank	Category	Total	2017	2018	2019	2020	2021	2022	2023	2024
50	computer vision	572	0	1	6	32	64	140	251	78
51	task analysis	567	0	0	3	20	49	143	245	107
52	5g mobile communication systems	545	0	1	7	36	75	141	216	69
53	bandwidth	529	0	1	7	50	71	123	207	70
54	current	523	0	0	6	15	62	142	229	69
55	signal processing	510	0	0	5	36	78	129	190	72
56	antennas	499	0	0	3	29	71	130	197	69
57	quality of service	495	0	0	13	41	57	134	184	66
58	diagnosis	491	0	0	1	8	59	83	265	75
59	stochastic systems	487	0	1	5	28	67	118	201	67
60	image enhancement	487	0	0	3	20	51	132	208	73
61	decision making	486	0	0	7	24	59	134	199	63
62	resource allocation	480	0	2	6	22	63	128	190	69
63	inference attacks	476	0	0	1	28	60	101	209	77
64	convergence	468	0	0	3	20	54	101	210	80
65	vehicles	466	0	1	4	23	55	121	181	81
66	intelligent vehicle highway systems	458	0	0	2	18	54	110	203	71
67	digital storage	446	0	0	6	29	47	127	185	52
68	cryptography	438	1	0	9	31	48	92	181	76
69	matrix algebra	433	1	0	7	24	51	116	166	68
70	reinforcement learning	427	0	0	5	20	49	111	178	64
71	risk assessment	424	0	0	6	23	60	94	199	42
72	intrusion detection	416	0	0	1	13	46	108	189	59
73	iid data	403	0	0	2	12	44	114	180	51
74	large scales	397	0	0	3	18	52	112	160	52
75	medical imaging	369	0	0	1	12	34	85	172	65
76	incentive mechanism	352	0	0	7	32	47	92	118	56
77	clustering	350	1	0	2	11	44	92	153	47
78	channel state information	342	0	0	4	21	55	104	118	40
79	gradient methods	334	0	0	6	24	49	87	129	39
80	Industrial Internet of Things	301	0	0	2	19	51	71	121	37
81	biological system modeling	288	0	0	0	5	26	59	140	58
82	speech recognition	273	0	0	1	26	30	84	99	33
83	real-time systems	241	0	1	5	18	35	53	94	35
84	game theory	232	0	0	6	16	23	57	90	40
85	graph neural networks	195	0	0	1	1	18	41	97	37
86	machine design	180	0	1	1	27	34	37	55	25
87	unmanned aerial vehicles (UAV)	176	0	0	0	7	30	43	70	26
88	spatial-temporal	174	0	0	1	10	15	49	73	26
89	labeled data	174	0	0	2	9	25	41	75	22
90	traffic congestion	166	0	0	1	9	24	40	65	27
91	quantization	164	0	0	0	6	24	40	61	33
92	sensor nodes	156	0	0	3	10	21	27	70	25
93	model compression	143	0	0	3	15	23	24	60	18
94	data sample	138	0	0	2	13	20	28	59	16
95	tumors	132	0	0	0	5	11	32	67	17
96	hyperparameter	128	0	0	4	12	20	30	50	12
97	synchronization	121	0	0	1	4	20	28	55	13
98	leaf disease	118	0	0	1	0	3	16	85	13
99	web services	90	0	0	3	11	9	36	25	6

References

1. Konečný, J.; McMahan, H.B.; Yu, F.X.; Richtárik, P.; Suresh, A.T.; Bacon, D. Federated Learning: Strategies for Improving Communication Efficiency. *arXiv* **2016**, arXiv:1610.05492.
2. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, L.; Polosukhin, I. Attention is all you need. In Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17, Red Hook, NY, USA, 4–9 December 2017; pp. 6000–6010.

3. Lo, K.; Wang, L.L.; Neumann, M.; Kinney, R.; Weld, D. S2ORC: The Semantic Scholar Open Research Corpus. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Online, 5–10 July 2020; pp. 4969–4983. [CrossRef]
4. Hashimoto, T.B.; Alvarez-Melis, D.; Jaakkola, T.S. Word embeddings as metric recovery in semantic spaces. *Trans. Assoc. Comput. Linguist.* **2016**, *4*, 273–286. [CrossRef]
5. Fredrikson, M.; Jha, S.; Ristenpart, T. Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures. In Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, CCS '15, New York, NY, USA, 12–16 October 2015; pp. 1322–1333. [CrossRef]
6. Zhang, L.; Xu, J.; Vijayakumar, P.; Sharma, P.K.; Ghosh, U. Homomorphic Encryption-Based Privacy-Preserving Federated Learning in IoT-Enabled Healthcare System. *IEEE Trans. Netw. Sci. Eng.* **2023**, *10*, 2864–2880. [CrossRef]
7. Schlegel, R.; Kumar, S.; Rosnes, E.; Amat, A.G.i. CodedPaddedFL and CodedSecAgg: Straggler Mitigation and Secure Aggregation in Federated Learning. *IEEE Trans. Commun.* **2023**, *71*, 2013–2027. [CrossRef]
8. Asad, M.; Shaukat, S.; Javanmardi, E.; Nakazato, J.; Bao, N.; Tsukada, M. Secure and Efficient Blockchain-Based Federated Learning Approach for VANETs. *IEEE Internet Things J.* **2024**, *11*, 9047–9055. [CrossRef]
9. Qiao, F.; Li, Z.; Kong, Y. A Privacy-Aware and Incremental Defense Method Against GAN-Based Poisoning Attack. *IEEE Trans. Comput. Soc. Syst.* **2024**, *11*, 1708–1721. [CrossRef]
10. Zhou, J.; Wu, N.; Wang, Y.; Gu, S.; Cao, Z.; Dong, X.; Choo, K.K.R. A Differentially Private Federated Learning Model Against Poisoning Attacks in Edge Computing. *IEEE Trans. Dependable Secur. Comput.* **2023**, *20*, 1941–1958. [CrossRef]
11. Dwork, C. Differential privacy. In Proceedings of the International Colloquium on Automata, Languages, and Programming, Venice, Italy, 10–14 July 2006; Springer: Berlin/Heidelberg, Germany, 2006; pp. 1–12.
12. Jiang, W.; Li, H.; Liu, S.; Ren, Y.; He, M. A flexible poisoning attack against machine learning. In Proceedings of the ICC 2019—2019 IEEE International Conference on Communications (ICC), Shanghai, China, 20–24 May 2019; IEEE: Piscataway, NJ, USA, 2019; pp. 1–6.
13. Gupta, P.; Yadav, K.; Gupta, B.B.; Alazab, M.; Gadekallu, T.R. A Novel Data Poisoning Attack in Federated Learning based on Inverted Loss Function. *Comput. Secur.* **2023**, *130*, 103270. [CrossRef]
14. Omran, A.H.; Mohammed, S.Y.; Aljanabi, M. Detecting Data Poisoning Attacks in Federated Learning for Healthcare Applications Using Deep Learning. *Iraqi J. Comput. Sci. Math.* **2023**, *4*, 225–237. [CrossRef]
15. Li, S.; Ngai, E.; Voigt, T. Byzantine-Robust Aggregation in Federated Learning Empowered Industrial IoT. *IEEE Trans. Ind. Inform.* **2023**, *19*, 1165–1175. [CrossRef]
16. Yang, M.; Cheng, H.; Chen, F.; Liu, X.; Wang, M.; Li, X. Model poisoning attack in differential privacy-based federated learning. *Inf. Sci.* **2023**, *630*, 158–172. [CrossRef]
17. Kalapaaking, A.P.; Khalil, I.; Yi, X. Blockchain-Based Federated Learning with SMPC Model Verification against Poisoning Attack for Healthcare Systems. *IEEE Trans. Emerg. Top. Comput.* **2024**, *12*, 269–280. [CrossRef]
18. Wang, Z.; Huang, Y.; Song, M.; Wu, L.; Xue, F.; Ren, K. Poisoning-Assisted Property Inference Attack against Federated Learning. *IEEE Trans. Dependable Secur. Comput.* **2023**, *20*, 3328–3340. [CrossRef]
19. Zhao, P.; Cao, Z.; Jiang, J.; Gao, F. Practical Private Aggregation in Federated Learning against Inference Attack. *IEEE Internet Things J.* **2023**, *10*, 318–329. [CrossRef]
20. Gong, X.; Chen, Y.; Wang, Q.; Kong, W. Backdoor Attacks and Defenses in Federated Learning: State-of-the-Art, Taxonomy, and Future Directions. *IEEE Wirel. Commun.* **2023**, *30*, 114–121. [CrossRef]
21. Lyu, X.; Han, Y.; Wang, W.; Liu, J.; Wang, B.; Liu, J.; Zhang, X. Poisoning with Cerberus: Stealthy and Colluded Backdoor Attack against Federated Learning. *Proc. AAAI Conf. Artif. Intell.* **2023**, *37*, 9020–9028. [CrossRef]
22. Lai, Y.C.; Lin, J.Y.; Lin, Y.D.; Hwang, R.H.; Lin, P.C.; Wu, H.K.; Chen, C.K. Two-phase Defense against Poisoning Attacks on Federated Learning-based Intrusion Detection. *Comput. Secur.* **2023**, *129*, 103205. [CrossRef]
23. Carrascosa, C.; Rincón, J.; Rebollo, M. Co-Learning: Consensus-based Learning for Multi-Agent Systems. In Proceedings of the Advances in Practical Applications of Agents, Multi-Agent Systems, and Complex Systems Simulation. The PAAMS Collection, L'Aquila, Italy, 13–15 July 2022; Dignum, F., Mathieu, P., Corchado, J.M., De La Prieta, F., Eds.; Springer: Cham, Switzerland, 2022; pp. 63–75.
24. Thennakoon, R.; Wanigasundara, A.; Weerasinghe, S.; Seneviratne, C.; Siriwardhana, Y.; Liyanage, M. Decentralized Defense: Leveraging Blockchain against Poisoning Attacks in Federated Learning Systems. In Proceedings of the 2024 IEEE 21st Consumer Communications & Networking Conference (CCNC), Las Vegas, NV, USA, 6–9 January 2024; pp. 950–955. [CrossRef]
25. Rebollo, M.; Rincon, J.A.; Hernández, L.; Enguix, F.; Carrascosa, C. Extending the Framework for Developing Intelligent Virtual Environments (FIVE) with Artifacts for Modeling Internet of Things Devices and a New Decentralized Federated Learning Based on Consensus for Dynamic Networks. *Sensors* **2024**, *24*, 1342. [CrossRef]
26. Sheng, T.; Shen, C.; Liu, Y.; Ou, Y.; Qu, Z.; Liang, Y.; Wang, J. Modeling global distribution for federated learning with label distribution skew. *Pattern Recognit.* **2023**, *143*, 109724. [CrossRef]
27. Yang, J.; Jiang, W.; Nie, L. Hypernetworks-Based Hierarchical Federated Learning on Hybrid Non-IID Datasets for Digital Twin in Industrial IoT. *IEEE Trans. Netw. Sci. Eng.* **2024**, *11*, 1413–1423. [CrossRef]
28. Sun, W.; Li, Z.; Wang, Q.; Zhang, Y. FedTAR: Task and Resource-Aware Federated Learning for Wireless Computing Power Networks. *IEEE Internet Things J.* **2023**, *10*, 4257–4270. [CrossRef]

29. Li, J.; Liu, X.; Mahmoodi, T. Federated Learning in Heterogeneous Wireless Networks with Adaptive Mixing Aggregation and Computation Reduction. *IEEE Open J. Commun. Soc.* **2024**, *5*, 2164–2182. [CrossRef]
30. Wu, Q.; Chen, X.; Ouyang, T.; Zhou, Z.; Zhang, X.; Yang, S.; Zhang, J. HiFlash: Communication-Efficient Hierarchical Federated Learning with Adaptive Staleness Control and Heterogeneity-Aware Client-Edge Association. *IEEE Trans. Parallel Distrib. Syst.* **2023**, *34*, 1560–1579. [CrossRef]
31. Chen, J.; Xue, J.; Wang, Y.; Huang, L.; Baker, T.; Zhou, Z. Privacy-Preserving and Traceable Federated Learning for data sharing in industrial IoT applications. *Expert Syst. Appl.* **2023**, *213*, 119036. [CrossRef]
32. Wu, H.T.; Li, H.; Chi, H.L.; Kou, W.B.; Wu, Y.C.; Wang, S. A hierarchical federated learning framework for collaborative quality defect inspection in construction. *Eng. Appl. Artif. Intell.* **2024**, *133*, 108218. [CrossRef]
33. Uddin, M.P.; Xiang, Y.; Cai, B.; Lu, X.; Yearwood, J.; Gao, L. ARFL: Adaptive and Robust Federated Learning. *IEEE Trans. Mob. Comput.* **2024**, *23*, 5401–5417. [CrossRef]
34. Yang, H.; Gu, D.; He, J. A Robust and Efficient Federated Learning Algorithm against Adaptive Model Poisoning Attacks. *IEEE Internet Things J.* **2024**, *11*, 16289–16302. [CrossRef]
35. Cao, Y.; Zhang, J.; Zhao, Y.; Su, P.; Huang, H. SRFL: A Secure & Robust Federated Learning framework for IoT with trusted execution environments. *Expert Syst. Appl.* **2024**, *239*, 122410. [CrossRef]
36. Hossain, M.B.; Shinde, R.K.; Oh, S.; Kwon, K.C.; Kim, N. A Systematic Review and Identification of the Challenges of Deep Learning Techniques for Undersampled Magnetic Resonance Image Reconstruction. *Sensors* **2024**, *24*, 753. [CrossRef]
37. Ghader, M.; Farahani, B.; Rezvani, Z.; Shahsavari, M.; Fazlali, M. Exploiting Federated Learning for EEG-based Brain-Computer Interface System. In Proceedings of the 2023 IEEE International Conference on Omni-Layer Intelligent Systems (COINS), Berlin, Germany, 23–25 July 2023; pp. 1–6. [CrossRef]
38. Mehta, S.; Kukreja, V.; Gupta, A. Next-Generation Wheat Disease Monitoring: Leveraging Federated Convolutional Neural Networks for Severity Estimation. In Proceedings of the 2023 4th International Conference for Emerging Technology (INCET), Belgaum, India, 26–28 May 2023; pp. 1–6. [CrossRef]
39. Pandianchery, M.S.; Sowmya, V.; Gopalakrishnan, E.A.; Ravi, V.; Soman, K.P. Centralized CNN–GRU Model by Federated Learning for COVID-19 Prediction in India. *IEEE Trans. Comput. Soc. Syst.* **2024**, *11*, 1362–1371. [CrossRef]
40. Bukhari, S.M.S.; Zafar, M.H.; Houran, M.A.; Moosavi, S.K.R.; Mansoor, M.; Muaaz, M.; Sanfilippo, F. Secure and privacy-preserving intrusion detection in wireless sensor networks: Federated learning with SCNN-Bi-LSTM for enhanced reliability. *Ad Hoc Netw.* **2024**, *155*, 103407. [CrossRef]
41. Kumbhare, S.; Kathole, A.B.; Shinde, S. Federated learning aided breast cancer detection with intelligent Heuristic-based deep learning framework. *Biomed. Signal Process. Control* **2023**, *86*, 105080. [CrossRef]
42. Deng, Z.; Qureshi, T.A.; Javed, S.; Wang, L.; Christodoulou, A.G.; Xie, Y.; Gaddam, S.; Pandol, S.J.; Li, D. FedRNN: Federated Learning with RNN-Based Aggregation on Pancreas Segmentation. In Proceedings of the Medical Imaging and Computer-Aided Diagnosis, San Diego, CA, USA, 19–23 February 2023; Su, R., Zhang, Y., Liu, H., Frangi, A.F., Eds.; Springer: Singapore, 2023; pp. 453–464.
43. Little, C.; Elliot, M.; Allmendinger, R. Federated learning for generating synthetic data: A scoping review. *Int. J. Popul. Data Sci.* **2023**, *8*. [CrossRef] [PubMed]
44. Cai, X.; Lan, Y.; Zhang, Z.; Wen, J.; Cui, Z.; Zhang, W. A Many-Objective Optimization Based Federal Deep Generation Model for Enhancing Data Processing Capability in IoT. *IEEE Trans. Ind. Inform.* **2023**, *19*, 561–569. [CrossRef]
45. Yan, R.; Qu, L.; Wei, Q.; Huang, S.C.; Shen, L.; Rubin, D.L.; Xing, L.; Zhou, Y. Label-Efficient Self-Supervised Federated Learning for Tackling Data Heterogeneity in Medical Imaging. *IEEE Trans. Med Imaging* **2023**, *42*, 1932–1943. [CrossRef] [PubMed]
46. Haggenmüller, S.; Schmitt, M.; Krieghoff-Henning, E.; Hekler, A.; Maron, R.C.; Wies, C.; Utikal, J.S.; Meier, F.; Hobelsberger, S.; Gellrich, F.F.; et al. Federated Learning for Decentralized Artificial Intelligence in Melanoma Diagnostics. *JAMA Dermatol.* **2024**, *160*, 303–311. [CrossRef] [PubMed]
47. Yu, Y.; Guo, L.; Gao, H.; He, Y.; You, Z.; Duan, A. FedCAE: A New Federated Learning Framework for Edge-Cloud Collaboration Based Machine Fault Diagnosis. *IEEE Trans. Ind. Electron.* **2024**, *71*, 4108–4119. [CrossRef]
48. Christidis, K.; Devetsikiotis, M. Blockchains and Smart Contracts for the Internet of Things. *IEEE Access* **2016**, *4*, 2292–2303. [CrossRef]
49. Tang, Y.; Zhang, Y.; Niu, T.; Li, Z.; Zhang, Z.; Chen, H.; Zhang, L. A Survey on Blockchain-Based Federated Learning: Categorization, Application and Analysis. *Comput. Model. Eng. Sci.* **2024**, *139*, 2451–2477. [CrossRef]
50. Wu, B.; Kang, H. Research on Federated Sharing Methods for Massive Data in Blockchain. In *Proceedings of the Smart Grid and Internet of Things*; Deng, D.J., Chen, J.C., Eds.; Springer: Cham, Switzerland, 2024; pp. 12–27.
51. Sumitra; Shenoy, M.V. HFedDI: A novel privacy preserving horizontal federated learning based scheme for IoT device identification. *J. Netw. Comput. Appl.* **2023**, *214*, 103616. [CrossRef]
52. Zhang, J.; Zhao, L.; Yu, K.; Min, G.; Al-Dubai, A.Y.; Zomaya, A.Y. A Novel Federated Learning Scheme for Generative Adversarial Networks. *IEEE Trans. Mob. Comput.* **2024**, *23*, 3633–3649. [CrossRef]
53. Nicolas-Alonso, L.F.; Gomez-Gil, J. Brain computer interfaces, a review. *Sensors* **2012**, *12*, 1211–1279. [CrossRef]
54. Liu, R.; Chen, Y.; Li, A.; Ding, Y.; Yu, H.; Guan, C. Aggregating intrinsic information to enhance BCI performance through federated learning. *Neural Netw.* **2024**, *172*, 106100. [CrossRef]

55. Liu, L.; Zhang, J.; Song, S.; Letaief, K.B. Hierarchical Federated Learning with Quantization: Convergence Analysis and System Design. *IEEE Trans. Wirel. Commun.* **2023**, *22*, 2–18. [CrossRef]
56. Fan, X.; Wang, Y.; Huo, Y.; Tian, Z. 1-Bit Compressive Sensing for Efficient Federated Learning over the Air. *IEEE Trans. Wirel. Commun.* **2023**, *22*, 2139–2155. [CrossRef]
57. Zhao, R.; Wang, Y.; Xue, Z.; Ohtsuki, T.; Adebisi, B.; Gui, G. Semisupervised Federated-Learning-Based Intrusion Detection Method for Internet of Things. *IEEE Internet Things J.* **2023**, *10*, 8645–8657. [CrossRef]
58. Jiang, Y.; Wang, S.; Valls, V.; Ko, B.J.; Lee, W.H.; Leung, K.K.; Tassiulas, L. Model Pruning Enables Efficient Federated Learning on Edge Devices. *IEEE Trans. Neural Netw. Learn. Syst.* **2023**, *34*, 10374–10386. [CrossRef]
59. Tang, Z.; Shi, S.; Li, B.; Chu, X. GossipFL: A Decentralized Federated Learning Framework with Sparsified and Adaptive Communication. *IEEE Trans. Parallel Distrib. Syst.* **2023**, *34*, 909–922. [CrossRef]
60. Vásquez-Correa, J.C.; Álvarez Muniain, A. Novel Speech Recognition Systems Applied to Forensics within Child Exploitation: Wav2vec2.0 vs. Whisper. *Sensors* **2023**, *23*, 1843. [CrossRef]
61. Hagos, D.H.; Tankard, E.; Rawat, D.B. A Scalable Asynchronous Federated Learning for Privacy-Preserving Real-Time Surveillance Systems. In Proceedings of the IEEE INFOCOM 2023—IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), New York, NY, USA, 17–20 May 2023; pp. 1–6. [CrossRef]
62. Zhang, L.; Zhu, T.; Xiong, P.; Zhou, W.; Yu, P.S. A Robust Game-Theoretical Federated Learning Framework with Joint Differential Privacy. *IEEE Trans. Knowl. Data Eng.* **2023**, *35*, 3333–3346. [CrossRef]
63. Shoham, Y.; Leyton-Brown, K. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*; Cambridge University Press: Cambridge, UK, 2008.
64. Luo, X.; Zhang, Z.; He, J.; Hu, S. Strategic Analysis of the Parameter Servers and Participants in Federated Learning: An Evolutionary Game Perspective. *IEEE Trans. Comput. Soc. Syst.* **2024**, *11*, 132–143. [CrossRef]