



A survey on Clustered Federated Learning: Taxonomy, Analysis and Applications

Michael Ben Ali, Omar El-Rifai, Imen Megdiche, André Peninou, Olivier Teste

► To cite this version:

Michael Ben Ali, Omar El-Rifai, Imen Megdiche, André Peninou, Olivier Teste. A survey on Clustered Federated Learning: Taxonomy, Analysis and Applications. 2025. <hal-04916238v3>

HAL Id: hal-04916238

<https://hal.science/hal-04916238v3>

Preprint submitted on 19 Jan 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

A survey on Clustered Federated Learning: Taxonomy, Analysis and Applications*

MICHAEL BEN ALI, Université Toulouse, UT3, IRIT, CNRS, France

OMAR EL-RIFAI, Université Toulouse, UT3, IRIT, CNRS, France

IMEN MEGDICHE, INU Champollion, ISIS Castres, IRIT, CNRS, France

ANDRÉ PENINO, Université Toulouse, UT2J, IRIT, CNRS, France

OLIVIER TESTE, Université Toulouse, UT2J, IRIT, CNRS, France

As Federated Learning (FL) expands, the challenge of non-independent and identically distributed (non-IID) data becomes critical. Clustered Federated Learning (CFL) addresses this by training multiple specialized models, each representing a group of clients with similar data distributions. However, the term "CFL" has increasingly been applied to operational strategies unrelated to data heterogeneity, creating significant ambiguity. This survey provides a systematic review of the CFL literature and introduces a principled taxonomy that classifies algorithms into Server-side, Client-side, and Metadata-based approaches. Our analysis reveals a distinct dichotomy: while theoretical research prioritizes privacy-preserving Server/Client-side methods, applied papers overwhelmingly favor Metadata-based approaches, prioritizing efficiency over privacy. Furthermore, we explicitly distinguish "Core CFL" (grouping clients for non-IID data) from "Clustered X FL" (operational variants for system heterogeneity). Finally, we outline lessons learned and future directions to bridge the gap between theoretical privacy and practical efficiency.

1 Introduction

Federated Learning (FL) was designed to collaboratively train a **single global** deep learning model — that is, one shared set of parameters collaboratively optimized by all participating clients — without sharing raw data across devices. In a typical FL setup, a central server distributes a model to participating clients, who train it locally using private data. The clients return their updated parameters¹ to the server, which aggregates them into a new global model. This complete cycle of distribution, local training, and aggregation constitutes a single **communication round**. The process iterates over multiple rounds until convergence. FL is designed for large-scale, distributed environments where communication is constrained, and data is decentralized, heterogeneous, and privacy-sensitive [81].

Nevertheless, subsequent studies have revealed several shortcomings in the initial approaches [56]. FL solutions are often inadequate when faced with data heterogeneity in many practical applications [91, 137]. Consequently, particular attention is given to the challenges of training a single global solution when datasets are heterogeneous (non-IID) across clients.

Zhao et al. [148] demonstrated the difficulty of this scenario by showing that the optimization direction of federated models on non-IID data diverges significantly from that of a hypothetical centralized model trained on the same data. This weight divergence creates a larger

*This manuscript is a major revision of the version previously deposited as <https://arxiv.org/abs/2501.17512> arXiv:2501.17512. Updates include a restructuring of the taxonomy, a literature review update through June 2025, and an adjustment to the authorship order to reflect contributions to this revised version.

¹For convenience, we refer to the model and its parameters interchangeably. We assume the model architecture is fixed; thus, references to the parameters define the model.

gap between the federated and optimal centralized weights than in homogeneous settings, resulting in a significant reduction in performance.

In response to this challenge, multiple variants of FL have emerged in the scientific literature [53, 67], including model aggregation techniques, knowledge distillation, regularization-based methods, and personalized FL. One such approach is known as *Clustered Federated Learning (CFL)* [33, 103], which relaxes the constraint of a single global model by training multiple models, each tailored to a group of clients with similar data distributions. This strategy builds upon the analysis of weight divergence in Zhao et al. [148], leveraging the observation that clients with similar local data distributions tend to update their models in similar directions. The CFL approach assumes the existence of numerous, distinct distributions of client data. It aims to capture this structure by grouping clients and learning *an independent model for each cluster*.

The goal of the paper is to conduct a Survey of Clustered Federated Learning (CFL). From our literature review, however, it is clear that the term CFL has not remained confined to its original meaning, that is, to address non-IID data by training multiple models. Yet, over time, other uses of the same terminology appeared in the literature. For example, several studies employ the term CFL to describe solutions using clustering to tackle system heterogeneity [27, 109, 120, 125, 133, 151]. However, these works retain the single global model objective rather than addressing non-IID data distributions through multiple models. Furthermore, significant terminological inconsistency exists in the literature: terms such as *Clustered FL*, *Clustering FL*, *Cluster-based FL*, and *Clustering-based FL* are often used interchangeably, and sometimes overlap with unrelated notions such as Federated Clustering.

One of the objectives of this survey is therefore to disentangle the different meanings and research trends of CFL in the literature. We will explicitly distinguish what we call core CFL: algorithms that cluster clients to tackle non-IID data by training multiple models. We also separate other uses we call **Clustered X FL**, where clustering supports goals such as communication efficiency, client scheduling, or system-level optimization without addressing non-IID data. This distinction structures our survey and ensures clarity in positioning CFL relative to other branches of the FL literature.

Organization of the Content of the Paper. The survey is structured to provide a unified framework for understanding CFL and its applications. First, Section 2 presents our literature search methodology. Section 3 then discusses the positioning of CFL within the broader field of Federated Learning. Section 4 introduces Core CFL, beginning with background on FL and the non-IID challenge (Subsection 4.1). Subsection 4.2 formalizes the CFL problem and presents our core taxonomy, classifying state-of-the-art domain agnostic algorithms into **Server-side**, **Client-side**, and **Metadata-based** approaches. This is supported by a classification table in Section 4.3 summarizing existing methods. We then highlight common design challenges, such as determining the number of clusters, scalability, soft clustering, dynamic environments, security, and privacy (section 4.3.1 to 4.3.7). Section 4.4 introduces a unified multi-tier workflow that abstracts the structure of CFL algorithms. Section 4.5 discusses CFL evaluation methodologies, covering the role of the non-IID data taxonomy in creating benchmark configurations. Section 4.6 reviews the application studies of Core CFL across domains like IoT, Mobility, Energy, and Healthcare. Section 5 distinguishes Core CFL from **Clustered X FL** variants, which apply clustering for operational optimization, often maintaining a single global model. Section 6 identifies open challenges and future research directions. Finally, Section 7 concludes the survey by summarizing key findings and highlighting the distinction between CFL uses.

2 Systematic Literature Search Methodology

To ensure a rigorous review of the CFL literature, we adopted a systematic methodology for paper selection, guided by clear inclusion criteria. The scope of the review was limited to publications released before June 2025. We constructed a search query targeting CFL terminology :

```
("Clustered" OR "Cluster-based" OR "Clustering-based" OR "Clustering" OR
Clustered FL" OR "Cluster-based FL" OR "Clustering-based FL" OR "Clustering FL")
AND "Federated Learning"
```

This query was applied across four major databases: Web of Science, Semantic Scholar, Google Scholar, and Dimensions.ai. These platforms were chosen both for their free accessibility and coverage of computer science and interdisciplinary research on FL.

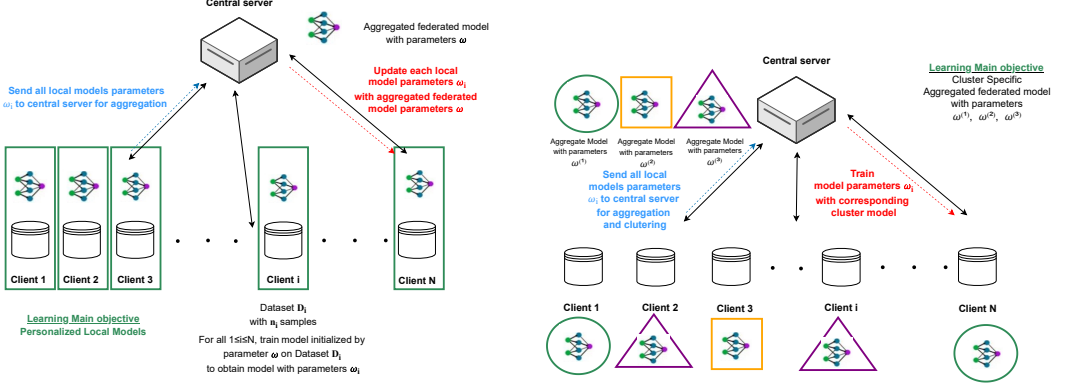
Only peer-reviewed articles published in journals or conferences indexed in Scopus or Web of Science were retained. Workshop papers and non-peer-reviewed sources, such as preprints, were discarded. To ensure quality, ranking filters were applied according to the year of publication: journal articles were required to appear in Q1-ranked journals based on the Scimago Journal Rank (SJR), while conference papers were selected only if classified as A or A* in the CORE Conference Ranking. For core CFL papers introducing new domain-agnostic algorithms (see Section 4.3), all selected works provide experimental evaluations against prior CFL approaches to ensure comparability and relevance. Exceptions were made for foundational papers [10, 32, 33, 78, 103] which established the CFL paradigm. In total, 529 articles were initially retrieved from queries. After removing duplicates, extended versions, false positives, and papers out of scope, 130 CFL articles were retained. Among these, 36 correspond to *domain agnostic core CFL* (Q1/A/A*) publications, while 75 are *application-oriented* and 19 classified as *Clustered X FL*.

3 Positioning of CFL within the Federated Learning Landscape

Non-IID data remains a major challenge in FL [53, 56, 67]. Several strategies have been proposed to mitigate heterogeneous client distributions. These include techniques that modify the optimization dynamics, impose regularization or model distillation (All highlighted in Ji et al. [53]). These approaches share the single global model assumption. When data distributions are only mildly heterogeneous, such strategies often suffice to stabilize training and preserve convergence. However, when client distributions are deeply distinct, the single-model assumption becomes restrictive. Updating one global model to fit multiple divergent updates leads to unstable optimization and poor performance.

CFL addresses this limitation by assuming the existence of multiple client groups, each following its own underlying distribution. The goal is not to correct clients so that they all fit a shared model but instead to identify coherent subsets of clients and train one model per group. CFL therefore replaces the single-model constraint with multiple models adapted to different distributions. While few categorize CFL as a subfield of Personalized Federated Learning (PFL) (a family of techniques that relaxes the single global model constraint by tailoring models to individual clients [99, 113]), most research treats CFL as a distinct approach [53, 67]. The confusion often arises from a personalization standpoint. In PFL, the personalization goal is typically associated with optimizing N models locally for N clients (Figure 1a) using one single global model. In CFL, by contrast, the objective is to optimize N models using only K unique models ($K \ll N$), since clients within the same cluster share a common model (Figure 1b). We suggest that CFL is better viewed as distinct from PFL, as the two approaches rely on different assumptions about the underlying data distribution across clients. While PFL typically assumes a single global distribution and adapts the model to individual clients to accommodate local empirical distribution differences, CFL assumes the existence of multiple distinct client groups, each characterized by its own distribution. This perspective aligns with how

some popular FL surveys [53, 67] present the field, where CFL and PFL are introduced as separate approaches. Moreover, some CFL algorithms [12, 77] integrate personalization techniques to tackle intra-cluster heterogeneity, optimizing both K global models and N local models. This further highlights the need to treat CFL and PFL as separate lines of research. Since the focus of this survey is CFL, we suggest readers look at Sabah et al. [99], Liu et al. [67] and Ji et al. [53] for more details on PFL.



(a) PFL: Each client trains a personalized model ω_i starting from shared parameters model ω . Unlike standard FL, the learning objective shift to local.

(b) CFL: Clients are grouped into clusters (triangle, circle, square). Each cluster trains separate parameters $\omega^{(k)}$, rather than a single global model.

Fig. 1. Comparison between Personalized Federated Learning (PFL) and Clustered Federated Learning (CFL).

Finally, it is also important not to confuse CFL with what is usually called Federated Clustering [53] (sometimes called unsupervised federated learning). While both involve clustering techniques, Federated Clustering usually refers to distributed clustering methods applied directly to unlabeled data across clients, where the task is unsupervised. On the other hand, the objective of the CFL is not a clustering task in itself; clustering is only used as a mechanism to group clients' models for more effective training. The CFL is designed as agnostic to the learning tasks. Its function does not modify the FL protocol making it compatible with all task compatible with standard FL.

4 Core Clustered Federated Learning for Non-IID Data

To establish a clear frame of reference for further discussions, we begin by outlining the FL framework described in [81] and by specifying the Non-IID data problem as mentioned in that study. Before proceeding, it is important to clarify our article selection criteria for this section. We focus here exclusively on what we call **core CFL**—that is, algorithms designed to handle non-IID data.

4.1 Background on Federated Learning and Data Heterogeneity

Federated Learning (FL) is an iterative method to train a model on a network of devices without sharing raw data. We consider a distributed setting with N clients², $N \geq 2$, coordinated by a central server. Each client $i \in I = \{1, \dots, N\}$ holds a local dataset D_i that cannot be shared with others. The goal of federated learning is to find a set of parameters $\omega \in \mathbb{R}^d$, where d is the number of model

²In this study, we use terms clients or devices interchangeably when speaking about FL participants.

Notations	Description
N	Number of clients in the FL ecosystem, $N \geq 2$
I	$I = \{1, \dots, N\}$ set of all clients
i	Index of client ($i \in I$)
D_i	Local dataset of client i noted as $D_i = \{(x_j^{(i)}, y_j^{(i)}) \mid 1 \leq j \leq n_i\}$, where $(x_j^{(i)})_{1 \leq j \leq n_i}$ and $(y_j^{(i)})_{1 \leq j \leq n_i}$ represents respectively features and target
n_i	Number of samples in D_i . We note $ D_i = n_i$
d	Number of parameters of considered neural networks (NN), $d \in \mathbb{N}^*$
ω	Federated neural network model parameters with $\omega \in \mathbb{R}^d$
ω_i	Local neural network parameters of client i with $\omega_i \in \mathbb{R}^d$
f	FL objective function
f_i	Local objective function for client i defined to optimize a local model on D_i
$(\Omega, \mathcal{F}, \mathcal{P})$	Probability space where Ω is the sample space, representing all possible outcomes of the random experiment associated with sampling from the dataset $D = \cup_{i=1}^N D_i$ \mathcal{F} the associate sigma algebra and \mathcal{P} the probability measure
X	Random variable of the possible outcome of features of D . In this case, each $(x_j^{(i)})_{1 \leq j \leq n_i}$ for $i \in I$ is a realization of X .
Y	Random variable of the possible outcome of targets of D . In this case, each $(y_j^{(i)})_{1 \leq j \leq n_i}$ for $i \in I$ is a realization of Y .
(X, Y)	Joint random variable of the possible outcome of features and targets of D . We have $(X, Y) : \Omega \mapsto D$
$\mathcal{P}_i(X)$	Marginal distribution of dataset D_i features
$\mathcal{P}_i(Y)$	Marginal distribution of dataset D_i target
$\mathcal{P}_i(X, Y)$	Joint distribution of features and targets of dataset D_i
$\mathcal{P}_i(X Y)$	Conditional distribution of features conditioned by target of dataset D_i
$\mathcal{P}_i(Y X)$	Conditional distribution of targets conditioned by features of dataset D_i
K	Number of clusters, $K \leq N$
C_k	k -th cluster with $k \in \{1, \dots, K\}$; a cluster C_k will be considered as a subset of I . Meaning that $C_k \subseteq I$ and $I = \bigcup_{k=1}^K C_k$
F_k	Objective function to optimize the federated model on cluster C_k
$\text{dist}(\cdot, \cdot)$	Generic distance metric
$\mathbb{1}_{i \in C_k}$	Equal to 1 if client i is in cluster C_k , 0 otherwise
μ_k	Cluster C_k representative point essential for cluster formation; a vector of size d calculated using weights ω_i for all $i \in C_k$
η_i	Metadata parameters of client i ; may correspond to local data statistics, distribution statistics, or any type of metadata representation of client i (see Section 4.2.3)
δ_k	Cluster C_k metadata-based representative point essential for cluster formation; This point share the same type as metadata parameters η_i for all clients $i \in C_k$ (see Section 4.2.3)

Table 1. Summary of notations

parameters, that minimizes a global objective function $f(\omega)$ based on all local datasets, all this without sharing D_i , $\forall i \in I$ to ensure data privacy. Letting $n_i = |D_i|$ denote the dataset size of client i and $n = \sum_{j=1}^N n_j$ the total number of samples across all clients. Formally, this global objective can be expressed as:

$$\min_{\omega \in \mathbb{R}^d} f(\omega) := \sum_{i=1}^N \frac{n_i}{n} f_i(\omega), \quad (1)$$

where $f_i(\omega)$ denotes the local objective function of client i . In supervised learning, one typically defines

$$f_i(\omega) = \mathbb{E}_{(X,Y) \sim D_i} [L_i(X, Y, \omega)], \quad (2)$$

that is, the expected loss L_i over samples (X, Y) drawn from the client's data distribution D_i . This formulation is general and applies to a wide range of machine learning models. To solve this objective without compromising privacy, McMahan et al. [81] introduced the FedAvg algorithm. In this protocol, during multiple iterative rounds, each client optimizes $f_i(\omega)$ locally to find parameters ω_i , and the server aggregates the resulting models through the weighted averaging defined in Equation (1). The FedAvg protocol propose that only ω_i (local model parameters) and n_i (local dataset size) can be shared by each client with the server to maintain privacy.

A FL network consists of multiple distributed clients, each owning data that reflects its local context (usage patterns or environment). As a result, the statistical properties of those datasets are inherently heterogeneous: the distributions of D_i across clients may differ significantly. This

heterogeneity, often referred to as non-IID data, is a central challenge for FL [56]. It arises naturally in massively distributed environments where each client captures only a narrow and biased view of the overall population distribution. This statistical heterogeneity has direct implications for the optimization process in FL. When clients train on data drawn from different distributions, their local updates ω_i may point in conflicting directions, making the aggregation of models into a single global parameter vector ω suboptimal, affecting performance. To tackle the non-IID problem, CFL proposes grouping clients with similar distributions into clusters and training a shared model for each cluster.

4.2 Clustered Federated Learning Formulation and taxonomy

We define a generic CFL optimization formulation. Let N clients $i \in \{1, \dots, N\}$ having each a dataset D_i containing $|D_i| = n_i$ data samples. We assume that those samples are derived from K different data probabilistic distributions $\mathcal{P}_1, \dots, \mathcal{P}_K$ with $1 \leq K < N$. This implies that the clients can be partitioned into K disjoint clusters C_1, \dots, C_K where the clients' dataset distributions are homogeneous intra-cluster. For the sake of simplicity in notations, each cluster will be considered as a partition of $I = \{1, \dots, N\}$ with the notation $i \in C_k$ meaning that the client of index i is inside cluster C_k (i.e.: $I = \bigcup_{k=1}^K C_k$). We would then have K different clusters and for $k \in \{1, \dots, K\}$ a corresponding objective function F_k to optimize the FL model on cluster C_k of the form:

$$\min_{\omega \in \mathbb{R}^d} F_k(\omega) := \sum_{i \in C_k} \frac{n_i}{\sum_{j \in C_k} n_j} f_i(\omega) \quad (3)$$

where $\sum_{j \in C_k} n_j$ corresponds to the samples-size's sum of all clients inside cluster C_k and :

$$f_i(\omega) = \mathbb{E}_{(X,Y) \sim D_i} [L_i(X, Y, \omega)] \quad \forall i \in C_k \quad (4)$$

that is the expected value of the local loss function L_i calculated with feature and target (X, Y) following the distribution of dataset D_i with model of parameters $\omega \in \mathbb{R}^d$. Once the clusters are formed, each cluster can independently run a standard FL protocol (e.g., FedAvg) restricted to its members, thereby producing K federated models, each tailored to the data distribution of its cluster.

Both the identification of cluster memberships C_k and the optimization of Equation 3 must be carried out while respecting the original assumptions of FL, where only model weights (or gradients) and dataset sizes are shared. However, it is important to note that some papers [69, 76] relax this assumption by allowing the exchange of additional encrypted metadata. Based on this, a taxonomy is proposed to classify CFL solutions found in the literature, namely:

- *Server-side CFL* (Section 4.2.1), where clustering is performed solely using shared model weights or gradients (Figure 2a).
- *Client-side CFL* (Section 4.2.2), where clustering membership is determined locally by the clients using client training loss (Figure 2b)
- *Metadata-based CFL* (Section 4.2.3), which allows the sharing of additional encrypted metadata with the server for clustering (Figure 2c).

4.2.1 Server-side CFL. The initial approach, referred in this paper as Server-side CFL (see Figure 2a), is used in [3, 10, 32, 80, 103] and proposes to solve Equation 3 by only using model weights or gradient information. This strategy has the advantage of requiring minimal additional setup compared to the baseline FL approach. In the most straightforward version, the central server calculates clusters by applying a k-means algorithm with ℓ_2 (Ghosh et al. [32]) to the locally trained model weights, after which cluster-specific federated models are optimized separately for each cluster using the desired FL protocol.

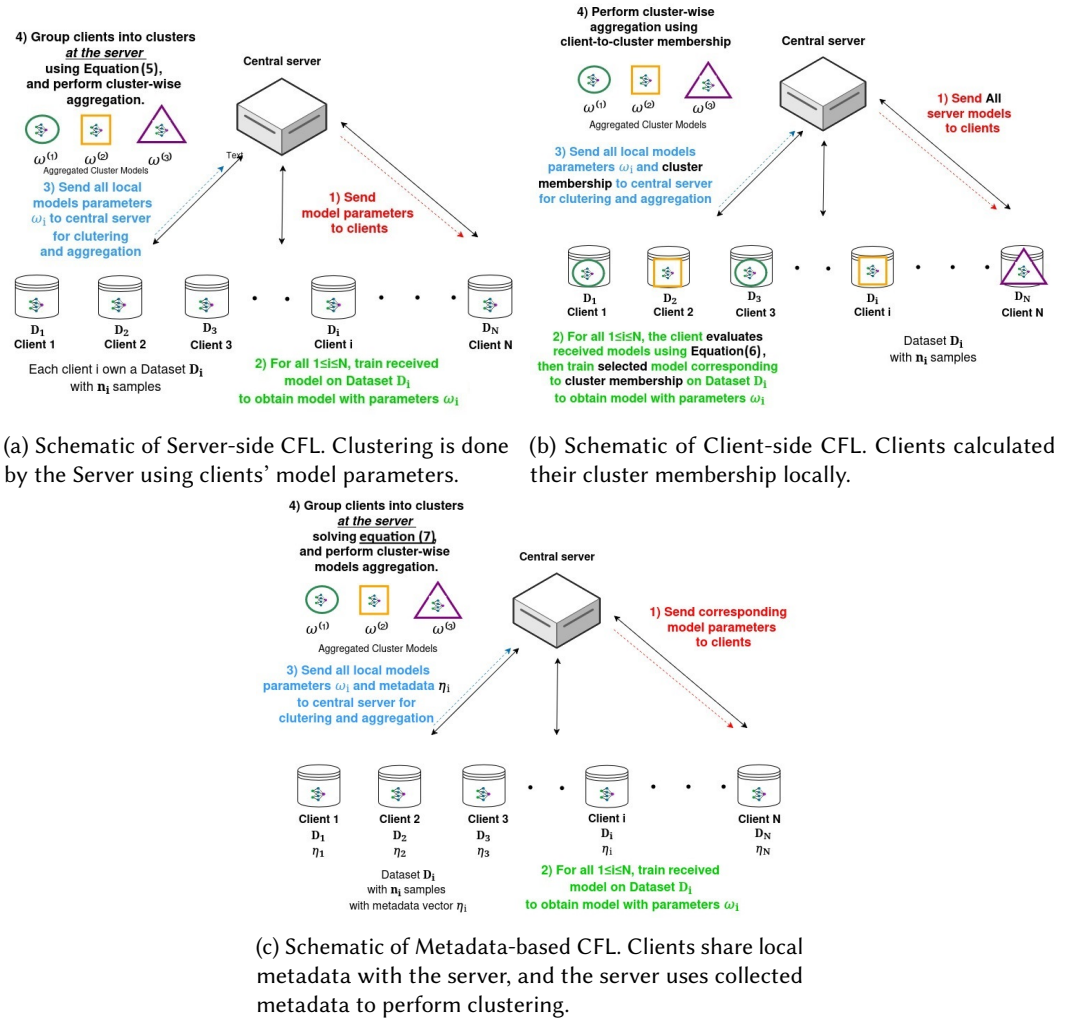


Fig. 2. Comparison of different CFL approaches from taxonomy. In each figure, steps 1 to 4 consist of a full communication round of CFL and can be repeated as needed.

To formalize this, we introduce for each cluster C_k a representative point $\mu_k \in \mathbb{R}^d$, which serves as the cluster centroid in the model-parameter space (i.e., a model representative of the cluster). The central server assigns each client $i \in \{1, \dots, N\}$ to a cluster by selecting the representative point closest to its local model parameters $\omega_i \in \mathbb{R}^d$. This clustering step can be expressed as the following k-means-like optimization problem:

$$\arg \min_{\{\mu_k\}_{k=1}^K} \sum_{k=1}^K \sum_{i=1}^N \mathbb{1}_{i \in C_k} \text{dist}(\omega_i, \mu_k) \quad (5)$$

This approach has been adopted in many studies, with **dist** denoting a generic distance function applied to model parameters or their gradients, and $\mathbb{1}_{i \in C_k} = 1$ if client i belongs to cluster C_k , and 0 otherwise. The distance can also be replaced by a similarity metric, in which case Equation 5

becomes a maximization problem. Early server-side CFL methods, such as those by Sattler et al. [103] (MTCFL), relied on cosine similarity within an iterative bipartitioning procedure, where the most similar pairs are grouped until two coherent subsets naturally emerge, and this process is repeated recursively to find clusters.

A central assumption behind these approaches is that locally trained model weights implicitly encode each client’s data distribution. MTCFL [103] formalizes this idea by treating model updates as surrogates for data heterogeneity and provides a theoretical guarantee that a meaningful bipartition exists when multiple underlying data distributions are present. In this view, model parameters act as high-level summaries of local data, enabling the server to infer clients’ similarity without accessing raw data. These works represent the pioneering approaches of CFL. Since the central server determines the cluster assignment, we refer to this family of methods as **server-side** throughout the remainder of this paper.

4.2.2 Client-side CFL. As an alternative, several studies delegated determination of cluster membership to the clients (see Figure 2b)[33, 34, 60, 72, 78, 98]. In these approaches, clients are presented with a set of candidate cluster representative models, and each decides locally—using its data—which matches best. In the original method IFCA by Ghosh et al. [33], this assignment is performed by selecting the representative model that minimizes the client’s training loss.

Formally, let $\mu_k \in \mathbb{R}^d$ denote the representative model of cluster C_k for all $k \in 1, \dots, K$, which plays a central role in cluster formation. In [33], these representatives are randomly initialized by the server. Each client $i \in 1, \dots, N$ then evaluates all candidates μ_k on its local dataset and selects the one that yields the lowest training loss:

$$\omega_i = \arg \min_{\{\mu_k\}_{k=1}^K} \mathbb{E}_{(X,Y) \sim D_i} [L_i(X, Y, \mu_k)], \quad \forall i \in 1, \dots, N. \quad (6)$$

Each client assigns itself to the μ_k that minimizes the local loss L_i over all samples of dataset D_i . After each client selects a cluster, the representatives are updated at each communication round by aggregating the clients’ trained models within the corresponding cluster [33]. This process can be repeated iteratively until cluster assignments stabilize.

The main advantage of this approach is that it leverages the distributed nature of FL for cluster assignment. As stated by Ghosh et al. [33], “one of the major advantages of [their] algorithm is that it does not require a centralized clustering algorithm, and thus significantly reduces the computational cost at the center machine.”

However, it also presents drawbacks. As argued in [72], these methods introduce higher communication and computation overhead, since clients must download and evaluate multiple global models during the assignment step. This additional burden can become problematic for resource-constrained clients.

Since clients determine the cluster assignment, we call those methods client-side CFL. For clarity, they refer strictly to methods where the clustering computation is carried out entirely on the client side, with no clustering operation performed on the server.

4.2.3 Metadata-based CFL. The third category (see Figure 2c) considered in this study relies on sharing additional clients’ metadata to the server for clustering purposes [20, 30, 69, 76, 115, 146]. For each client $i \in \{1, \dots, N\}$, we define a metadata vector η_i , which may include statistical summaries of the local dataset (e.g., label distributions, feature statistics, or centroids of local data clusters). For each cluster C_k , we then compute a representative point δ_k (typically the centroid of metadata vectors) for all $k \in \{1, \dots, K\}$. Clients membership is determined through a formulation analogous to Equation (5):

$$\arg \min_{\{\delta_k\}_{k=1}^K} \sum_{k=1}^K \sum_{i=1}^N \mathbb{1}_{i \in C_k} \text{dist}(\eta_i, \delta_k). \quad (7)$$

Here, **dist** denotes a generic distance between metadata vectors, and $\mathbb{1}_{i \in C_k} = 1$ if client i is assigned to cluster C_k , and 0 otherwise. Since the cluster assignment is driven by metadata, we refer to these approaches as *metadata-based*. Although Equation (7) is expressed using a distance, an equivalent formulation can be obtained with a similarity function, with the clustering objective becoming a maximization problem.

In terms of design, metadata-based CFL remains structurally similar to server-side approach, as the server still performs the grouping. The key distinction lies in replacing or supplementing model weights, updates, or gradients with metadata explicitly provided by clients. When well-chosen, such metadata can substantially improve clustering quality and reduce computational cost, as metadata vectors are typically far lower-dimensional than model parameters. Another practical advantage is that, unlike model weights or gradients that evolve at every round, metadata can often remain static, which reduces communication overhead since clients do not need to transmit them at every iteration.

However, these benefits come with weaker privacy guarantees, especially when the shared metadata consists of clients' dataset statistics, which reveal sensitive information. Even when protected using security mechanisms such as differential privacy [76], metadata introduces potential privacy and security risks [56]. Such trade-offs depend on the application domain and whether users can explicitly consent to metadata sharing. In sensitive contexts, even aggregated statistics may be infeasible to share due to regulatory or ethical constraints, limiting the applicability of such approaches.

Early examples include Dennis et al. [20], who proposed clustering data locally and sharing the resulting cluster centroids as metadata, and Zhang et al. [146], who combined local label-distribution vectors with model distances. These approaches highlight how leveraging informative, low-dimensional metadata can significantly improve the efficiency and quality of clustering.

4.2.4 Summary of Strength and Weakness of CFL Categories. The comparison between the three CFL categories illustrates that each approach offers a different balance between privacy, computational cost, and clustering effectiveness. **Server-side** approaches align most closely with the original FL assumptions, preserving strong privacy and imposing no additional computation burden on clients. However, centralized clustering on high-dimensional model weights is computationally expensive, leading to scalability issues with large models or many clients. **Client-side** methods also preserve strong privacy while distributing the clustering computation across clients, reducing server computational burden. They impose a higher burden on resource-constrained devices and may increase communication overhead. **Metadata-based** solutions enhance clustering efficiency and effectiveness by operating on lightweight descriptors, reducing clustering costs compared to using model weights. Additionally, unlike model parameters, statistical descriptors are often invariant over time, ensuring stable cluster formation metrics across communication rounds. This efficiency, however, comes at the cost of weaker privacy guarantees, as sharing metadata exceeds the strict constraints of the standard FL exchange protocol.

The choice of an approach thus depends heavily on the application domain and the acceptable trade-off between privacy, computational efficiency, and cluster formation metrics. To summarize these contrasting criteria, Table 2 provides a compact view of the relative performance of each CFL category across five high-level metrics.

CFL Category	Privacy	Server Computation Burden	Client Computation Burden	Communication Efficiency	Cluster metrics Stability
Server-side	++	++	–	+	–
Client-side	++	–	++	–	–
Metadata-based	–	+	–	++	++

Table 2. Expanded comparison of CFL categories across privacy, computation burden, communication efficiency, and metrics stability.

4.3 Review and Categorization of core CFL Solutions

In this section, we focus exclusively on **core CFL approaches that are domain agnostic**—that is, algorithms designed to handle non-IID data without being tied to any task, model architecture, or application domain. All selected articles follow the guidelines described in Section 2. Our objective is to provide a classification of those CFL approaches according to our taxonomy. To clarify this classification, we propose a diagram in figure 3. First, it determines whether clustering relies on metadata shared with the server beyond standard FL assumptions, and secondly, whether the clustering computation is performed entirely by clients or involves the server.




Fig. 3. Taxonomy Diagram for classification of CFL methods.

We provide a high-level synthesis of reviewed papers in Table 3, ordered by publication date. The column “Algorithm” lists the acronym of the approach, either as defined in the original paper or adopted in subsequent works. “CFL Category” summarizes the procedure used for clustering and assigns each method to a category within the taxonomy introduced in Section 4.2. In server-side and metadata-based approaches, it typically refers to a classical clustering algorithm, whereas in client-side methods, it describes how clients determine their cluster membership. The “Metric” column specifies the criterion for similarity between clients in server-side, which determines cluster membership in client-side CFL, and which metadata is used in metadata-based. It is worth noting that several algorithms share similar clustering mechanisms or metrics in Table 3, yet differ in other aspects of the CFL mechanisms discussed in Sections 4.3.2 to 4.3.6. To accurately capture the historical development of CFL, foundational studies are classified in chronological order, according to the date of publication of their first preprint release, as many were cited before formal publication, while later works follow their official publication date. This chronological organization situates the evolution of CFL within a unified timeline, while providing a taxonomy-based categorization and highlighting the key mechanisms of each method.

The earliest theoretical formulation of CFL we identified is by Ghosh et al. [32] in 2019 with a server-side approach. Their approach tackles a non-IID FL setting in the presence of adversarial machines by applying variants of k-means clustering on locally trained models to form clusters during the initial FL round, followed by robust optimization within each cluster to perform federated learning. The first explicit use of the term CFL appears in Sattler et al. [103] in 2020. The authors formalized CFL as a model-agnostic framework where clients are partitioned into groups based on

Table 3. Chronological timeline of CFL approaches with clustering methods (server-side, client-side, or metadata-based) and corresponding evaluation metrics.

Reference	Algorithm	CFL Category			Metrics
		Server-side	Client-side	Metadata-based	
Ghosh et al. 2019 [32] ³	N/A	Trimmed k-means			L_2 norm of models weights
Sattler et al. 2020 [103] ²	MTCFL	Iterative bi-partitioning			Cosine similarity of models gradients
Mansour et al. 2020 [78] ²	HypCluster		Cluster selection by minimum		Local training loss
Briggs et al. 2020 [10] ²	FL+HC	Hierarchical clustering (HC)			L_1/L_2 /Cosine similarity of models update
Ghosh et al. 2020 [33] ²	IFCA		Cluster selection by minimum		Local training loss
Dennis et al. 2021 [20]	k-Fed			K-means	L_2 over local clients data centroids
Li et al. 2021 [60]	FLSC		Cluster selection by minimum		Local training loss
Agrawal et al. 2021 [3]	GeneticCFL			DBSCAN	clients learning rate/batch-size
Marfoq et al. 2021 [80]	FedEM		Expectation-Maximization		Local samples contribution for each model
Duan et al. 2021 [24]	FlexCFL	HC/K-means			MADD-based metric on models update
Jamali-Rad et al. 2022 [51]	FLT			HC	Centroids of Local Autoencoder outputs
Gong et al. 2022 [34]	AutoCFL		Voting scheme		L_2 Model similarity matrix shared with clients
Ruan et al. 2022 [98]	FedSoft		Cluster selection by minimum		Local loss per-sample score
Long et al. 2023 [72]	FeSEM		Cluster selection by minimum		Local model distance to cluster models
Jothimurugesu al 2023 [55]	FedDrift		Cluster selection by minimum		Local training loss
Jingyu et al. 2023 [44]	ASCFL			K-means	Custom metric over clients loss and weights
Yihan et al. 2023 [132]	ICFL	Iterative bipartitioning			SVD of cosine similarity matrix of models weights
Tun et al. 2023 [114]	CP-CFL		Cluster selection by minimum		Loss
Vahidian et al. 2023 [115]	PACFL			HC	Principal angle from SVD-reduced features
Cai et al. 2023 [12]	FedCE		Cluster selection by minimum		Historical loss score (Euclidean)
Morafah et al. 2023 [84]	FLIS	Hierarchical Clustering			Cosine similarity on inferences over public dataset
Ma et al. 2023 [77]	FedCAM		Cluster selection by minimum		Same as FeSEM or IFCA
Liang et al. 2023 [64]	FedEOC	K-means			Cosine similarity of model weights
Mehta et al. 2023 [82]	FLACC	HC			Cosine similarity of model weights
Yu et al. 2024 [143]	FPFC	Hard Thresholding			Penalized Euclidean distance of model weights
Zhang et al. 2024 [146]	MD-ICFL			Min Distance	L_1 norm over weighted softmax outputs
Gao et al. 2024 [30]	FedPEC			GMM	Clients local Shapley values
Zhang et al. 2024 [147]	FedGK	HC			Cosine similarity of weight updates
Liu et al. 2024 [69]	CFLGT	Affinity Propagation			Gradient trajectories (pull/push forces)
Guo et al. 2024 [39]	FedRC		Bi-level optimization		Likelihood-weighted importance + loss
Luo et al. 2024 [76]	pp-CFL			COPRA	Noisy label distribution
Gong et al. 2024 [35]	HiCFL	Iterative Bi-partitioning			Cumulated update cosine similarity per rounds
Ren et al. 2025 [97]	TDCFL			HC	Distance of SVD decomposition of data
Cheng et al. 2025 [16]	SnapCFL	K-Medoids / HC / DBSCAN			Classifier two-sample test
Zeng et al. 2025 [144]	StoCFL	Agglomerative Clustering			Frozen model's Gradient cosine similarity
Guo et al. 2025 [40]	HCFL+		Bi-level optimization		Custom loss function

the cosine similarity of their model updates. The server builds a similarity matrix and applies an iterative bipartitioning process that recursively splits clusters whenever their internal similarity falls below a predefined threshold. Within each cluster, the algorithm merges the most similar client pairs until only two groups remain, forming the partitions. Each resulting cluster is then treated as an independent FL instance, and the partitioning is re-evaluated at every communication round. This approach is often referred to simply as MTCFL (Multi-Task CFL), as each cluster model can be viewed as an independent task. Together, these two early methods established the foundations of server-side CFL.

Later in 2020, HypCluster [78] and IFCA [33], both introduced similar ideas, laying the foundation of client-side CFL. Both propose to clients multiple models which locally select the model that minimizes its training loss (As explained in section 4.2.2). Both papers also introduce theoretical advances for CFL: HypCluster provides a generalization bound showing that learning multiple

cluster-specific models yields a provable trade-off between global generalization and local specialization, whereas IFCA establishes a convergence-rate analysis for its algorithm. However, a key limitation of these methods is their sensitivity to initialization. Poorly chosen initial models may lead to empty clusters, sometimes requiring multiple runs to obtain stable partitions.

In the same year, Briggs et al. [10] proposed FL+HC which proposes the use of **Hierarchical Clustering (HC)** with model update similarity. Different similarities (L_1 , L_2 , cosine) and linkage (rule to measure inter-cluster distance) were evaluated. Instead of iteratively calculating clusters at each communication round, FL+HC proposes to apply a one-shot clustering after a prior predefined number of standard FL rounds, making it more computationally efficient. Another major contribution of this paper is the first explicit use of a non-IID data taxonomy to create an evaluation benchmark for CFL, something we will discuss more thoroughly in Section 4.5. The five aforementioned papers have given a solid foundation for CFL approach to build upon and remain relevant even today. All following papers in the timeline will have at least an experimental comparison to one of these previous CFL articles (see Figure 8 that we will discuss in Section 4.5).

The first domain-agnostic metadata-based CFL approach was introduced by Dennis et al. [20] in 2021 with k-FED, which achieves competitive results to IFCA. In this method, each client performs local k-means clustering on their dataset and sends the resulting centroids to the server, which are interpreted as metadata summarizing local data distributions. The server then applies k-means on those centroids to form groups of clients. Agrawal et al. [2] proposed GeneticCFL, a metadata-based CFL algorithm that clusters clients according to local hyperparameters such as learning rate and batch size. The authors' intuition is that optimal hyperparameters reflect characteristics of data distributions. Clients are first assigned several candidate hyperparameter configurations and select the one that best fits their data. Over time, these hyperparameters evolve through genetic optimization, and the server clusters clients using DBSCAN computed from their Euclidean distances.

From this point onward, the overall framework of CFL becomes more stable, and research gradually shifts from proposing entirely new paradigms to refining and extending existing ones. Most subsequent studies in Table 3 can be seen as evolutions of the foundational IFCA for client-side approaches and MTCFL for server-side ones, or as adaptations of MTCFL that incorporate metadata in the case of metadata-based CFL. These works build on earlier principles to address broader challenges of CFL. Their goals include determining the appropriate number of clusters, improving computational and communication efficiency, introducing soft client assignments or local adaptation mechanisms to better handle diverse non-IID scenarios, and enhancing adaptability to dynamic environments such as newcomer arrival or data drift. Rather than redefining the CFL paradigm, these later contributions focus on making it more practical, flexible, and robust.

4.3.1 Number of Clusters. Like most clustering algorithms, the earliest CFL methods required the number of clusters k to be predefined by the user. In practice, k is often chosen based on domain knowledge or empirically by selecting the value that yields the best performance on the FL task. However, since the number of underlying client data distributions is unobservable, determining this value is challenging and incurs a substantial communication cost. Server-side and metadata-based algorithms benefit from decades of clustering research, as they can directly apply standard clustering techniques that do not require specifying the number of clusters beforehand. Several CFL methods employ algorithms such as HC [10, 16, 115], iterative bipartitioning [103, 132], or DBSCAN [3], which implicitly determine the number of clusters using alternative hyperparameters. However, client-side all relies on a fixed number of clusters except HCFL+ [40], which inputs an initial number of clusters, then uses metrics thresholding to divide or merge clusters.

Nevertheless, all CFL algorithms depend on hyperparameters that influence the number of clusters either directly or indirectly. Thus, although recent methods have reduced reliance on a fixed k , fully automatic and stable cluster determination in CFL remains an open research challenge.

4.3.2 Scalability and Efficiency Enhancement Strategies. Several works focus on improving scalability and efficiency, addressing both computational and communication costs. The problem becomes increasingly severe as the number of clients and model parameters grows. Computing pairwise similarities and performing clustering quickly becomes resource-intensive. These strategies aim to reduce the computational burden, accelerate convergence, and simplify client participation.

It is interesting to note that early CFL approaches such as FL+HC [10] and k-FED [20] proposed one-shot clustering after several rounds of standard FL, whereas most CFL methods are iterative. This approach is computationally efficient since clustering is performed only once. However, in non-metadata-based settings, identifying optimal clusters in a one-shot manner is challenging because the features used for clustering, such as local model weights or losses, evolve across rounds.

In server-side CFL, scalability issues often arise because clustering is performed directly on model, which may contain millions of parameters. To address this **high-dimensional low-sample-size (HDLSS)** challenge, Duan et al. proposed FlexCFL [24]. It adopts MADD-based similarity metrics tailored for HDLSS data. Additionally, truncated **singular value decomposition (SVD)** reduces the dimensionality of the similarity matrix, significantly lowering computational complexity. Similarly, ICFL [132] applies the same decomposition and iterative bipartitioning to improve clustering.

Other approaches focus on computational speed by adjusting clustering dynamics. For instance, Mehta et al. introduced FLACC [82], a greedy clustering strategy that groups clients before full model convergence, reducing training time while maintaining competitive performance. Likewise, FedEOC [64] improves scalability by comparing models through a single representative layer rather than full model parameters.

Client-side CFL also integrates lightweight mechanisms to enhance efficiency. AutoCFL [34] reduces computation by calculating inter-model distances using only partial model weights. FeSEM [72] adopts layer-wise matching aggregation, inspired by [116], to achieve faster convergence and lower training overhead.

Finally, some metadata-based CFL methods improve scalability by operating in a reduced feature space. Examples include ASCFL [44], CP-CFL [114], and PACFL [115], which rely on server-trained or pretrained autoencoders (often on public datasets) to minimize local computational costs while preserving meaningful data representations. As they extract low-dimensional embeddings from client datasets through autoencoders separate from the task model. However, training such models can still be costly from a CFL perspective.

4.3.3 Soft clustering. CFL assumes that devices can be divided into distinct data distributions. Unfortunately, real-world data are often more complex. In some cases, each client's dataset may originate from a mixture of overlapping distributions rather than a single one. Interestingly, none of the server-side approaches studied uses **soft clustering (SC)**. The only metadata-based SC approach is FLT [51], uses a soft HC algorithm. However, the use of SC is technically applicable in both approaches using standard SC algorithms.

For client-side, Li et al. [60] highlight that algorithms like IFCA assume non-overlapping clusters, which leads to inefficient utilization of local information, since a client's knowledge contributes to only one cluster per round. They propose FLSC, a SC version of IFCA. Locally, each client selects several of the best-performing models instead of just one, averages them, and trains the resulting model before sending it back to the server, thereby improving results over IFCA in their evaluations.

FedEM [80] introduces an approach based on the **Expectation-Maximization (EM)** algorithm. The server maintains multiple global models and broadcasts them to all clients. Each client is

associated with all clusters simultaneously and, during the local EM step, computes the contribution of each data sample to each cluster model. These contributions are then used to weight local updates for all models before being sent back to the server, which aggregates them to refine each global model. Similarly, several other algorithms [12, 39, 40, 76, 84, 98] rely on client-derived cluster importance scores to simulate soft clustering.

4.3.4 CFL Dynamics: Client Selection, Newcomers, Distribution Drift, and Re-Clustering. Early CFL strategies, such as Sattler et al. MTCFL [103] assumes full client participation during clustering rounds. However, FL inherently operates in a dynamic environment, making this assumption often unrealistic. Several studies [24, 33, 34, 51, 82, 132] incorporated and evaluated random client selection in their designs. For example, AutoCFL [34] and ICFL [132] propose incremental clustering strategies that preserve past models and cluster memberships of inactive clients, allowing clusters to evolve progressively.

Another important scenario, distinct from standard client selection, is the appearance of newcomers. A newcomer is a client who joins late or even near the end of training. Fortunately, MTCFL has already proposed a newcomer strategy. Since it relies on iterative bipartitioning, the complete division tree of clients is preserved, allowing newcomers to traverse the tree and find their most suitable cluster. Newcomer strategies generally aim to assign new clients to existing clusters. Server-side approaches such as FlexCFL [23] assign newcomers by evaluating their similarity to existing cluster aggregates. FlexCFL also includes a client migration mechanism to handle data distribution drift, using the Wasserstein distance between previous and current local distributions to detect significant shifts and treat affected clients as newcomers. Metadata-based methods such as pp-CFL [76] and TDCFL [97] instead assign newcomers to clusters based on data feature similarity.

Jothimurugesan et al. [55] demonstrated that client-side methods can naturally extend their local cluster-selection mechanisms to handle newcomers dynamically. However, this becomes more challenging when newcomers exhibit novel data distributions. A similar issue arises when existing clients experience data distribution drift, requiring re-clustering. Such clients can be treated as newcomers or may form a new cluster that is later made available to other clients in a typical client-side fashion.

Finally, features used for clustering in both server-side and client-side approaches, such as model weights, gradients, or losses, are not static between communication rounds. This makes CFL fundamentally different from standard clustering, which operates on fixed data points. One consequence is the emergence of misclustered clients or incorrect cluster counts—clients may end up in unsuitable clusters, or clusters themselves may need to be split or merged. In such cases, migration mechanisms like those in FlexCFL [23] and FedDrift [55] allow clients to move between clusters. Furthermore, Guo et al. (2025) [40] proposed HCFL+, a multi-tier generalized CFL architecture encompassing various dynamic re-clustering mechanisms. Proposed approaches can be summarized by highlighting that clusters with large intra-cluster distances can be split, as in MTCFL and ICFL [132], while those with small inter-cluster distances can be merged, as demonstrated in StoCFL [144].

4.3.5 Local Optimization, Regularization and Personalization. As aforementioned, CFL aims to identify and group clients originating from similar probabilistic data distributions. However, clients within the same cluster are rarely perfectly homogeneous, as they are subject to various influences. This lack of intra-cluster homogeneity can still affect training stability and local model adaptation. To mitigate such issues, several studies have integrated local optimization, regularization, and personalization strategies directly into CFL design.

In Gong et al. AutoCFL [34], the authors observe that such heterogeneity may arise from varying dataset sizes. They propose a local training adjustment mechanism that adapts the number of

training epochs according to dataset size. Clients with larger datasets perform more local updates, while those with smaller train for fewer epochs to accelerate convergence.

Similarly, FeSEM [72] extends IFCA by redefining the assignment of clients to clusters based on the distance between local and cluster models, rather than on training loss. It introduces a personalized local objective function with a regularization term that penalizes deviations from the assigned cluster model, thereby balancing personalization and shared cluster learning. Additionally, FedCAM [77] introduces an additive modeling structure for IFCA and FESEM in which each client's model is decomposed into a global component, a cluster-specific component, and a personalized residual. FedCE [12] adopts a similar personalization approach, where the local aggregation of the cluster models to which a client belongs is added as a regularization component.

FedEOC [64] adopts a local optimization strategy where each client splits its local dataset into a support set and a query set. The client first updates its model parameters using the support set, then computes new gradients on the query set using the adapted model, which are sent to the server for clustering. This allows the global model to form efficient one-shot clustering quickly. FPFC [143] reformulates CFL as an optimization problem (ADMM [43] scheme with a penalty regularizes pairwise model distances). This mechanism acts as a form of regularized local optimization that balances model fusion and personalization.

More recently, CFL frameworks such as FedRC [39] and HCFL+ [40] formulate clustering and optimization as a unified bi-level problem. These SC CFL methods jointly optimize locally cluster importance score (section 4.3.3) and model parameters through regularization. These approaches bridge the gap between clustered and personalized FL, showing that personalization and regularization can enhance intra-cluster stability and robustness.

4.3.6 Security Concerns. As data privacy lies at the core of FL, ensuring the security of exchanged information remains crucial. Fortunately, since CFL does not alter the underlying FL communication protocol, standard FL security mechanisms can be directly applied [85]. However, only a few core CFL studies reviewed explicitly address security, although some works have explored it.

Ghosh et al. MTCFL [32] already proposed a defense against Byzantine clients in CFL by filtering out outliers by using a trimmed K-means based on the geometric median. Moreover, they demonstrated that rotation-invariant metrics such as cosine similarity or L_2 norm can be made compatible with encryption schemes like those introduced by Bonawitz et al. [9].

It is worth noting that, as shown in Table 4, metadata-based CFL methods are inherently more vulnerable to privacy leakage. A promising direction is the integration of differential privacy. Luo et al. introduced pp-CFL [76], which shares clients' differentially private label distribution vectors for clustering, achieving improved privacy. Additionally, Ren et al. TDCFL [97] notes that existing privacy-preserving techniques, such as multiparty secure encryption, homomorphic encryption, and differential privacy, can be further integrated to enhance protection against potential risks. However, secure clustering over encrypted data has been studied (Jaschke et al. [52]); its adoption remains impractical because performing encrypted clustering introduces high computational cost, making it incompatible with many real-world CFL setups.

4.3.7 A Privacy-Vulnerability Risk Analysis of Metadata-based CFL. As aforementioned, metadata-based CFL relies on low-dimensional descriptors to encode clients' data distribution. While this is communication-efficient and may improve grouping, it also introduces new privacy concerns. Metadata may leak sensitive information or increase vulnerability to inference attacks. With strict confidentiality, sharing metadata may even be prohibited.

We propose to categorize metadata privacy risk in Table 4. It proposes three risk levels reflecting how much shared metadata exposes the structure of the underlying local dataset and how easily it can be exploited for inference or reconstruction attacks.

Level I: Low Risk (Indirect Metadata). Those represent relatively low risk, such as hyperparameter selection or local model losses. These values mainly indicate computational choices or overall model fit, rather than exposing information about the data.

Level II: Significant Risk (Direct Statistical Signatures). Clients share raw statistical descriptors such as label distributions, centroids, or Shapley values. These act as a direct representation of the dataset and can support property-inference attacks, where an adversary tests whether data points or attributes are present in the training data (Melis et al. [83]). Furthermore, Shapley values explicitly quantify the contribution of each feature to the model’s decision for specific data points. We argue that this could be exploited to infer membership or reveal sensitive information.

In FL, where gradients are already exchanged, releasing such statistics further strengthens privacy leakage. They provide strong information on data, thus significantly improving gradient inversion attacks such as those in Geiping et al. [31].

Level III: High Risk (Reduced Feature Representations).

This is the most privacy-sensitive category. Autoencoder embeddings, SVD projections, or other reduced representations of the full client dataset preserve the topological structure of the raw data to ensure quality clustering. However, research demonstrates that these low-dimensional embeddings are highly susceptible to inversion attacks [22]. Those vectors have a risk of being inverted to recover input samples.

Table 4. Classification of Metadata-based CFL Approaches by Privacy Risk

Risk Level	Shared Metadata Type	Representative Approaches
I. Low	Local Hyperparameters; Local models training loss	GeneticCFL [3], ASCFL [44]
II. Moderate	Data Centroids; Local Sample Contribution Weights; Shapley Values; Noisy labels distribution vectors	k-Fed [20], MD-ICFL [146], FedPEC [30], pp-CFL [76]
III. High	Autoencoder Feature Embeddings; SVD-based Reduced Features	FLT [51], PACFL [115], TDCFL [97]

4.4 Overview of CFL Workflow design

Designing a CFL algorithm involves multiple interacting components that are often presented as a single integrated mechanism in the literature. To motivate a more structured view, Guo et al. (2025) [40] propose a multi-tier CFL system that integrates several mechanisms from IFCA, FedEM, and FedRC. Their architecture combines soft clustering with dynamic strategies such as those discussed in Section 4.3.4.

Intending to provide a unified perspective that encompasses server-side, client-side, and metadata-based CFL approaches for handling non-IID data, and inspired by this architecture, we introduce a generalized CFL workflow that extends their formulation to cover the broader methodological landscape of CFL (see Figure 4). This workflow is designed to clarify and simplify the key aspects that must be considered when designing or analyzing CFL solutions.

This framework consists of five layers, each divided into one or more tiers, and each tier is denoted by a letter for ease of reference:

- **Tier 0 addresses the pre-clustering phase.** A corresponds to a cold start (no specific initialization), B to warm start (clustering begins after several convergence iterations), and C to external initialization (for instance, transfer learning or metadata-guided initialization).
- **Tier 1 defines the global CFL strategy (see Section 4.2).** D denotes server-side clustering based on model similarity, E denotes client-side clustering through local selection among

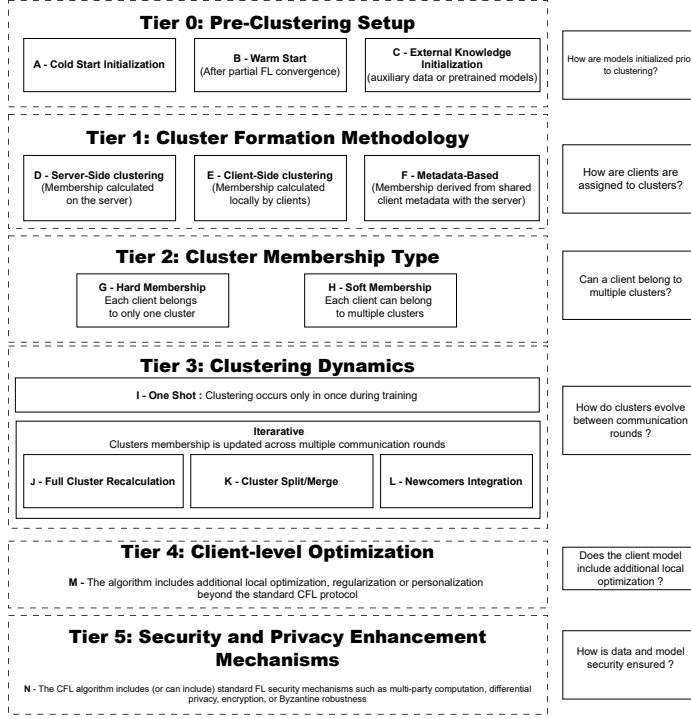


Fig. 4. Generalized CFL System workflow: Each tier captures a specific dimension of the CFL pipeline, enabling structured analysis and comparison of different CFL strategies.

multiple models, and F denotes metadata-based clustering using external features and similarity metrics.

- **Tier 2 defines the global CFL strategy (see Section 4.3.3).** G denotes Hard Clustering Membership, meaning each client belongs to only one cluster, H denotes Soft Clustering Membership, meaning each client can belong to multiple clusters.
- **Tier 3 concerns clustering dynamics (see Section 4.3.4).** Tier I corresponds to one-shot clustering, J to Full Reclustering, K to cluster splitting or merging, and L to newcomer integration.
- **Tier 4 concerns client-level optimization,** and is denoted by M, some client-level optimization mechanism (see Section 4.3.5).
- **Tier 5 concerns security and privacy.** It is denoted by N. Although such mechanisms are rarely included explicitly in CFL studies, they remain essential to the FL paradigm. As discussed in Section 4.3.6, CFL remains fully compatible with standard FL security practices.

Table 5 summarizes studied CFL approaches with their corresponding tier classifications; each algorithm is annotated with the letters corresponding to the mechanisms it implements. It complements Table 3 and aims to highlight which challenges each approach specifically addresses or requires further investigation. Unsurprisingly, tiers 4 and 5 are the least studied in the reviewed CFL approaches, as they remain somewhat complementary to the inter-client separate data distribution, which is the main focus of CFL.

Reference	Algorithm	Tiers
Ghosh et al. 2019 [32]	N/A	A-D-G-I-N
Sattler et al. 2020 [103]	CFL / MTCFL	A-D-G-K-L-N
Mansour et al. 2020 [78]	HypCluster	A-E-G-J
Briggs et al. 2020 [10]	FL+HC	B-D-G-I
Ghosh et al. 2020 [33]	IFCA	A-E-G-J
Dennis et al. 2021 [20]	k-Fed	A-F-G-I-L
Li et al. 2021 [60]	FLSC	A-E-H-J
Agrawal et al. 2021 [3]	GeneticCFL	A-F-G-J-M
Marfoq et al. 2021 [80]	FedEM	A-E-J
Duan et al. 2021 [24]	FlexCFL	A-D-G-J-L
Jamali-Rad et al. 2022 [51]	FLT	A-F-H-J
Gong et al. 2022 [34]	AutoCFL	A-E-G-J-M
Ruan et al. 2022 [98]	FedSoft	A-E-H-J
Long et al. 2023 [72]	FedSEM	A-E-H-J-M
Jothimurugesan et al. 2023 [55]	FedDrift	A-E-G-J
Jingyu et al. 2023 [44]	ASCFL	A-F-G-J
Yihan et al. 2023 [132]	ICFL	A-D-G-K
Tun et al. 2023 [114]	CP-CFL	C-E-G-J

Reference	Algo.	Tiers
Vahidian et al. 2023 [115]	PACFL	A-F-G-I-L
Cai et al. 2023 [12]	FedCE	A-E-H-J-M
Morafah et al. 2023 [84]	FLIS	C-D-G-J
Ma et al. 2023 [77]	FedCAM	A-E-H-J-M
Liang et al. 2023 [64]	FedEOC	A-D-G-I-M
Mehta et al. 2023 [82]	FLACC	A-D-G-J-M
Yu et al. 2024 [143]	FPFC	B-D-G-J-M
Zhang et al. 2024 [146]	MD-ICFL	A-F-G-J
Gao et al. 2024 [30]	FedPEC	A-F-G-J
Zhang et al. 2024 [147]	FedGK	C-D-G-K
Liu et al. 2024 [69]	CFLGT	A-D-G-J
Guo et al. 2024 [39]	FedRC	A-E-H-J-M
Luo et al. 2024 [76]	pp-CFL	A-F-H-J-L-N
Gong et al. 2024 [35]	HiCFL	A-D-G-K
Ren et al. 2025 [97]	TDCFL	A-F-G-J-L
Cheng et al. 2025 [16]	SnapCFL	C-D-G-I
Zeng et al. 2025 [144]	StoCFL	A-D-G-K-M
Guo et al. 2025 [40]	HCFL+	A-E-H-J-K-M

Table 5. CFL approaches with corresponding tier classification...

4.5 Evaluating CFL methods

The evaluation of CFL algorithms requires a rigorous experimental framework comprising three key components: representative non-IID data scenarios, diverse benchmarks, and appropriate performance metrics. In this section, we analyze how the state-of-the-art addresses these challenges.

4.5.1 Formalization of Non-IID Data in FL. Non-Independent and Identically Distributed (non-IID) data give rise to a broad range of statistical challenges. The taxonomy of non-IID heterogeneity, as proposed by Kairouz et al. [56], is often used as the foundation to form interpretable non-IID scenarios for evaluating CFL benchmarks. We therefore rely on the five categories of non-IID data proposed by this taxonomy. Before introducing it formally, we provide an intuitive overview of the five major forms of data heterogeneity in FL. These categories describe how clients' data differ from each other. First, clients may observe the same labels but with inputs that differ visually or structurally (e.g., rotations or style differences), which corresponds to Concept shift on feature. Second, clients may apply different labeling functions to identical inputs, resulting in Concept shift on targets. Clients can also differ in their feature statistics (e.g., brightness or texture, features distribution skew) or in the frequency of classes they observe (label distribution skew). Finally, clients may simply have unequal dataset sizes, a form of heterogeneity known as quantity skew. These five types cover most non-IID situations encountered in practice and clarify how heterogeneity affects learning stability and cluster formation in CFL.

To formalize these situations, we consider a general supervised task where each client $i \in I = \{1, \dots, N\}$ has a local dataset D_i represented as a set of samples $D_i = \{(x_j^{(i)}, y_j^{(i)}) \mid 1 \leq j \leq n_i\}$ where $(x_j^{(i)})_{1 \leq j \leq n_i}$ represents the features and $(y_j^{(i)})_{1 \leq j \leq n_i}$ represents the target and with $|D_i| = n_i$.

Using this notation, we define the random variables X representing the possible outcome of features and Y representing the possible outcome of the targets of all client datasets $(D_i)_{i \in I}$. That means that, in this case, for all $i \in I$, $(x_j^{(i)})_{1 \leq j \leq n_i}$ are realizations of X and $(y_j^{(i)})_{1 \leq j \leq n_i}$ are realizations of Y .

If we denote the probability space as (Ω, \mathcal{F}, P) where Ω is the sample space, representing all possible outcomes of the random experiment associated with sampling from the dataset $D = \cup_{i=1}^N D_i$, \mathcal{F} its associate sigma algebra and P the probability measure, then the joint random variable (X, Y) can be defined as $(X, Y) : \Omega \mapsto D = \cup_{i=1}^N D_i$.

Each client's dataset D_i is represented by its local joint distribution $\mathcal{P}_i(X, Y)$, which we can factor as follows using Bayes' rule: $\mathcal{P}_i(X, Y) = \mathcal{P}_i(Y)\mathcal{P}_i(X|Y) = \mathcal{P}_i(X)\mathcal{P}_i(Y|X)$ where $\mathcal{P}_i(X)$ is the

marginal distribution of the features (respectively $\mathcal{P}_i(Y)$ the target) and $\mathcal{P}_i(Y|X)$ is the distribution of the target conditioned by the features (respectively $\mathcal{P}_i(X|Y)$ is the distribution of features conditioned by the target) for client $i \in I$.




Fig. 5. Illustration of non-IID categories for two clients i and j using MNIST samples.

While this taxonomy was originally developed for classification problems, where the target variable Y is discrete, it naturally extends to general supervised learning settings. However, interpreting these categories becomes less direct when Y is continuous or structured. To make each situation concrete, we illustrate the five categories using MNIST (Figure 5), where different clients' image data exhibit distinct forms of heterogeneity. For two clients $i, j \in I$ with $i \neq j$, datasets D_i and D_j can be considered non-IID if one or more of the following conditions apply:

- a. **Concept shift⁴ on features:** (same target, different features) This occurs when $\mathcal{P}_i(X|Y) \neq \mathcal{P}_j(X|Y)$ even though $\mathcal{P}_i(Y) = \mathcal{P}_j(Y)$. In other words, clients share the same target space data, still the feature representations differ. For example, in Figure 5(a), both clients share the same label distribution $\mathcal{P}(Y)$ (possessing both “3”s and “9”s). However, the conditional distribution $\mathcal{P}(X|Y)$ differs because Client j 's images are rotated relative to Client i 's, meaning the feature representation X for the same label Y has changed.
- b. **Concept shift⁵ on targets:** (same features, different target) This occurs when $\mathcal{P}_i(Y|X) \neq \mathcal{P}_j(Y|X)$ even though $\mathcal{P}_i(X) = \mathcal{P}_j(X)$. This means that, even when clients share the same feature inputs, the target mapping is different. For instance, in Figure 5(b), both clients observe identical feature inputs X (e.g., the image of a “7”). However, their mapping $\mathcal{P}(Y|X)$ diverges: Client i correctly labels it as “7”, while Client j labels the same feature as “1” (label swapping).
- c. **Feature distribution skew:** This occurs when $\mathcal{P}_i(X) \neq \mathcal{P}_j(X)$ even if $\mathcal{P}_i(Y|X) = \mathcal{P}_j(Y|X)$. In this case, clients share the same underlying concepts, but the overall distribution of feature values is skewed. For instance, in Figure 5(c), both clients agree on the concept (Both “0” and “8” look like the actual digits). However, the marginal feature distribution $\mathcal{P}(X)$ differs: Client i exclusively holds images drawn with a thick stroke, while Client j holds images with a thin stroke, representing a skew in the feature domain (e.g., writing style).
- d. **Target⁴ distribution skew:** $\mathcal{P}_i(Y) \neq \mathcal{P}_j(Y)$ even though $\mathcal{P}_i(X|Y) = \mathcal{P}_j(X|Y)$. Here, clients share similar feature-target relationships but differ in the frequencies of target values. For instance, in Figure 5(d), the marginal label distribution $\mathcal{P}(Y)$ varies significantly: Client i 's dataset is dominated by the label “1”, whereas Client j possesses almost exclusively “4”s.
- e. **Quantity skew:** The number of samples varies significantly among clients, such that $|D_i| \ll |D_j|$ or vice versa. This imbalance, illustrated in Figure 5(e), causes clients to contribute unequally to the global model update, potentially biasing the aggregation towards clients with larger local datasets $|D_j|$.


⁴Also referred as “Concept Drift” [39, 55, 56]

⁵On labels in [56]

This taxonomy offers a structured framework to characterize how client data diverges. Since real-world FL data are typically private and inherently non-IID, early CFL algorithms like IFCA [33] already relied on this taxonomy to synthetically generate heterogeneous scenarios from centralized datasets. By partitioning and transforming data, such as applying rotations to image datasets like MNIST or CIFAR, researchers can simulate client groups exhibiting distinct yet interpretable forms of heterogeneity. For instance, assigning rotations of 0° , 90° , 180° , and 270° to different client groups creates a synthetic benchmark for evaluating CFL algorithms in concept shift on features scenarios.

In these settings, traditional FL methods often produce biased models that fail to generalize, whereas CFL algorithms can effectively cluster clients with similar data distributions without direct access to raw samples. An optimal CFL approach should thus separate these distributions into coherent clusters, as illustrated in Figure 6. To make this clearer, the figure contrasts two clustering outcomes under rotation-induced heterogeneity. In the left panel, clients sharing the same rotation are grouped consistently, leading to well-separated clusters that reflect the true underlying data distributions. In the right panel, however, clients with the same rotation are scattered across different clusters, illustrating a poor assignment that fails to capture the structure of the data. This mismatch exemplifies the core challenge in CFL: ensuring that clients with similar distributions are assigned to the same cluster.

Nevertheless, several studies highlight nuanced challenges. Mehta et al. [82] observed that ambiguous cluster boundaries may emerge when feature distribution skew dominates, making cluster separation unclear. Similarly, Guo et al. [39] noted that clustered FL methods sometimes group clients primarily by target(label) similarity, which can lead to overfitting local distributions, particularly in the presence of target(label) distribution skew combined with other heterogeneities.



(a) Optimal clients-to-clusters assignment, groups reflecting underlying data heterogeneity.

(b) Poor clients-to-clusters assignment, failing to reflect underlying heterogeneity.

Fig. 6. Illustration of client grouping in CFL under rotation-induced data heterogeneity: optimal vs. poor clustering

Even though these concepts are useful for classification, extending them beyond discrete targets introduces ambiguity. In tasks such as regression, forecasting, or other supervised problems, the target space is often continuous or multidimensional. As a result, differences in $\mathcal{P}_i(Y)$ or $\mathcal{P}_i(Y|X)$ cannot be categorized discretely.

Developing non-IID benchmarks for such domains may require domain-specific formulations or relies on natural heterogeneity, as seen in distributed IoT setups (e.g. geographical influence or similar usage patterns). In these cases, evaluating CFL becomes more complex, and benchmarks tend to focus mainly on improving performance metrics. Identifying optimal clusters is challenging, even in controlled benchmarks where CFL algorithms are tested with available ground-truth data.

Nevertheless, this taxonomy remains a practical foundation for understanding statistical heterogeneity in CFL. It provides a consistent basis for designing controlled and reproducible experimental setups across various tasks.

4.5.2 Datasets and Tasks. The CFL solutions reviewed in Section 4.3 all provide experimental evaluations to benchmark their approaches. As mentioned in Section 4.5.1, the data heterogeneity taxonomy plays an important role in assessing CFL performance. In this section, we examine the datasets, tasks, and types of heterogeneity used to benchmark these solutions. Specifically, we analyze which datasets have been employed and which heterogeneity categories the proposed methods were evaluated on. This categorization highlights gaps in the literature and points toward potential directions for future research.

Although the non-IID taxonomy is commonly used to transform and distribute centralized datasets into federated settings, many studies do not explicitly map their experimental designs to this taxonomy. Nevertheless, in most cases, their configurations align well with it. Some papers lack a unified terminology to describe similar concepts. For instance, [82] uses the terms concept drift and label distribution skew, while [39] refers to concept shift and label distribution shift, respectively. Such discrepancies may introduce ambiguity for the reader.

To reduce this ambiguity and establish a consistent basis for comparison, Table 6 summarizes the different benchmarks found in the literature, aligns them with the non-IID taxonomy defined in Section 4.5.1. Additionally, it explains how heterogeneity was created for each dataset. We believe that providing a unified framework across studies facilitates clearer comparisons.

In Table 6, the column "*Taxonomy*" shows each non-iid scenario from Kairouz et al. [56]. The "*Task*" column specifies the learning task addressed (e.g., classification, prediction, regression). "*Datasets*" lists the public datasets used (e.g., MNIST, FEMNIST, Shakespeare). The "*Heterogeneity creation*" column describes how data heterogeneity is introduced into datasets, while the "*Papers*" column enumerates the corresponding algorithm/publications.

Table 6. Taxonomy-based view of Tasks, Datasets, and Heterogeneity creation methods in CFL literature.

Taxonomy	Task	Datasets	Heterogeneity Creation	Papers
Concept Shift on Features	Image Class.	MNIST, EMNIST, F-MNIST, CIFAR-10	Image Rotations	IFCA [33], k-FED [20], AutoCFL [34], ASCFL [44], ICFL [132], CP-CFL [114], FedCE [12], FLACC [82], MD-ICFL [146], HiCFL [35], SnapCFL [16], StoCFL [144]
	Image Class.	MNIST, USPS, SVHN, SIGN	Separate Datasets per client	PACFL [115], MD-ICFL [146]
	Regression	Housing, Bodyfat	Separate Datasets per client	FPFC [143]
Concept Shift on Targets	Image Class.	MNIST, F-MNIST, CIFAR-10, CIFAR-100, Tiny-ImageNet, FEMNIST	Label Swapping	MT-CFL [103], FPFC [143], StoCFL [144], FL+HC [10], FedGK [147], FedRC [39], HCFL+ [40], FLACC [82], MD-ICFL [146]
	Rating Pred.	ml-100K, ml-1M	Split by User (Rating bias)	FEDEOC [64]
Feature Distribution Skew	Image Class.	EMNIST, FEMNIST, Fed-CelebA	Natural Split (Users/Writers)	HypCluster [78], k-FED [20], FL+HC [10], IFCA [33], Flex-CFL [24], FedEM [80], FeSEM [72], FLACC [82], FedPEC [30], StoCFL [144]
	Image Class.	EMNIST	Uppercase vs. Lowercase	FedSoft [98], FedCE [12], FLACC [82]
	Image Class.	F-MNIST, CIFAR-10, CIFAR-100, Tiny-ImageNet	Image Corruption	FedRC [39], HCFL+ [40]
	NLP / HAR / Sent. Analysis	Shakespeare, AG News, USCHAD, Sentiment140	Natural Split (Actors/Topic/Users)	k-FED [20], FedEM [80], MTCFL [103], HiCFL [35], FedPEC [30]
Target Distribution Skew	Image Class.	MNIST, EMNIST, FEMNIST, F-MNIST, CIFAR-10, CIFAR-100, STL-10, SVHN, Tiny-ImageNet	Different Label Distribution for each client	FLSC [60], FlexCFL [24], GeneticCFL [3], AutoCFL [34], FLT [51], FPFC [143], MD-ICFL [146], pp-CFL [76], HiCFL [35], SnapCFL [16], FedEM [80], k-FED [20], PACFL [115], FLIS [84], FedCAM [77], FedGK [147], CFLGT [69], FedRC [39], HCFL+ [40], TDCFL [97], CP-CFL [114]
Quantity Skew	Image Class.	MNIST, EMNIST, FEMNIST, F-MNIST, CIFAR-10	Random Sample Size	AutoCFL [34], FLACC [82]

Observations from Table 6 provide insights into how different forms of heterogeneity are generated. Concept shift on features is typically induced by rotating images across clients for image classification tasks or by assigning different but related datasets to different clients (e.g., the Housing and Body Fat datasets). Concept shift on targets often arises from label swapping in classification

problems, except for rating prediction tasks such as ml-100K and ml-1M, where each user naturally exhibits unique preferences. Feature distribution skew is generally inherent to the dataset, stemming from user- or topic-specific distributions. The only artificial case is the use of image corruption by FedRC [39] and HCFL+ [40]. Conversely, target distribution skew is introduced by assigning clients different label distributions within classification tasks. Finally, quantity skew appears in a few works, where clients receive varying numbers of samples without further heterogeneity.

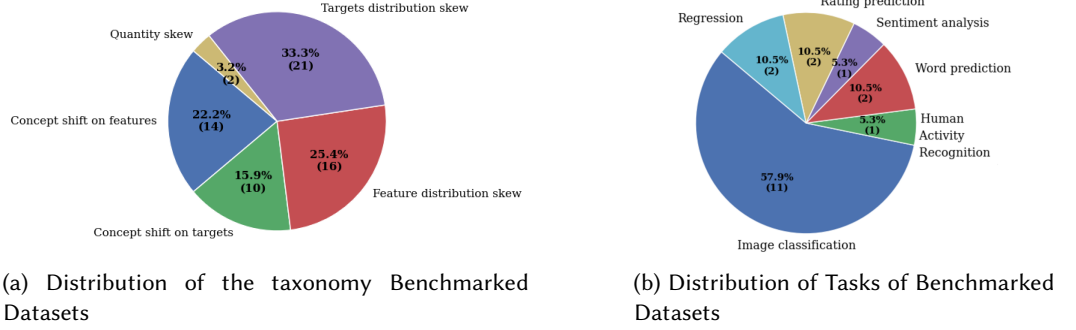


Fig. 7. Distribution of experimental designs across reviewed core CFL papers from Table 3.


Figure 7a illustrates the distribution of heterogeneity types (taxonomy) explored across surveyed CFL algorithms. Each category—such as feature distribution skew, target distribution skew, or concept shift—is represented proportionally to the number of algorithms addressing it. The visualization shows that most studies focus on target or feature distribution skew, while quantity skew remains largely unexplored, having been evaluated only partially in AutoCFL [34] and FLACC [82].

Figure 7b shows the distribution of learning task, emphasizing that image classification dominates as the primary benchmark. Only a few works consider regression, recommendation, or text-based prediction tasks. This imbalance reflects the convenience of applying heterogeneity taxonomies to image classification and highlights the need for broader empirical validation across task types. This trend largely stems from the ease of introducing heterogeneity in image data, such as rotations (concept shift on features) or corruptions (feature distribution skew), as shown for MNIST and CIFAR variants. In contrast, text, time-series, and tabular datasets appear less frequently, indicating a limited exploration of non-vision domains.

4.5.3 Performance and Clustering Metrics. It is interesting to note that the majority of CFL algorithms rely primarily on task-specific metrics—such as average accuracy or RMSE across clients. Additionally, classical convergence indicators like average loss evolution curves are often used to demonstrate communication improvements over baselines. Very few studies introduce metrics specific to the clustering process itself.

Some approaches, such as FeSEM [72], CP-CFL [114], and FedCAM [77], distinguish between *micro* (weighted average based on data amount) and *macro* (simple average across devices) aggregation of these results, though this distinction applies to FL in general rather than CFL specifically.

Regarding clustering-specific behaviors, a limited number of works analyze specific CFL criteria. For instance, FlexCFL [24] tracks the *average discrepancy* (the distance between local and global cluster models), demonstrating that successful clustering lowers this metric over rounds. Similarly, AutoCFL [34] analyzes the variance metrics across clients to validate the reduced heterogeneity within formed clusters.



However, explicit evaluation of cluster structural quality (against a ground truth) remains rare, likely because clustering in CFL is often treated as a heuristic to improve target metrics rather than an end goal. Notable exceptions include ICFL [132], FPFC [143], MD-ICFL [146], and FedPEC [30], which report the **Adjusted Rand Index (ARI)** or **Cluster Purity**, and TDCFL [97] which utilizes the **Silhouette Score**. Importantly, these studies highlight a strong correlation between structural metrics and model performance; for example, Yan et al. [132] observes that algorithms achieving a perfect Purity (perfect recovery of ground-truth partitions) consistently yield the best performances.

4.5.4 Timeline of CFL experimental comparison. To better understand the evolution and benchmarking practices in CFL, we summarize in Figure 8 the timeline of experimental comparisons of algorithms listed in Table 3. Each timeline focuses on a key algorithm (highlighted in red) that has served as a reference point for subsequent works. As shown in the figure, both MTCFL (challenged by 16 follow-up studies) and IFCA (respectively by 25) serve as benchmarks for the majority of algorithms, underlining their major influence and pioneering roles in the development of server-side and client-side CFL strategies. It is noteworthy that no metadata-based methods are challenged in these timelines, likely because such approaches are highly context-dependent, less domain-agnostic, and may raise additional privacy concerns. Among foundational works, HypCluster [78] has only two direct competitors—an understandable outcome given its conceptual similarity to IFCA and its status as a non-peer-reviewed paper—while FL+HC [10] is compared against only four methods.

Among later contributions, FeSEM [72] is the only subsequent method that distinguishes itself, being confronted by eight algorithms. As an early pioneer of soft clustering and personalized CFL, it demonstrated the value of integrating both mechanisms into CFL frameworks. Unfortunately, after that, no algorithm clearly stands out as a reference to be surpassed by subsequent CFL approaches. This illustrates one of the field’s main current challenges: despite incremental progress, as CFL research has matured, the proliferation of specialized algorithms targeting narrow contexts has made it increasingly difficult to identify clear leaders or universally effective methods. Another limitation also emerges when examining experimental comparisons. While many studies benchmark CFL algorithms against previous baselines, the distinction between server-side, client-side, or metadata-based approaches is rarely made explicit. As a result, CFL methods often compare performance metrics such as accuracy or convergence speed without addressing the cost asymmetry between these paradigms. Consequently, the abundance of state-of-the-art solutions built on distinct assumptions and scenarios complicates the selection of an appropriate CFL algorithm for both researchers and practitioners.

4.6 Review of CFL Applications

This section examines how CFL has been applied across specific domains, including IoT, mobility, energy, and healthcare. Unlike domain agnostic methods reviewed in section 4.3, which are predominantly server-side or client-side, applications overwhelmingly favor metadata-based clustering. As shown in Figure 9, metadata-driven approaches are substantially more prevalent, highlighting their practicality in domain-specific settings. This preference for efficiency over strict privacy often results in solutions that are difficult to generalize.

Given the large number of domain-specific studies, a high-level summary and classification according to our taxonomy is provided in a supplement material A.




Fig. 9. Comparison of clustering strategies between domain agnostic (articles reviewed in section 4.3) and applied domain-specific CFL studies (section 4.3).

4.6.1 Internet of Things and Networking. In IoT and networking environments, when device heterogeneity and resource constraints are critical, CFL aims to balance system and statistical heterogeneity. Many methods rely on structural metadata—such as resource efficiency, network topology, spectrum characteristics, or geographic information [4, 59, 111, 121, 130, 136]—to guide clustering decisions. While some works use traditional server-side mechanisms with K-Means or DBSCAN on model weights [13, 18, 25, 37, 45, 142] or IFCA-like approaches [47, 101], client-side solutions remain rare due to the high computational demand on edge devices. Although domain-specific tasks like network anomaly detection and reinforcement learning are explored [7, 25, 100, 124], many applied studies still use conventional image classification benchmarks like MNIST or CIFAR-10 [18, 28, 37, 41, 42, 45, 47, 66, 101, 111, 130, 142, 149]. This domain demonstrates CFL’s adaptability to resource-constrained scenarios, but highlights the tension between system efficiency and models performance.

4.6.2 Mobility and Intelligent Transportation. The mobility domain, which includes vehicular networks and ITS, is characterized by dynamic connectivity and spatiotemporal variability [90]. Metadata-based approaches dominate. They leverage shared information like geographic location, trajectory similarity, or mobility patterns to form clusters for tasks such as traffic forecasting and POI prediction [65, 70, 112, 145]. Some other works also employ embeddings [8, 36] as metadata. Server-side methods often rely on K-means or MTCFL [75, 88, 89, 103, 117, 139, 145]. However, no client-side approaches were identified, even though vehicular and ITS clients may possess higher computational capabilities than standard IoT devices. Although domain-specific datasets (e.g., time-series traffic flow) and tasks (e.g., fault classification, HAR) are used [8, 62, 73], a significant gap remains between the theoretical challenge of mobility and current CFL implementations. Although mobility inherently induces continuous statistical distribution drift (Data distribution change over time), none of the reviewed works employs dynamic mechanisms described in Section 4.3.4. Instead, adaptability in current literature is largely reactive to network topology (like in Taik et al. [112]) changes rather than statistical, highlighting a critical direction for future research in mobility-aware CFL.

4.6.3 Energy. CFL is well-suited to the energy sector, where diverse user profiles naturally emerge. Clients typically correspond to different companies, households, or smart grids, each with its own specific consumption behaviors. CFL can exploit contextual similarities to create specific models.

Most studies use metadata-based CFL with **time-series (TS)** forecasting. Geographical coordinates, daily load profiles, or statistical information are well suited as metadata to detect profiles [5, 21, 46, 68, 104, 107, 119, 140]. For instance, Liue et al. [68] employs privacy-preserving **Dynamic Time Warping (DTW)** for energy resource forecasting, while Ali et al. [5] introduces an explainable CFL framework based on shared local dataset Shapley values for solar power prediction. Other uses CFL on load profiles or Markov transition matrices for short-term energy forecasting and load monitoring [21, 107, 140].

Server-side is less frequent, mostly applied to fault diagnosis and demand forecasting [11, 19, 93, 150]. Client-side remains rare, only emerging in a load monitoring framework (Wang et al. [118]). Metadata-based once again overwhelmingly dominates, underscoring reliance on contextual information to optimize energy models—though often at the cost of reduced privacy.

4.6.4 Healthcare. Healthcare is among the most privacy-sensitive and heterogeneous domains, making CFL a promising paradigm for learning across distributed medical centers [29, 79, 87, 92]. Interestingly, metadata-based approaches remain overrepresented, relying on sharing low-dimensional feature representations for clustering, such as latent representations from local encoders (e.g., ClusterGAN [54]) or dataset-level embeddings [15, 48, 96]. A smaller number of studies

explore server-side clustering strategies inspired by MTCFL [38, 79, 105, 129]. While medical image classification remains a common task [13, 29, 54, 92], CFL has been extended to complex tasks like molecular graph generation [38], 3D medical image segmentation [79], and temporal physiological data analysis [15, 95, 96, 105, 134], demonstrating its flexibility across different healthcare data types.

4.6.5 Summary and Outlook. Beyond the major fields discussed above, CFL is also finding applications in advertising, blockchain, robotics, and speech processing [26, 57, 58, 110, 123, 126, 128, 141]. The broad applicability of CFL across these diverse domains reinforces the real-world utility of the paradigm. However, a key observation is the methodological fragmentation: application-oriented studies overwhelmingly rely on metadata-based clustering strategies, while client-side approaches are almost nonexistent. This preference prioritizes communication efficiency and computational cost over privacy, which is particularly concerning in sensitive domains like healthcare. While these metadata-based approaches provide valuable, context-dependent insights, their reliance on domain-specific auxiliary data makes them difficult to generalize. Future research should encourage the adoption of client-side methods in settings where device computational capacity is sufficient, as they offer competitive performance with stronger privacy guarantees.

5 Clustered X Federated Learning: Beyond Statistical Heterogeneity

One key characteristic of Core Clustered Federated Learning (CFL) is its primary objective: mitigating non-IID data distributions by training multiple specialized models. In contrast, several recent works introduce the notion of **Clustered X Federated Learning** (borrowing the terminology from Liu et al. [53]), where X refers to architectural departures (e.g., decentralized, hierarchical, split, or resource-aware). These variants are fundamentally driven by architectural needs or system efficiency, not by the need to model distinct statistical distributions. We review these variants here to provide a clear disambiguation from Core CFL.

The following categories utilize clustering for operational purposes, contrasting them with the Core CFL objective of producing multiple specialized models.

5.1 Clustered Decentralized FL

These adopt decentralized designs where clients communicate through direct interactions without the presence of a server. As the FL paradigm is changed completely, cluster formation is implicit and emergent, driven by iterative similarity among neighboring models rather than explicit server-side partitioning. Examples include GTV-CFL [102] and FedCBO [14], which encourage local agreements; ACNM [63], which adjusts communication links based on proximity; and blockchain-based coordination methods [131]. While they result in multiple specialized models, their decentralized architecture distinguishes them from the centralized Core CFL paradigm.

5.2 Clustered Hierarchical FL

This category combines clustering with Hierarchical Federated Learning (HFL), where client results are first aggregated locally at an intermediate node before being sent to the global server. Although these approaches borrow CFL-related terminology, the role of clustering is fundamentally different from core CFL: clusters are created to optimize network performance rather than to separate client groups with distinct data distributions. Consequently, these methods universally maintain a **single global model** shared across the hierarchy, failing the multi-model requirement of Core CFL. Clustering in HFL is used to manage operational aspects within the multi-tier architecture. For example, Zhou et al. [151] and Wang et al. [120] use clustering based on social-context information or privacy needs to facilitate participant selection or local aggregation, but the final output is a

single global model. Similarly, Feng et al. [27] group clients based on mobility patterns to stabilize connectivity, while Wen et al. [125] use dynamic clustering to adapt to rapidly changing network topology in Internet-of-Vehicles scenarios. The objective in all these cases is efficient aggregation and scheduling at the edge layer, with no intent to specialize models based on statistical heterogeneity. Overall, these Clustered Hierarchical FL methods do not leverage clustering to model statistical heterogeneity or to produce multiple specialized models.

5.3 Clustered Split FL

This category combines CFL with Split Federated Learning (SFL). In SFL, the neural network is split into two parts: one part runs on each client device, and the other part runs on the server. During training, clients process their data through their portion of the model and send the resulting intermediate output to the server, which continues the forward and backward passes. This setup differs significantly from standard Federated Learning, where each client trains the entire model locally.

Because of this architectural split, these methods fall outside Core CFL. Their purpose is not to identify groups of clients needing different models, but rather to make SFL more efficient. Clustering here is mainly operational: grouping clients helps reduce communication delays, avoid slow devices (stragglers), or better organize how intermediate outputs are exchanged. In all cases, these approaches still aim to train one single global model, not several specialized ones.

Clustered SFL differs from Core CFL in two essential ways. First, the learning process depends on model partitioning rather than full-model aggregation, which changes communication patterns and optimization behavior. Second, clustering is used to manage practical limitations of SFL rather than to address statistical differences among clients. For example, Cheng et al. [17] optimize network topology so that intermediate outputs can be transmitted efficiently, while Wazzeh et al. [122] cluster clients according to their hardware resources to prevent stragglers. Even Arafeh et al. [6], although targeting non-IID IoT environments, still use clustering mainly to make the split architecture more efficient while training a single global model.

5.4 Clustered Resource-Aware FL

This grouping of Clustered X variants focuses specifically on optimizing operational logistics. These methods are explicitly identified by terms such as Resource-Aware, Asynchronous, or Straggler-Aware FL, where the clustering criteria are driven by system-level constraints (e.g., resource availability, timing, or network topology) rather than statistical data distribution. The objective is to uniformly improve the scalability and stability of training a single global model. Several approaches focus on hardware and network limitations. For instance, Abdulrahman et al. [1] introduce CRAS-FL, designed for vehicular networks, where clustering groups vehicles based on communication stability. Similarly, Mughal et al. [86] propose a framework to dynamically assign resource-constrained IoT devices to edge nodes based on resource availability. Other works address timing and scheduling challenges: Liu et al. [71] introduce Straggler-Aware FL, using adaptive clustering to group clients by computation speed to mitigate delays. Qin et al. [94] CAFL groups medical sensors based on communication quality to ensure reliable transmission times. Finally, Lu et al. [74] apply clustering within an Auction-Based incentive mechanism to manage participation costs, while Wu et al. [127] employ Clustered Topology-Based edge selection to optimize data routing paths. Similarly, Hu et al. [49] propose SRP-pFed, which utilizes clustering to optimize communication efficiency within Clustered personalized FL framework. This introduces a 'Sparsified Partial' update strategy where clients are clustered based on their optimal model 'update rate' (the ratio of parameters shared with the server). By grouping clients that require similar levels

of global information, the system significantly reduces communication costs while maintaining a **single global model**.

All the variants reviewed in the section, despite using the term "clustered FL", are fundamentally driven by architectural needs or system efficiency and are not primarily designed to solve the statistical heterogeneity problem through model specialization. By maintaining a **single global model** or by fundamentally altering the standard FL architecture, these Clustered X FL variants serve crucial roles in solving system heterogeneity but lie once again outside the scope of Core CFL for non-IID data specialization.

6 Lessons Learned and Future Directions

Our review reveals a significant gap between CFL theory and practice. While theoretical studies prioritize server-side and client-side clustering for data privacy, real-world applications in IoT, Mobility, and Energy overwhelmingly favor metadata-based approaches because practitioners prioritize efficiency and stability over strict privacy constraints. This separation means general algorithms are rarely tested under real-world constraints, and applications experiments frequently compare new algorithms only to older foundational methods like MTCFL [32] or IFCA [33], rather than modern baselines. Furthermore, ambiguity persists in the terminology, where "Clustered X FL" is used for operational grouping, rather than to solve the core non-IID data problem.

Several key research directions emerge. First, evaluation must evolve toward a multi-criteria comparison framework aligned with the proposed taxonomy. Performance alone is insufficient, as each category imposes distinct trade-offs: server-side methods incur heavy centralized computation, client-side approaches exacerbate communication costs and the straggler effect, and metadata-based methods compromise privacy. Comparative studies must explicitly acknowledge and quantify these trade-offs rather than relying solely on performance metrics. Second, future evaluations must incorporate richer heterogeneity scenarios, with particular regard to Quantity Skew and Target Distribution Skew. Recent studies by Guo et al. [39, 40] highlight that existing CFL approaches struggle with label imbalance, leading to clusters that ignore underlying conceptual relationships. Including personalization methods in CFL may offer a more suitable alternative, and future research should investigate this possibility rigorously. Third, we emphasize the need to decouple the multi-tier CFL framework. Current methods are introduced as indivisible packages, simultaneously specifying initialization, clustering, dynamics, and personalization. Future work should conduct extensive ablation studies to isolate the impact of individual tiers, enabling both more interpretable research and the construction of modular CFL systems tailored to specific settings. Fourth, application-oriented CFL research, which currently centers heavily on metadata-based approaches, must establish a transparent performance–privacy trade-off analysis. Authors should explicitly quantify the privacy implications of each metadata component to justify its use. Conversely, client-side CFL approaches deserve more attention in application domains with substantial device computational capacity, as distributing the clustering burden can substantially improve scalability. Lastly, we encourage authors working on system-oriented variants (Hierarchical FL, Split FL, Resource-Aware FL) to adopt precise terminology rather than the ambiguous Clustered X FL label. Preserving terminological clarity is essential to maintain the historical definition of Core CFL as methods specifically designed to mitigate non-IID heterogeneity by producing multiple specialized models.

7 Conclusion

This survey provides a unified conceptual foundation and a principled taxonomy for Clustered Federated Learning (CFL). We defined Core CFL as the set of methods explicitly designed to mitigate statistical heterogeneity by training multiple specialized models. The proposed taxonomy organizes existing methods into server-side, client-side, and metadata-based categories, detailing

their respective trade-offs: server-side methods preserve strong privacy but lack scalability; client-side approaches distribute computation but burden devices; and metadata-based methods achieve efficiency but compromise privacy. We reviewed complementary CFL challenges, analyzed how heterogeneity taxonomies structure evaluation, and clarified the critical distinction between Core CFL and Clustered X FL variants. The latter, encompassing Hierarchical FL, Split FL, and Resource-Aware FL, employ clustering for architectural or system efficiency reasons while retaining a single global model objective, placing them outside the scope of Core CFL. We hope this unified framework supports the development of the next generation of rigorous and practically impactful CFL research.

References

- [1] Sawsan AbdulRahman, Ouns Bouachir, Safa Otoum, and Azzam Mourad. Cras-fl: Clustered resource-aware scheme for federated learning in vehicular networks. *Vehicular Communications*, 47:100769, 2024.
- [2] Shaashwat Agrawal, Sagnik Sarkar, Mamoun Alazab, Praveen Kumar Reddy Maddikunta, Thippa Reddy Gadekallu, and Quoc-Viet Pham. Genetic cfl: Hyperparameter optimization in clustered federated learning. *Computational Intelligence & Neuroscience*, 2021.
- [3] Shaashwat Agrawal, Sagnik Sarkar, Mamoun Alazab, Praveen Kumar Reddy Maddikunta, Thippa Reddy Gadekallu, and Quoc-Viet Pham. Genetic cfl: Hyperparameter optimization in clustered federated learning. *Computational Intelligence and Neuroscience*, 2021(1):7156420, 2021.
- [4] Amjad Ali and Yevgeni Koucheryavy. Adaptive beam tracking in 5g/6g mmwave networks: A clustered federated learning approach. *IEEE Access*, 2025.
- [5] Syed Saqib Ali, Mazhar Ali, Dost Muhammad Saqib Bhatti, and Bong Jun Choi. Explainable clustered federated learning for solar energy forecasting. *Energies*, 18(9):2380, 2025.
- [6] Mohamad Arafah, Mohamad Wazzeh, Hani Sami, Hakima Ould-Slimane, Chamseddine Talhi, Azzam Mourad, and Hadi Otrouk. Efficient privacy-preserving ml for iot: Cluster-based split federated learning scheme for non-iid data. *Journal of Network and Computer Applications*, 236:104105, 2025.
- [7] KyeongDeok Baek and In-Young Ko. Multifedrl: Efficient training of service agents for heterogeneous internet of things environments. *IEEE Internet of Things Journal*, 2025.
- [8] Asma Belhadi, Youcef Djenouri, Ahmed Nabil Belbachir, Tomasz Michalak, and Gautam Srivastava. Knowledge guided visual transformers for intelligent transportation systems. *IEEE Transactions on Intelligent Transportation Systems*, 2025.
- [9] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In *proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1175–1191, 2017.
- [10] Christopher Briggs, Zhong Fan, and Peter Andras. Federated learning with hierarchical clustering of local updates to improve training on non-iid data. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–9. IEEE, 2020.
- [11] Christopher Briggs, Zhong Fan, and Peter Andras. Federated learning for short-term residential load forecasting. *IEEE open access journal of power and energy*, 9:573–583, 2022.
- [12] Luxin Cai, Naiyue Chen, Yuanzhouhan Cao, Jiahuan He, and Yidong Li. Fedce: Personalized federated learning method based on clustering ensembles. In *Proceedings of the 31st ACM international conference on multimedia*, pages 1625–1633, 2023.
- [13] Yulin Cao, Jianping Ma, Zaobo He, and Yingshu Li. Ap-cfl: Clustered federated learning through dynamic clustering and adaptive participation in heterogeneous iot. *IEEE Internet of Things Journal*, 2025.
- [14] Jose A Carrillo, Nicolas Garcia Trillos, Sixu Li, and Yuhua Zhu. Fedcbo: Reaching group consensus in clustered federated learning through consensus-based optimization. *Journal of machine learning research*, 25(214):1–51, 2024.
- [15] Muhammad Asaad Cheema, Muhammad Zohaib Sarwar, Daniel Cantero, and Pierluigi Salvo Rossi. Clustered federated learning for population-based structural health monitoring. *IEEE Internet of Things Journal*, 2025.
- [16] Yujun Cheng, Weiting Zhang, Zhewei Zhang, Jiawen Kang, Qi Xu, Shengjin Wang, and Dusit Niyato. Snapcfl: A pre-clustering-based clustered federated learning framework for data and system heterogeneities. *IEEE Transactions on Mobile Computing*, 2025.
- [17] Zhipeng Cheng, Xiaoyu Xia, Minghui Liwang, Xuwei Fan, Yanglong Sun, Xianbin Wang, and Lianfen Huang. Cheese: Distributed clustering-based hybrid federated split learning over edge networks. *IEEE Transactions on Parallel and Distributed Systems*, 34(12):3174–3191, 2023.
- [18] Wei-Che Chien, Chih-Hsun Lin, Tianli Zhu, Cheng Dai, Sahil Garg, and Amrit Mukherjee. Density-clustering aggregation for personalized federated learning with ai-enabled aerial and edge computing in uavs. *IEEE Internet of*

Things Journal, 2025.

- [19] Tianxu Cui, Ying Shi, Bo Lv, Rijia Ding, and Xianqiang Li. Federated learning with sarima-based clustering for carbon emission prediction. *Journal of cleaner production*, 426:139069, 2023.
- [20] Don Kurian Dennis, Tian Li, and Virginia Smith. Heterogeneity for the win: One-shot federated clustering. In *International Conference on Machine Learning*, pages 2611–2620. PMLR, 2021.
- [21] Atharvan Dogra, Ashima Anand, and Jatin Bedi. Consumers profiling based federated learning approach for energy load forecasting. *Sustainable Cities and Society*, 98:104815, 2023.
- [22] Alexey Dosovitskiy and Thomas Brox. Inverting visual representations with convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4829–4837, 2016.
- [23] Moming Duan, Duo Liu, Xinyuan Ji, Renping Liu, Liang Liang, Xianzhang Chen, and Yajuan Tan. Fedgroup: Efficient federated learning via decomposed similarity-based clustering. In *2021 IEEE Intl Conf on Parallel & Distributed Processing with Applications, Big Data & Cloud Computing, Sustainable Computing & Communications, Social Computing & Networking (ISPA/BDCloud/SocialCom/SustainCom)*, pages 228–237. IEEE, 2021.
- [24] Moming Duan, Duo Liu, Xinyuan Ji, Yu Wu, Liang Liang, Xianzhang Chen, Yajuan Tan, and Ao Ren. Flexible clustered federated learning for client-level data distribution shift. *IEEE Transactions on Parallel and Distributed Systems*, 33(11): 2661–2674, 2021.
- [25] Jiamin Fan, Kui Wu, Guoming Tang, Yang Zhou, and Shengqiang Huang. Taking advantage of the mistakes: Rethinking clustered federated learning for iot anomaly detection. *IEEE Transactions on Parallel and Distributed Systems*, 35(6): 862–876, 2024.
- [26] Bahar Farahani, Shima Tabibian, and Hamid Ebrahimi. Toward a personalized clustered federated learning: A speech recognition case study. *IEEE Internet of Things Journal*, 10(21):18553–18562, 2023.
- [27] Chenyuan Feng, Howard H Yang, Deshun Hu, Zhiwei Zhao, Tony QS Quek, and Geyong Min. Mobility-aware cluster federated learning in hierarchical wireless networks. *IEEE Transactions on Wireless Communications*, 21(10): 8441–8458, 2022.
- [28] Reza Fotohi, Fereidoon Shams Aliee, and Bahar Farahani. A lightweight and secure deep learning model for privacy-preserving federated learning in intelligent enterprises. *IEEE Internet of Things Journal*, 11(19):31988–31998, 2024.
- [29] Tianrun Gao, Keyan Liu, Yuning Yang, Xiaohong Liu, Ping Zhang, and Guangyu Wang. Fedpc: An efficient prototype-based clustered federated learning on medical imaging. *IEEE Journal of Biomedical and Health Informatics*, 2025.
- [30] Yuan Gao, Ziyue Lin, Maoguo Gong, Yuanqiao Zhang, and Yihong Zhang. Enhancing federated learning with pattern-based client clustering. *IEEE Internet of Things Journal*, 2024.
- [31] Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. Inverting gradients-how easy is it to break privacy in federated learning? *Advances in neural information processing systems*, 33:16937–16947, 2020.
- [32] Avishek Ghosh, Justin Hong, Dong Yin, and Kannan Ramchandran. Robust federated learning in a heterogeneous environment. *arXiv preprint arXiv:1906.06629*, 2019.
- [33] Avishek Ghosh, Jichan Chung, Dong Yin, and Kannan Ramchandran. An efficient framework for clustered federated learning. *Advances in Neural Information Processing Systems*, 33:19586–19597, 2020.
- [34] Biyao Gong, Tianzhang Xing, Zhidan Liu, Wei Xi, and Xiaojiang Chen. Adaptive client clustering for efficient federated learning over non-iid and imbalanced data. *IEEE Transactions on Big Data*, 2022.
- [35] Biyao Gong, Tianzhang Xing, Zhidan Liu, Wei Xi, and Xiaojiang Chen. Towards hierarchical clustered federated learning with model stability on mobile devices. *IEEE Transactions on Mobile Computing*, 23(6):7148–7164, 2023.
- [36] Shreya Gummadi, Mateus V Gasparino, Deepak Vasisht, and Girish Chowdhary. Fed-ec: Bandwidth-efficient clustering-based federated learning for autonomous visual robot navigation. *IEEE Robotics and Automation Letters*, 2024.
- [37] Jialin Guo, Zhetao Li, Anfeng Liu, Xiong Li, and Ting Chen. Rec-fed: a robust and efficient clustered federated system for dynamic edge networks. *IEEE Transactions on Mobile Computing*, 2024.
- [38] Yan Guo, Yongqiang Gao, and Jiawei Song. Molcfl: A personalized and privacy-preserving drug discovery framework based on generative clustered federated learning. *Journal of biomedical informatics*, 157:104712, 2024.
- [39] Yongxin Guo, Xiaoying Tang, and Tao Lin. Fedrc: tackling diverse distribution shifts challenge in federated learning by robust clustering. In *Proceedings of the 41st International Conference on Machine Learning, ICML’24*. JMLR.org, 2024.
- [40] Yongxin Guo, Xiaoying Tang, and Tao Lin. Enhancing clustered federated learning: Integration of strategies and improved methodologies. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [41] Moqbel Hamood, Abdullatif Albaser, Mohamed Abdallah, and Ala Al-Fuqaha. Clustered and multi-task federated distillation for heterogeneous and resource constrained industrial iot applications. *IEEE Internet of Things Magazine*, 6(2):64–69, 2023.

- [42] Moqbel Hamood, Abdullatif Albaseer, Mohamed Abdallah, and Ala Al-Fuqaha. Efficient data labeling and optimal device scheduling in hwns using clustered federated semi-supervised learning. *IEEE Transactions on Communications*, 2024.
- [43] De-Ren Han. A survey on some recent developments of alternating direction method of multipliers. *Journal of the Operations Research Society of China*, 10(1):1–52, 2022.
- [44] Jingyi He, Biyao Gong, Jiadi Yang, Hai Wang, Pengfei Xu, and Tianzhang Xing. Ascl: Accurate and speedy semi-supervised clustering federated learning. *Tsinghua Science and Technology*, 28(5):823–837, 2023.
- [45] Wenji He, Haipeng Yao, Tianle Mai, Fu Wang, and Mohsen Guizani. Three-stage stackelberg game enabled clustered federated learning in heterogeneous uav swarms. *IEEE Transactions on Vehicular Technology*, 72(7):9366–9380, 2023.
- [46] Yu He, Fengji Luo, Mingyang Sun, and Gianluca Ranzi. Privacy-preserving and hierarchically federated framework for short-term residential load forecasting. *IEEE Transactions on Smart Grid*, 14(6):4409–4423, 2023.
- [47] Zaobo He, Lintao Wang, and Zhipeng Cai. Clustered federated learning with adaptive local differential privacy on heterogeneous iot data. *IEEE Internet of Things Journal*, 11(1):137–146, 2023.
- [48] Wei-Tung Hsu, Chin-Po Chen, Yun-Shao Lin, and Chi-Chun Lee. A cluster-based personalized federated learning strategy for end-to-end asr of dementia patients. In *Proc Interspeech*, volume 2024, pages 2450–2454, 2024.
- [49] Xinyi Hu, Zihan Chen, Chenyuan Feng, Geyong Min, Tony QS Quek, and Howard H Yang. Sparsified random partial model update for personalized federated learning. *IEEE Transactions on Mobile Computing*, 2024.
- [50] Xiaoge Huang, Zhi Chen, Qianbin Chen, and Jie Zhang. Federated learning based qos-aware caching decisions in fog-enabled internet of things networks. *Digital Communications and Networks*, 9(2):580–589, 2023.
- [51] Hadi Jamali-Rad, Mohammad Abdizadeh, and Anuj Singh. Federated learning with taskonomy for non-iid data. *IEEE transactions on neural networks and learning systems*, 34(11):8719–8730, 2022.
- [52] Angela Jäschke and Frederik Armknecht. Unsupervised machine learning on encrypted data. In *International conference on selected areas in cryptography*, pages 453–478. Springer, 2018.
- [53] Shaoxiong Ji, Yue Tan, Teemu Saravirta, Zhiqin Yang, Yixin Liu, Lauri Vasankari, Shirui Pan, Guodong Long, and Anwar Walid. Emerging trends in federated learning: From model fusion to federated x learning. *International Journal of Machine Learning and Cybernetics*, 15(9):3769–3790, 2024.
- [54] Shiyi Jiang, Farshad Firouzi, and Krishnendu Chakrabarty. Low-overhead clustered federated learning for personalized stress monitoring. *IEEE Internet of Things Journal*, 11(3):4335–4347, 2023.
- [55] Ellango Jothimurugesan, Kevin Hsieh, Jianyu Wang, Gauri Joshi, and Phillip B Gibbons. Federated learning under distributed concept drift. In *International Conference on Artificial Intelligence and Statistics*, pages 5834–5853. PMLR, 2023.
- [56] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2):1–210, 2021.
- [57] Timm Koppelman, Luca Becker, Alexandru Nelus, Rene Glitz, Lea Schönherr, and Rainer Martin. Clustering-based wake word detection in privacy-aware acoustic sensor networks. In *Interspeech* 2022, pages 719–723, 2022.
- [58] Chae-Won Lee, Jae-Hong Lee, and Joon-Hyuk Chang. Language model personalization for speech recognition: A clustered federated learning approach with adaptive weight average. *IEEE Signal Processing Letters*, 2024.
- [59] Jaewook Lee and Haneul Ko. Energy and distribution-aware cooperative clustering algorithm in internet of things (iot)-based federated learning. *IEEE Transactions on Vehicular Technology*, 72(10):13799–13804, 2023.
- [60] Chengxi Li, Gang Li, and Pramod K Varshney. Federated learning with soft clustering. *IEEE Internet of Things Journal*, 9(10):7773–7782, september 2021.
- [61] Pei Li, Zhijun Liu, Jialiang Peng, Lili Zhou, Tongjun Liu, Yi Wu, and Ahmed A Abd El-Latif. Personalized federated learning for goods classification in wireless sensor-assisted smart warehouses. *IEEE Transactions on Consumer Electronics*, 70(1):4178–4186, 2024.
- [62] Xiaochen Li, Sicong Liu, Zimu Zhou, Bin Guo, Yuan Xu, and Zhiwen Yu. Echopfl: Asynchronous personalized federated learning on mobile devices with on-demand staleness control. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 8(1):1–22, 2024.
- [63] Zexi Li, Jiaxun Lu, Shuang Luo, Didi Zhu, Yunfeng Shao, Yinchuan Li, Zhimeng Zhang, Yongheng Wang, and Chao Wu. Towards effective clustered federated learning: A peer-to-peer framework with adaptive neighbor matching. *IEEE Transactions on Big Data*, 10(6):812–826, 2022.
- [64] Tingting Liang, Cheng Yuan, Cheng Lu, Youhuizi Li, Junfeng Yuan, and Yuyu Yin. Efficient one-off clustering for personalized federated learning. *Knowledge-Based Systems*, page 110813, 2023.
- [65] Zhuhua Liao, Shoubin Li, Yijiang Zhao, Yizhi Liu, Wei Liang, and Shaohua Wan. Predicting ride-hailing passenger demand: a poi-based adaptive clustering federated learning approach. *Future Generation Computer Systems*, 156: 168–178, 2024.

- [66] Yushen Lin, Kaidi Wang, and Zhiguo Ding. Rethinking clustered federated learning in noma enhanced wireless networks. *IEEE Transactions on Wireless Communications*, 2024.
- [67] Bingyan Liu, Nuoyan Lv, Yuanchun Guo, and Yawen Li. Recent advances on federated learning: A systematic survey. *Neurocomputing*, 597:128019, 2024.
- [68] Linbin Liu, June Li, and Juan Wang. Cfl-iccv: Clustered federated learning framework with an intra-cluster cross-validation mechanism for der forecasting. *Applied Energy*, 377:124699, 2025.
- [69] Ruiqi Liu, Songcan Yu, Linsi Lan, Junbo Wang, Krishna Kant, and Neville Calleja. A remedy for heterogeneous data: Clustered federated learning with gradient trajectory. *Big Data Mining and Analytics*, 7(4):1050–1064, 2024.
- [70] Yi Liu, JQ James, Jiawen Kang, Dusit Niyato, and Shuyu Zhang. Privacy-preserving traffic flow prediction: A federated learning approach. *IEEE Internet of Things Journal*, 7(8):7751–7763, 2020.
- [71] Yi-Jing Liu, Gang Feng, Hongyang Du, Zheng Qin, Yao Sun, Jiawen Kang, Xiaoqian Li, and Dusit Niyato. Adaptive clustering-based straggler-aware federated learning in wireless edge networks. *IEEE Transactions on Communications*, 72(12):7757–7771, 2024.
- [72] Guodong Long, Ming Xie, Tao Shen, Tianyi Zhou, Xianzhi Wang, and Jing Jiang. Multi-center federated learning: clients clustering for better personalization. *World Wide Web*, 26(1):481–500, 2023.
- [73] Hao Lu, Adam Thelen, Olga Fink, Chao Hu, and Simon Laflamme. Federated learning with uncertainty-based client clustering for fleet-wide fault diagnosis. *Mechanical Systems and Signal Processing*, 210:111068, 2024.
- [74] Renhao Lu, Weizhe Zhang, Yan Wang, Qiong Li, Xiaoxiong Zhong, Hongwei Yang, and Desheng Wang. Auction-based cluster federated learning in mobile edge computing systems. *IEEE Transactions on Parallel and Distributed Systems*, 34(4):1145–1158, 2023.
- [75] Shoupeng Lu, Hao Yin, Peng Wang, Tianli Zhu, Cheng Dai, Shengxin Dai, and Bing Guo. Clustered federated learning with adaptive pruning for 6g edge-intelligent transportation. *IEEE Transactions on Intelligent Transportation Systems*, 2025.
- [76] Guixun Luo, Naiyue Chen, Jiahuan He, Bingwei Jin, Zhiyuan Zhang, and Yidong Li. Privacy-preserving clustering federated learning for non-iid data. *Future Generation Computer Systems*, 154:384–395, 2024.
- [77] Jie Ma, Tianyi Zhou, Guodong Long, Jing Jiang, and Chengqi Zhang. Structured federated learning through clustered additive modeling. *Advances in Neural Information Processing Systems*, 36:43097–43107, 2023.
- [78] Yishay Mansour, Mehryar Mohri, Jae Ro, and Ananda Theertha Suresh. Three approaches for personalization with applications to federated learning. *arXiv preprint arXiv:2002.10619*, 2020.
- [79] Matthis Manthe, Stefan Duffner, and Carole Lartizien. Federated brain tumor segmentation: An extensive benchmark. *Medical Image Analysis*, 97:103270, 2024.
- [80] Othmane Marfoq, Giovanni Neglia, Aurélien Bellet, Laetitia Kameni, and Richard Vidal. Federated multi-task learning under a mixture of distributions. *Advances in neural information processing systems*, 34:15434–15447, 2021.
- [81] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017.
- [82] Manan Mehta and Chenhui Shao. A greedy agglomerative framework for clustered federated learning. *IEEE Transactions on Industrial Informatics*, 2023.
- [83] Luca Melis, Congzheng Song, Emiliano De Cristofaro, and Vitaly Shmatikov. Exploiting unintended feature leakage in collaborative learning. In *2019 IEEE symposium on security and privacy (SP)*, pages 691–706. IEEE, 2019.
- [84] Mahdi Morafah, Saeed Vahidian, Weijia Wang, and Bill Lin. Flis: Clustered federated learning via inference similarity for non-iid data distribution. *IEEE Open Journal of the Computer Society*, 4:109–120, 2023.
- [85] Virraji Mothukuri, Reza M Parizi, Seyedamin Pouriyeh, Yan Huang, Ali Dehghantanha, and Gautam Srivastava. A survey on security and privacy of federated learning. *Future Generation Computer Systems*, 115:619–640, 2021.
- [86] Fahad Razaque Mughal, Jingsha He, Bhagwan Das, Fayaz Ali Dharejo, Nafei Zhu, Surbhi Bhatia Khan, and Saeed Alzahrani. Adaptive federated learning for resource-constrained iot devices through edge intelligence and multi-edge clustering. *Scientific Reports*, 14(1):28746, 2024.
- [87] Ke Niu, Wenjuan Tai, Jiuyun Cai, Yuhang Zhou, and Heng Li. Fedcgp: Cluster-based gradual personalization for federated medical image segmentation. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2025.
- [88] Xiaomin Ouyang, Zhiyuan Xie, Jiayu Zhou, Jianwei Huang, and Guoliang Xing. Clusterfl: a similarity-aware federated learning system for human activity recognition. In *Proceedings of the 19th annual international conference on mobile systems, applications, and services*, pages 54–66, 2021.
- [89] Xiaomin Ouyang, Zhiyuan Xie, Jiayu Zhou, Guoliang Xing, and Jianwei Huang. Clusterfl: A clustering-based federated learning system for human activity recognition. *ACM Transactions on Sensor Networks*, 19(1):1–32, 2022.
- [90] Jiaming Pei and Wei Li. Unveiling the effects of slightly skewed labels on traffic data analysis. *IEEE Transactions on Intelligent Transportation Systems*, 2024.

- [91] Kilian Pfeiffer, Martin Rapp, Ramin Khalili, and Jörg Henkel. Federated learning for computationally-constrained heterogeneous devices: A survey. *ACM Computing Surveys*, 2023.
- [92] Adnan Qayyum, Kashif Ahmad, Muhammad Ahtazaz Ahsan, Ala Al-Fuqaha, and Junaid Qadir. Collaborative federated learning for healthcare: Multi-modal covid-19 diagnosis at the edge. *IEEE Open Journal of the Computer Society*, 3: 172–184, 2022.
- [93] Dalin Qin, Guobing Liu, Zengxiang Li, Weicheng Guan, Shubao Zhao, and Yi Wang. Federated deep contrastive learning for mid-term natural gas demand forecasting. *Applied Energy*, 347:121503, 2023.
- [94] Peng Qin, Yifei Wang, and Xianchao Zhang. Collaborative cloud-edge computing with wpt for air-ground integrated iomt network: A urllc aware cafl-based approach. *IEEE Transactions on Communications*, 2025.
- [95] Xihang Qiu, Wanyong Qiu, Ye Zhang, Kun Qian, Chun Li, Bin Hu, Björn W Schuller, and Yoshiharu Yamamoto. Fedkdc: Consensus-driven knowledge distillation for personalized federated learning in eeg-based emotion recognition. *IEEE Journal of Biomedical and Health Informatics*, 2025.
- [96] Houda Rafi, Yannick Benezeth, Fan Yang Song, Philippe Reynaud, Emmanuel Arnoux, and Cedric Demonceaux. Tree-based personalized clustered federated learning: A driver stress monitoring through physiological data case study. *IEEE Internet of Things Journal*, 2024.
- [97] Tian Ren, Siyao Cheng, Hao Zhang, and Jie Liu. Dynamic clustered federated learning via adaptive distribution similarity computation. *Computer Networks*, page 111302, 2025.
- [98] Yichen Ruan and Carlee Joe-Wong. Fedsoft: Soft clustered federated learning with proximal local updating. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8124–8131, 2022.
- [99] Fahad Sabah, Yuwen Chen, Zhen Yang, Muhammad Azam, Nadeem Ahmad, and Raheem Sarwar. Model optimization techniques in personalized federated learning: A survey. *Expert Systems with Applications*, 243:122874, 2024.
- [100] Xabier Sáez-de Cámara, Jose Luis Flores, Cristóbal Arellano, Aitor Urbieto, and Urko Zurutuza. Clustered federated learning architecture for network anomaly detection in large scale heterogeneous iot networks. *Computers & Security*, 131:103299, 2023.
- [101] Hasin Us Sami and Bacsak Guler. Over-the-air clustered federated learning. *IEEE Transactions on Wireless Communications*, 23(7):7877–7893, 2023.
- [102] Yasmin SarcheshmehPour, Yu Tian, Linli Zhang, and Alexander Jung. Clustered federated learning via generalized total variation minimization. *IEEE Transactions on Signal Processing*, 71:4240–4256, 2023.
- [103] Felix Sattler, Klaus-Robert Müller, and Wojciech Samek. Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints. *IEEE transactions on neural networks and learning systems*, 32(8): 3710–3722, 2020.
- [104] Marco Savi and Fabrizio Olivadese. Short-term energy consumption forecasting at the edge: A federated learning approach. *IEEE access*, 9:95949–95969, 2021.
- [105] Thanveer Shaik, Xiaohui Tao, Lin Li, Niall Higgins, Raj Gururajan, Xujuan Zhou, and Jianming Yong. Clustered fedstack: Intermediate global models with bayesian information criterion. *Pattern Recognition Letters*, 177:121–127, 2024.
- [106] Yao Shan, Yu Yao, Xiaoming Zhou, Tong Zhao, Bo Hu, and Lei Wang. Cfl-ids: An effective clustered federated learning framework for industrial internet of things intrusion detection. *IEEE Internet of Things Journal*, 11(6):10007–10019, 2023.
- [107] Yunchuan Shi, Wei Li, Xiaomin Chang, Ting Yang, Yaojie Sun, and Albert Y Zomaya. On enabling collaborative non-intrusive load monitoring for sustainable smart cities. *Scientific Reports*, 13(1):6569, 2023.
- [108] Chitraranjan Singh, Rahul Mishra, Hari Prabhat Gupta, and Garvit Banga. A federated learning-based patient monitoring system in internet of medical things. *IEEE Transactions on Computational Social Systems*, 10(4):1622–1628, 2022.
- [109] Faranaksadat Solat, Tae Yeon Kim, and Joohyung Lee. A novel group management scheme of clustered federated learning for mobile traffic prediction in mobile edge computing systems. *Journal of Communications and Networks*, 25(4):480–490, 2023.
- [110] Caiyu Su, Jinri Wei, Yuan Lei, Hongkun Xuan, and Jiahui Li. Empowering precise advertising with fed-gancc: A novel federated learning approach leveraging generative adversarial networks and group clustering. *Plos one*, 19(4): e0298261, 2024.
- [111] Wen Sun, Shiyu Lei, Lu Wang, Zhiqiang Liu, and Yan Zhang. Adaptive federated learning and digital twin for industrial internet of things. *IEEE Transactions on Industrial Informatics*, 17(8):5605–5614, 2020.
- [112] Afaf Taik, Zoubair Mlika, and Soumaya Cherkaoui. Clustered vehicular federated learning: Process and optimization. *IEEE Transactions on Intelligent Transportation Systems*, 23(12):25371–25383, 2022.
- [113] Alysa Ziyang Tan, Han Yu, Lizhen Cui, and Qiang Yang. Towards personalized federated learning. *IEEE transactions on neural networks and learning systems*, 34(12):9587–9603, 2022.
- [114] Ye Lin Tun, Minh NH Nguyen, Chu Myaet Thwal, Jinwoo Choi, and Choong Seon Hong. Contrastive encoder pre-training-based clustered federated learning for heterogeneous data. *Neural Networks*, 165:689–704, 2023.

- [115] Saeed Vahidian, Mahdi Morafah, Weijia Wang, Vyacheslav Kungurtsev, Chen Chen, Mubarak Shah, and Bill Lin. Efficient distribution similarity identification in clustered federated learning via principal angles between client data subspaces. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 10043–10052, 2023.
- [116] Hongyi Wang, Mikhail Yurochkin, Yuekai Sun, Dimitris Papailiopoulos, and Yasaman Khazaeni. Federated learning with matched averaging. *arXiv preprint arXiv:2002.06440*, 2020.
- [117] Siyang Wang, Haitao Zhao, Wanli Wen, Wenchao Xia, Bin Wang, and Hongbo Zhu. Contract theory based incentive mechanism for clustered vehicular federated learning. *IEEE Transactions on Intelligent Transportation Systems*, 25(7): 8134–8147, 2024.
- [118] Tianjing Wang and ZhaoYang Dong. Blockchain-based clustered federated learning for non-intrusive load monitoring. *IEEE Transactions on Smart Grid*, 15(2):2348–2361, 2023.
- [119] Tianjing Wang, Chao Ren, Zhao Yang Dong, and Christine Yip. Domain-adaptive clustered federated transfer learning for ev charging demand forecasting. *IEEE Transactions on Power Systems*, 40(2):1241–1254, 2024.
- [120] Yuntao Wang, Zhou Su, Yanghe Pan, Tom H Luan, Ruidong Li, and Shui Yu. Social-aware clustered federated learning with customized privacy preservation. *IEEE/ACM Transactions on Networking*, 2024.
- [121] Małgorzata Wasilewska, Hanna Bogucka, and Adrian Kliks. Federated learning for 5g radio spectrum sensing. *Sensors*, 22(1):198, 2021.
- [122] Mohamad Wazzeih, Mohamad Arafeh, Hani Sami, Hakima Ould-Slimane, Chamseddine Talhi, Azzam Mourad, and Hadi Otrok. Crsfl: Cluster-based resource-aware split federated learning for continuous authentication. *Journal of Network and Computer Applications*, 231:103987, 2024.
- [123] Xiao-Xiang Wei and Hua Huang. Edge devices clustering for federated visual classification: A feature norm based framework. *IEEE Transactions on Image Processing*, 32:995–1010, 2023.
- [124] Zongpu Wei, Jinsong Wang, Zening Zhao, and Kai Shi. Toward data efficient anomaly detection in heterogeneous edge–cloud environments using clustered federated learning. *Future Generation Computer Systems*, 164:107559, 2025.
- [125] Lihang Wen, Weijing Qi, Yufeng Lin, Qingyang Song, Lei Guo, and Abbas Jamalipour. A dynamic clustering-based hierarchical federated learning scheme in internet of vehicles. *IEEE Communications Letters*, 2024.
- [126] Honghu Wu, Xiangrong Zhu, and Wei Hu. A blockchain system for clustered federated learning with peer-to-peer knowledge transfer. *Proceedings of the VLDB Endowment*, 17(5):966–979, 2024.
- [127] Jiang Wu, Yunchao Yang, Miao Hu, Yipeng Zhou, and Di Wu. Fcer: A federated cloud-edge recommendation framework with cluster-based edge selection. *IEEE Transactions on Mobile Computing*, 2024.
- [128] Peng Xiao, Chuang Wang, Ze Lin, Ying Hao, Gang Chen, and Longhan Xie. Knowledge-based clustering federated learning for fault diagnosis in robotic assembly. *Knowledge-Based Systems*, 294:111792, 2024.
- [129] Han Xie, Jing Ma, Li Xiong, and Carl Yang. Federated graph classification over non-iid graphs. *Advances in neural information processing systems*, 34:18839–18852, 2021.
- [130] Bo Xu, Wenchao Xia, Haitao Zhao, Yongxu Zhu, Xinghua Sun, and Tony QS Quek. Clustered federated learning in internet of things: Convergence analysis and resource optimization. *IEEE Internet of Things Journal*, 11(2):3217–3232, 2023.
- [131] Xiaofeng Xue, Haokun Mao, Qiong Li, and Xin Guan. Clustered federated learning enhanced by dag-based blockchain with adaptive tip selection algorithm. *Internet of Things*, 31:101573, 2025.
- [132] Yihan Yan, Xiaojun Tong, and Shen Wang. Clustered federated learning in heterogeneous environment. *IEEE Transactions on Neural Networks and Learning Systems*, 35(9):12796–12809, 2023.
- [133] Guang Yang, Zhijia Yang, Shuping Cui, Chunhe Song, Jizhou Wang, and Haodong Wei. Clustering federated learning for wafer defects classification on statistical heterogeneous data. *IEEE Transactions on Instrumentation and Measurement*, 2024.
- [134] Xinyi Yang and Juan Li. A clustering-based federated deep learning approach for enhancing diabetes management with privacy-preserving edge artificial intelligence. *Healthcare Analytics*, 7:100392, 2025.
- [135] Zhikai Yang, Yaping Liu, Shuo Zhang, and Keshen Zhou. Personalized federated learning with model interpolation among client clusters and its application in smart home. *World Wide Web*, 26(4):2175–2200, 2023.
- [136] Lin Yao, Jia Li, Jing Deng, and Guowei Wu. Detection of cache pollution attack based on federated learning in ultra-dense network. *Computers & Security*, 124:102965, 2023.
- [137] Mang Ye, Xiuwen Fang, Bo Du, Pong C Yuen, and Dacheng Tao. Heterogeneous federated learning: State-of-the-art and research challenges. *ACM Computing Surveys*, 56(3):1–44, 2023.
- [138] Wenhui Ye, Ke Zheng, Yuanyu Wang, and Yuliang Tang. Federated double deep q-learning-based computation offloading in mobility-aware vehicle clusters. *IEEE Access*, 11:114475–114488, 2023.
- [139] Ziming Ye, Xiao Zhang, Xu Chen, Hui Xiong, and Dongxiao Yu. Adaptive clustering based personalized federated learning framework for next poi recommendation with location noise. *IEEE Transactions on Knowledge and Data Engineering*, 36(5):1843–1856, 2023.

- [140] Eungeun Yoo, Haneul Ko, and Sangheon Pack. Fuzzy clustered federated learning algorithm for solar power generation forecasting. *IEEE Transactions on Emerging Topics in Computing*, 10(4):2092–2098, 2022.
- [141] Sanshi Lei Yu, Qi Liu, Fei Wang, Yang Yu, and Enhong Chen. Federated news recommendation with fine-grained interpolation and dynamic clustering. In *Proceedings of the 32nd ACM international conference on information and knowledge management*, pages 3073–3082, 2023.
- [142] Xinlei Yu, Zhipeng Gao, Zijian Xiong, Chen Zhao, and Yang Yang. Ddpq-adaptconfig: A deep reinforcement learning framework for adaptive device selection and training configuration in heterogeneity federated learning. *Future Generation Computer Systems*, 163:107528, 2025.
- [143] Xue Yu, Ziyi Liu, Wu Wang, and Yifan Sun. Clustered federated learning based on nonconvex pairwise fusion. *Information Sciences*, 678:120956, 2024.
- [144] Dun Zeng, Xiangjing Hu, Shiyu Liu, Yue Yu, Qifan Wang, and Zenglin Xu. Stocfl: A stochastically clustered federated learning framework for non-iid data with dynamic client participation. *Neural Networks*, 187:107278, 2025.
- [145] Chenhan Zhang, Lei Cui, Shui Yu, and James JQ Yu. A communication-efficient federated learning scheme for iot-based traffic forecasting. *IEEE Internet of Things Journal*, 9(14):11918–11931, 2021.
- [146] Mao Zhang, Tie Zhang, Yifei Cheng, Changcun Bao, Haoyu Cao, Deqiang Jiang, and Linli Xu. Communication-efficient clustered federated learning via model distance. *Machine learning*, 113(6):3869–3888, 2024.
- [147] Wenjun Zhang, Xiaoli Liu, and Sasu Tarkoma. Fedgk: Communication-efficient federated learning through group-guided knowledge distillation. *ACM Transactions on Internet Technology*, 24(4):1–21, 2024.
- [148] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018.
- [149] Zhongyuan Zhao, Jingyi Wang, Wei Hong, Tony QS Quek, Zhiguo Ding, and Mugen Peng. Ensemble federated learning with non-iid data in wireless networks. *IEEE Transactions on Wireless Communications*, 23(4):3557–3571, 2023.
- [150] Rui Zhou, Yanting Li, and Xinhua Lin. A clustered federated learning framework for collaborative fault diagnosis of wind turbines. *Applied Energy*, 377:124532, 2025.
- [151] Xiaokang Zhou, Xiaozhou Ye, I Kevin, Kai Wang, Wei Liang, Nirmal Kumar C Nair, Shohei Shimizu, Zheng Yan, and Qun Jin. Hierarchical federated learning with social context clustering-based participant selection for internet of medical things applications. *IEEE Transactions on Computational Social Systems*, 10(4):1742–1751, 2023.

A Supplementary Material: High-level summary and classification of applicative CFL papers by domains

In the following tables, we detail the clustering information and methodology used in each applied CFL paper. The 'CFL classification' column classifies each approach according to our taxonomy (Section 4.2). For server-side and client-side approaches, the 'Clustering Strategy/Metadata' column briefly indicates the mechanism used and whether the method follows a known CFL baseline from Section 4.3. For metadata-based CFL, this column specifies the metadata shared with the server, such as geographical coordinates, resource availability or feature statistics. Finally, the 'Datasets & Task' column lists the datasets, data types, and tasks evaluated.

Table 7. Applicative CFL papers in the IoT/Networking domain.

CFL classification	Clustering Strategy/metadata	Datasets & Task	Reference
Server-side	MTCFL-like [103]	Gotham IoT Testbed – Anomaly Detection (Time Series)	Saez et al. 2023 [100]
Server-side	MTCFL-like [103]	CIFAR-10 – Image Classification	Hamood et al. 2023 [41]
Server-side	FlexCFL-like [24]	FEMNIST – Image Classification	He et al. 2023 [45]
Server-side	Coalition Formation Game with model cosine similarity	MNIST, FEMNIST, CIFAR-10, California Housing – Classification/Regression	Zhao et al. 2023 [149]
Server-side	Affinity Propagation + Cosine Similarity	Sent140, Fashion-MNIST, FEMNIST, CIFAR-10 – Text/Image Classification	Fotohi et al. 2024 [28]
Server-side	MTCFL-like [103]	FEMNIST, CIFAR-10 – Image Classification	Hamood et al. 2024 [42]
Server-side	K-means on weights	MNIST, CIFAR-10, EMNIST – Image Classification	Guo et al. 2024 [37]
Server-side	K-means on weights	IoT-23, IoTAnalytics – Anomaly Detection (Time Series)	Fan et al. 2024 [25]
Server-side	K-means on weights	MNIST, FMNIST, CIFAR-10 – Image Classification	Yu et al. 2025 [142]
Server-side	DBSCAN on weights	FEMNIST, CIFAR-10, Shakespeare – Image/Text Classification	Chien et al. 2025 [18]
Client-side	IFCA-like [33]	MNIST, FMNIST, FEMNIST – Image Classification	He et al. 2023 [47]
Client-side	IFCA-like [33]	MNIST, CIFAR-10 – Image Classification	Sami et al. 2023 [101]
Metadata-based	Computing power	MNIST – Image Classification	Sun et al. 2020 [111]
Metadata-based	Spectrum Sensing Metrics	Synthetic – Image Classification	Wasilewska et al. 2021[121]
Metadata-based	Load/Geographical data	Simulation Data – Attack Classification	Yao et al. 2023 [136]
Metadata-based	Quality of Service / popularity	MovieLens – Recommendation	Huang et al. 2023 [50]
Metadata-based	Utility Balance	MNIST, CIFAR-10, CIFAR-100 – Image Classification	Xu et al. 2023 [130]
Metadata-based	Edge-association	MNIST – Image Classification	Lee et al. 2023 [59]
Metadata-based	Temporal features	Gas Pipeline ICS, UNSW-NB15 – Intrusion Detection	Shan et al. 2023 [106]
Metadata-based	Soft prediction vector	Fashion-MNIST, CIFAR-10, SVHN – Image Classification	Li et al. 2024 [61]
Metadata-based	Local model performance and labels balance	UNSW-NB15, CIC-IDS2017/18/19, BCCC-DDoS2024 – Anomaly Detection	Wei et al. 2025 [124]
Metadata-based	Route, movement type	Ray-tracing / MATLAB Comm. Toolbox – Channel frame prediction / beam tracking	Ali et al. 2025 [4]
Metadata-based	Environment/service type	IoT Simulation – Reinforcement Learning	Baek et al. 2025 [7]

Table 8. Applicative CFL papers in the Mobility domain.

CFL classification	Clustering Strategy/metadata	Datasets & Task	Reference
Server-side	Kmeans-like on weights	HARBox; UWB; IMU; Depth gesture – HAR	Ouyang et al. 2021 [88]
Server-side	FL+HC-like [10]	PeMS, METR-LA – Traffic Forecasting	Zhang et al. 2021 [145]
Server-side	MTCFL-like [103]	HAR datasets – HAR	Ouyang et al. 2022 [89]
Server-side	Adaptive soft clustering using Gumbel-Softmax gating	GeoLife, Gowalla – Next POI Recommendation	Ye et al. 2023 [139]
Server-side	Kmeans on weights	MNIST, CIFAR-10, Belgium TSC – Classification	Wang et al. 2024 [117]
Server-side	FL+HC-like [10]	Fashion-MNIST; CIFAR-10; SVHN – Classification	Lu et al. 2025 [75]
Metadata-based	Geographic location	PeMS – Traffic Forecasting (Time Series)	Liu et al. 2020 [70]
Metadata-based	Transmit power, link lifetime, channel state informations, velocity, dataset size and diversity (using gradients)	MNIST, FMNIST – Image Classification	Taik et al. 2022 [112]
Metadata-based	Mobility factor	Simulation data – Reinforcement learning	Ye et al. 2024 [138]
Metadata-based	POI feature vectors	Chengdu – POI-based Forecasting	Liao et al. 2024 [65]
Metadata-based	DBSCAN on mean embedding vectors	Outdoor robot navigation – Traversability Prediction	Gummadi et al. 2024 [36]
Metadata-based	Predictive-uncertainty scores	CWRU, Paderborn PU, ISU – Fault Diagnosis (Time Series)	Lu et al. 2024 [73]
Metadata-based	Data GMM	FMNIST, CIFAR-10, CelebA, BC, NER, ADE, SSM4H, Wisconsin, Yeast1, Vehicle, Shuttle, Car, KR/Classification	Pei et al. 2024 [90]
Metadata-based	Local centroid	CIFAR-10; HAR-UCI; Ubisound/Classification	Li et al. 2024 [62]
Metadata-based	Contrastive embeddings	CityPersons, BDD100K, CCD – Pedestrian / anomaly detection	Belhadi et al. 2025 [8]

Table 9. Applicative CFL papers in the Energy domain.

CFL classification	Clustering Strategy/metadata	Datasets & Task	Reference
Server-side	FL+HC-like [10]	Low Carbon London – Forecasting (Time series)	Briggs et al. 2022 [11]
Server-side	Kmeans on weights	Chinese gas companies – Forecasting (Time series)	Qin et al. 2023 [93]
Server-side	Kmeans on weights	Windfarm data – Fault diagnosis (Time series)	Zhou et al. 2025 [150]
Client-side	IFCA-like [33]	REFIT, REDD smart-meter – Forecasting (Time series)	Wang et al. 2023[118]
Metadata-based	Customers’ behavioral and socio-economic metadata	London Smart Meter Energy Data – Forecasting (Time series)	Savi et al. 2021 [104]
Metadata-based	Geographical metadata	Solar + weather data – Forecasting (Time series)	Yoo et al. 2022 [140]
Metadata-based	Data centroid and avg load	Australian SGSC smart-meter – Forecasting (Time series)	He et al. 2023 [46]
Metadata-based	local SARIMA-parameters	China Carbon Accounting Database – Forecasting (Time series)	Cui et al. 2023 [19]
Metadata-based	Data-derived profiles	REFIT - Regression (Time series)	Shi et al. 2023 [107]
Metadata-based	Signal/metadata extraction from TS	BC Hydro residential smart-meter – Forecasting (Time series)	Dogra et al. 2023 [21]
Metadata-based	Resource + data distribution statistics	FMNIST, CIFAR-10 – Image Classification	Yang et al. 2023 [135]
Metadata-based	DTW-based centroids	BDG2 – Forecasting (Time series)	Liu et al. 2025 [68]
Metadata-based	SHAP feature-contribution vectors	German Solar Farm – Forecasting (Time series)	Ali et al. 2025 [5]
Metadata-based	Dataset Encoder output	Palo Alto EVCS; Boulder EVCS – Forecasting (Spatio-temporal Data)	Wang et al. 2024 [119]

Table 10. Applicative CFL papers in Health.

CFL classification	Clustering Strategy/metadata	Datasets & Task	Reference
Server-side	Stoer–Wagner global min-cut (graph built with pairwise cosine similarities of client models)	MUTAG; BZR; COX2; DHFR; PTC_MR; AIDS; NCI1; ENZYMES; DD; PROTEINS; COLLAB; IMDB-BINARY; IMDB-MULTI – Graph Classification	Xie et al. 2021 [129]
Server-side	MTCFL-like [103]	ESOL; QM9 – Generative molecular design (Graphs)	Guo et al. 2024 [38]
Server-side	K-means on weights	FeTS2022 (BraTS2021) – 3D brain tumor segmentation (3D image)	Manthe et al. 2024 [79]
Server-side	K-means/GMM on weights - HC+FL-like [10]	PPG-DaLiA; UCI Drug Review – Medical HAR classification; Drug rating classification	Shaik et al. 2024 [105]
Server-side	MTCFL-like [103]	BraTS2018; Retinal vessel datasets (FIVES, STARE, DRIVE, IOSTAR, CHASE_DB1, RAVIR) – Image segmentation	Niu et al. [87]
Server-side	FlexCFL-like [24] with affinity propagation	MNIST; USPS; BloodMNIST; OrganAMNIST; OrganCMNIST; OrganSMNIST – Image classification	Cao et al. 2025 [13]
Metadata-based	Based on modality	COVID-19 image data collection; Ultrasound COVID dataset (POCUS) – Image Classification	Quayyum et al. 2022 [92]
Metadata-based	Bandwidth + dataset size + freshness of data	CASAS smart home sensor datasets – Classification (HAR)	Singh et al. 2022 [108]
Metadata-based	Features distribution	MIT Drivedb; AffectiveRoad – Binary driver-stress classification	Rafi et al. 2024 [96]
Metadata-based	Clients token-frequency embeddings	ADReSS (DementiaBank) – Automatic Speech Recognition (Audio)	Hsu et al. 2024 [48]
Metadata-based	Dataset extracted Embeddings	WESAD wearable stress dataset – Classification (Time Series)	Jiang et al. 2023 [54]
Metadata-based	Daily carbohydrate (CHO) intake	UVA/Padova T1D Simulator – Forecasting (Time Series)	Yang et al. 2025 [134]
Metadata-based	PCA vectors of clients features	Damage Simulation Data – Classification	Cheema et al. 2025 [15]
Metadata-based	Dataset prototype vectors	MedMNIST v2 (BloodMNIST, DermaMNIST) – Image Classification	Gao et al. 2025 [29]
Metadata-based	Clients Soft predictions	SEED; SEED-IV; SEED-FRA; SEED-GER – Emotion classification	Qiu et al. 2025 [95]

Table 11. Applicative CFL papers in Advertising, Recommendation, Speech, and Robotics domains.

Approach	Clustering Metric / Strategy	Datasets & Task	Reference
Server-side	MTCFL-like [103]	MNIST; CIFAR-10; iNaturalist-Geo; FEMNIST; Structured Non-IID FEMNIST; VisDA2017; iWildCam – Image classification	Wei et al. 2023 [123]
Server-side	K-means on weights	CPHPD (Persian dialect speech) – Speech recognition classification (Audio)	Farahni et al. 2023 [26]
Server-side	Kmeans-like on weights	Ali Tianchi Display Ad Click; Kaggle CTR – Click rate prediction	Su et al. 2024 [110]
Server-side	DBSCAN on weights	LibriSpeech; UserLibri – Speech recognition	Lee et al. 2024 [58]
Metadata-based	Clients Autoencoder outputs	Speech recognition	Koppelman et al. 2022 [57]
Metadata-based	Clients local click history (news)	Adressa; MIND – News click prediction / recommendation ranking	Yu et al. 2023 [141]
Metadata-based	Clients dataset statistics + training-time	CIFAR-10; CIFAR-100; Shakespeare; Cora; CiteSeer – Image classification; character prediction; graph node classification	Wu et al. 2024 [126]
Metadata-based	Validation accuracy score vectors	Proprietary robotic-assembly dataset – Fault diagnosis	Xiao et al. 2024 [128]