Accelerating Federated Learning via Sequential Training of Grouped Heterogeneous Clients

IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received December 14, 2023. Revised and re-submitted March 11, 2024.

Digital Object Identifier 10.xxxx/ACCESS.2023.40950

Accelerating Federated Learning via Sequential Training of Grouped Heterogeneous Clients

ANDREA SILVI*, ANDREA RIZZARDI*, DEBORA CALDAROLA*, BARBARA CAPUTO, and MARCO CICCONE

All the authors are with Dipartimento di Automatica e Informatica (DAUIN), Politecnico di Torino, 10129 Turin, Italy.
*Equal contribution.

Corresponding author: Debora Caldarola (e-mail: debora.caldarola@polito.it).

This study was carried out within the FAIR - Future Artificial Intelligence Research and received funding from the European Union Next-GenerationEU (PIANO NAZIONALE DI RIPRESA E RESILIENZA (PNRR) – MISSIONE 4 COMPONENTE 2, INVESTIMENTO 1.3 – D.D. 1555 11/10/2022, PE00000013). This manuscript reflects only the authors' views and opinions, neither the European Union nor the European Commission can be considered responsible for them.

ABSTRACT Federated Learning (FL) allows training machine learning models in privacy-constrained scenarios by enabling the cooperation of edge devices without requiring local data sharing. This approach raises several challenges due to the different statistical distribution of the local datasets and the clients' computational heterogeneity. In particular, the presence of highly non-i.i.d. data severely impairs both the performance of the trained neural network and its convergence rate, increasing the number of communication rounds required to reach centralized performance. As a solution, we propose *FedSeq*, a novel framework leveraging the sequential training of subgroups of heterogeneous clients, *i.e.*, *superclients*, to learn more robust models before the server-side averaging step. Given a fixed budget of communication rounds, we show that FedSeq outperforms or match several state-of-the-art federated algorithms in terms of final performance and speed of convergence. Our method can be easily integrated with other approaches available in the literature, and empirical results show that combining existing algorithms with FedSeq further improves its final performance and convergence speed. We evaluate our method across multiple FL benchmarks, establishing its effectiveness in both i.i.d. and non-i.i.d. scenarios. Lastly, we highlight that the sequential training introduced here does not introduce additional privacy concerns when compared to the de facto standard, FedAvg.

INDEX TERMS Federated learning, distributed learning, privacy-preserving machine learning, statistical heterogeneity, deep learning.

I. INTRODUCTION

Federated Learning (FL) [1] is a Machine Learning (ML) method designed to train models in distributed systems while preserving the privacy of participants (the *clients*). Such a paradigm eliminates the need for clients to disclose their private data with a central authority, thus ensuring compliance with regulations in force. Federated training involves multiple communication rounds, during which a shared global model is trained independently on selected devices. The updated parameters are then aggregated by the server into a new model. While usually effective in homogeneous scenarios [1]–[3], where clients have access to similar data, in realistic settings they observe data by breaking the conventional assumption of independence and identical distribution (*i.i.d.*) of classic ML systems. In these scenarios, users collect data from an underlined global distribution based on preferences [4], [5] or geographic position [6]–[8], forming a heterogeneous distribution of local datasets. Several works have shown that heterogeneity generally results in slow and unstable convergence of FL algorithms [3], [9], [10], hampering final performance [11], [12] because of local gradients diverging towards different minima [1], [3], and the difficulty to merge specialized models, a phenomenon called *client drift* [13]. This results in a biased and suboptimal global solution compared to the actual minimum [14]. In this work, we focus on heterogeneity caused by i) *label skew*, *i.e.* given an instance-label pair $(x, y) \sim P_k(x, y)$, $P_k(y)$ varies across clients k while $P(y|x)$ is identical, ii) *features shift*, *i.e.* $P_k(x)$ varies while $P(y|x)$ is identical, and iii) *different local dataset cardinality*.

To mitigate the effects of the client drift, many methods focus on regularizing the local objective to bring it closer to the global one [10], [13], [15], or on improving the gener-

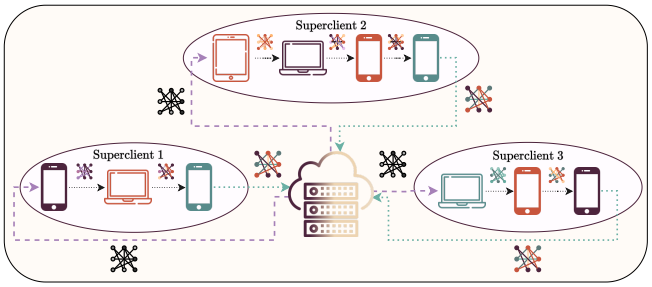

FIGURE 1. To mitigate statistical heterogeneity in FL, FedSeq forms *superclients* by grouping clients with distinct local data distributions (different colors), creating simulated larger and homogeneous datasets. Sequential training takes place within the selected superclients at each round. The current global model is received by the first client in the chain and sent back by the last one.

alization of the learned model [3], [16], [17]. Multitask FL views each local distribution as an individual *task* and aims to train separate but related models concurrently [18]–[20], while [21]–[26] cluster clients with similar tasks together and assign a specific model each. Data sharing approaches instead leverage fine-tuning of the model over small quantities of public or synthesized i.i.d. data to ensure a more balanced representation of the overall distribution on the server side [11]. However, most of these methods fall short of replicating the performance of centralized scenarios or struggle with extremely skewed heterogeneous distributions [3]. In contrast, **this paper approaches heterogeneity from a different angle, focusing on the training orchestration rather than the training objectives**, to learn more robust models before the server-side averaging step, resulting in reduced noise and achieving centralized performance.

This work presents *Federated Learning via Sequential Superclients Training* (**FedSeq**), a novel approach effectively addressing the issue of statistical heterogeneity in FL. FedSeq employs *sequential* training among clients, carried out in parallel across distinct client groups to harness the distributed setting's parallelism. By allowing the model to access a larger portion of data before the averaging step, the negative effects of data heterogeneity are mitigated, speeding up the training and moving closer to the desired minimum. By grouping clients having diverse local distributions together in a *superclient*, we simulate the existence of a larger, homogeneous dataset while maintaining data privacy, as illustrated in Fig. 1. Clients within the same superclient form a chain and train the received model in a sequential manner. The final updates are sent from the last client to the server and merged there. Intuitively, **this scheme emulates the training dynamics observed on devices with more extensive and evenly distributed datasets**, resulting in a favorable setting for FL.

Communication is known to be the main bottleneck in federated training, *e.g.* due to the clients' unavailability and unreliability [27]. While sequential training provides robustness against data heterogeneity, it can potentially result in slower training progress. This occurs when slower clients end up in the same superclient, leading to increased waiting times on the server side. To overcome this limitation, we present **FedAsyncSeq**, a novel approach that introduces **asynchronous client-server communication by implementing sequential training among superclients**. Rather than merging updates at the end of each training round, FedAsyncSeq allows the model updated by one superclient to be sent directly to another one, enabling faster groups of clients to complete multiple training iterations before merging their updates. At regular intervals of every R rounds, the updates received by the server, potentially stemming from varying numbers of training iterations, are aggregated. This approach not only reduces the number of aggregation and synchronization steps with the server but also allows the model to be trained on a larger number of clients before the averaging step. Consequently, this brings the model closer to a centralized scenario, as ideally, it encounters all superclients before being merged.

Extension Details

This work represents an extension of the previous manuscript [28] in several aspects, as summarized in Fig. 2. ① Firstly, we introduce a **novel metric for estimating local distribution similarity without compromising user privacy or needing additional public data**, outperforming the performance of the previously presented approximation techniques. Based on *Task2Vec* [29], it captures both taxonomic and semantic representations of each task, *i.e.*, the local data of each client. In addition, we note that sequential training can suffer from saturation in the later stages of training due to overfitting [30], which can further slow down the overall training process. ② Building upon [30], **FedSeq2Par** is presented as a solution **to increase parallelism as training moves on**. FedSeq2Par dynamically updates the number of superclients and their corresponding client assignments as the rounds progress. It prioritizes sequentiality, *i.e.* larger groups, during the initial stages, and gradually transitions to parallelism, emphasizing smaller groups in the later stages. This allows easier adaptation to varying numbers of devices and distributions.

Empirical analyses demonstrate the superior performance and convergence speed of the introduced methods compared to the current state-of-the-art algorithms in federated scenarios with various data distributions and tasks. ③ To provide a comprehensive evaluation, **the experimental benchmark is extended to include vision datasets that exhibit feature shifts in addition to label skew, as well as Natural Language Processing (NLP) datasets**.

④ Lastly, the robustness of the algorithm against threats posed by potentially malicious participants is a crucial aspect in the FL paradigm [31]–[33]. Malevolent clients may attempt to disrupt the training process by manipulating their input data [34], [35], or even try to infer private information of other clients by exploiting the received global model [36], [37]. To assess whether our novel *client-to-client* sequential training approach introduces any privacy vulnerabilities, in this extension we conduct tests against these attacks and ob-

serve that **FedSeq often exhibits higher privacy resistance compared to the widely-used FedAvg** [1], considered the de-facto standard algorithm for Federated Learning.

Contributions

To summarize, our main contributions are the following:

- We introduce FedSeq, a new FL algorithm that learns from groups of sequentially-trained clients (*superclients*).
- Several lightweight procedures to compare the clients' probability distributions, analyzing their impact on the creation of superclients. Extending the previous [28], this work proposes to estimate the distribution of the clients' tasks via *Task2Vec*, eliminating the need for external public data.
- Three grouping strategies are evaluated and compared with the naïve random assignment, showing the impact of group quality on the algorithm convergence.
- To speed up training, we introduce FedAsyncSeq to decrease the need for synchronization between superclients and server, and FedSeq2Par to increase parallelism.
- The extensive empirical analyses and tests demonstrate that the developed approaches outperform the state of the art in terms of convergence performance and speed in both i.i.d. and non-i.i.d. scenarios. This paper further extends the benchmark's scope by incorporating both vision and language datasets.
- We prove the resistance of FedSeq to common attacks against clients' privacy, showing its robustness.

II. RELATED WORKS

A. DATA HETEROGENEITY IN FL

Federated Learning (FL) [1] is a framework to learn a global model distributedly while preserving the users' data privacy [2], [27], [38]. Training [1] is based on rounds of communication between clients and a server, and the global model is built as the weighted average of the updated parameters (FedAvg) obtained via client-side training on local data. In realistic scenarios, a major challenge is posed by the presence of non-i.i.d. and unbalanced clients' data, often referred to as *statistical heterogeneity* [39], [40]. In those settings, the local optimization objectives drift from each other [13], leading to unstable trends [3] and slower convergence rates [6], [10], [39], [41], [42]. This work preserves FedAvg for the server-side updates aggregation while focusing on enhancing convergence performance.

Client-side approaches. Several works addressed data heterogeneity issues via client-side training regularization. SCAFFOLD [13] mitigates the effects of the clients' drifts through control variates, while [10], [43] minimize the gap between local and global model parameters to limit the impact of the clients' updates. FedAlign [44] aligns the Lipschitz constants of the models' last blocks to promote smooth optimization and global consistency. Other approaches leverage the *Alternating Direction Method of Multipliers* to asymptotically align local and global objectives [45]–[47]. Among those, FedDyn [15] dynamically modifies the local loss function so that local and global stationary points coincide at convergence. Another recent direction studies generalization through the lens of the loss landscape. By seeking flat minima during local training, FedSAM [3], [16], FedSpeed [17], and FedSMOO [48] improve the poor generalization of models in heterogeneous scenarios. Since FedSeq does not modify the client-side training, it can be applied alongside any of these approaches.

Server-side approaches. While client-side regularization is effective for mitigating client drift and enhancing local model learning, server-side aggregation is not to be overlooked. Using server-side optimizers [6], [40], [49], [50] allows coping with FedAvg's lack of adaptivity. Server-side momentum [40], [51], [52] addresses the bias towards recently observed clients by leveraging the memory of past updates. Other studies utilize global momentum to guide local updates [53]–[56]. The proposed revised orchestration of the training process can effortlessly integrate with any of these techniques.

As the learned local model under-represents the deducible patterns from the missing classes, [11] shows how sharing a small set of public data among the clients leads to notable improvements. A similar approach is followed by [57], where the public data enables knowledge distillation. [58] leverages public unlabelled data to learn a general representation robust to domain shifts. Similarly, FedSeq keeps the public data on the server-side, with the different purpose of using it to estimate the clients' data distribution in a privacy-compliant way. Unlike [11], [57], [58], such data is never used for training.

Federated Multi-Task Learning. Another line of works tackles the problem from a multitask perspective [59], where each client with its own data distribution is seen as a different task [18], [60]. In [21]–[25], clients with similar tasks are clustered together and a specific model is assigned to each cluster. [8] leverages the local style information to identify and group clients having similar distributions. Other works build clusters based on the edge systems complexity and available resources [26], [61]. Following the same approach of [21], [23], [24], FedSeq approximates the clients' data distribution via the locally trained models, which is later used to build groups of dissimilar clients. [19] uses the relatedness among clients' tasks to improve weight aggregation. Here, Task2Vec [29] embeddings, based on the Fisher information matrix (FIM) of fine-tuned local models, are leveraged to capture similarities among the clients' tasks without needing external public datasets.

Anti-clustering FL. FedSeq uses clustering techniques to effectively group clients with distant distributions, resulting in a homogeneous underlying dataset within each group. This approach relates to the "*anti-clustering*" literature [62], [63], whose goal is to build similar groups from dissimilar elements [64]. Building upon the strategy of FedGSP [30],

which dynamically expands the number of client groups in each round to enhance parallelism, we propose merging this approach with the data estimation strategies and grouping techniques unique to FedSeq. This combination results in FedSeq2Par, an approach that further increases parallelism.

B. BEYOND SYNCHRONOUS SERVER-CLIENT FEDERATED LEARNING

Peer-to-peer FL. Peer-to-peer (p2p) FL [65], [66] is a decentralized approach where clients communicate directly with each other in order to learn global models, eliminating the need for a central server. In particular, FedSeq shares some common traits with [67], which introduces two network topologies based on cyclic model parameters exchange among clients to enhance performance in heterogeneous scenarios, namely FedCyclic and FedStar. Similarly, our approach enables model sharing among clients within the same superclient. Unlike such works, FedSeq retains the central server as a proxy between clients, while ensuring communication costs equivalent to those of FedAvg.

Asynchronous FL. Asynchronous FL was introduced to handle stragglers and heterogeneous latency [68], [69]. In this scenario, the server does not wait for all devices to send back their updates but keeps aggregating the models as they arrive. Several approaches rely on a parameter accounting for staleness [68], [70], leverage gradient compression techniques to reduce the communication latency [71], or store the early updates for a given timeframe and discard the late ones [72]. Similarly to these methods, the server in FedAsyncSeq does not wait for all superclients to return their updates but allows faster ones to continue training for multiple rounds before the averaging step. Differently from the standard asynchronous approach, updates are not averaged as they come, but after a fixed window of rounds so that the exchanged models are trained on a larger portion of data.

C. PRIVACY IN FEDERATED LEARNING

A primary objective within the FL framework is to ensure the algorithm's resilience against potential threats posed by malicious participants. However, it's essential to recognize that the FL paradigm, in its current form, is not entirely impervious to threats [73], [74]. Malevolent clients may *poison* the training process by altering the input data [34], [35], worsening the capacity of the global model to acquire new useful knowledge. An attacker might *reconstruct clients' private data* by exploiting the incoming update model [36]. As an example, [75] leverages a GAN (Generative Adversarial Network) [76] to reconstruct other users' personal data, while [77] uses a GAN to ensure the quality of the data that the attacker aims to reconstruct and [78] tries to infer characteristics of the clients with ad-hoc classifiers. Additionally, a malicious server might put in place *label and feature fishing* attacks by intentionally modifying some parameters of the global model [79]. For an in-depth discussion on threats and attacks in FL, we refer the readers to [32], [33]. Common defense techniques involve using differential privacy [80], [81], or mixing fragments of the local updates before sending them to the server [82]. This work explores some of those attacks against FedSeq and shows that the proposed novel *client-to-client* sequential training approach is robust in terms of privacy compared to FedAvg.

TABLE 1. Summary of main introduced symbols.

Notation	Description
f	Global model
k	Client index
θ	Global model parameters
θ_t	Global model parameters at round t
θ_t^k	k-th client updated model parameters
C	Set of clients
C_t	Set of clients selected at round t
K	Number of clients ($\lvert C\rvert$)
C	Fraction of clients selected at each round
$\mathcal{D}_k$	k-th client local dataset
$\tilde{\mathcal{D}}_k$	Estimate of k-th client local dataset
n_k	Number of local samples ($\lvert\mathcal{D}_k\rvert$)
T	Number of rounds
E_k	Local training epochs
i	Superclient index
$\mathcal{S}$	Set of superclients
$\mathcal{S}_t$	Set of superclients selected at round t
N_S	Number of superclients ($\lvert\mathcal{S}\rvert$)
S_i	i-th superclient
C_S	Set of clients belonging to superclient S
K_{S_i}	Number of clients in superclient S_i ($\lvert C_S\rvert$)
$K_{S,max}$	Maximum number of clients in each superclient
ψ	Clients distribution estimator
ϕ	Grouping method
τ	Distance metric
e	Pretraining number of epochs
E_S	Loops within each superclient

III. METHOD

This section details the problem formulation (Sec. III-A), and the components of the proposed method, distinguishing between FedSeq, FedAsyncSeq and FedSeq2Par (Secs. III-B and III-C). The overall procedure is outlined in Fig. 2, highlighting the extension's additions, while the used symbols are summarized in Table 1.

A. PROBLEM FORMULATION

The objective of FL is to learn a global model $f_\theta : \mathcal{X} \to \mathcal{Y}$, where $\mathcal{X}$ and $\mathcal{Y}$ are the input and output space respectively, and $\theta \in \mathbb{R}^d$ the model parameters. The server communicates with a subset of active clients sampled uniformly at random from a set of users C across T rounds. In cross-device settings [27], the number of clients $K := |C|$ is in the order of millions. Each client $k \in C$ has access to a local private dataset $\mathcal{D}_k$ of the form $\{x_i, y_i\}_{i=1}^{n_k}$ where $x_i \in \mathcal{X}$ is the input data point, $y_i \in \mathcal{Y}$ its label and $n_k = |\mathcal{D}_k|$ the local dataset cardinality. Due to communication constraints [2], only a fraction C_t of clients is randomly selected without replacement for training at round $t \in [T]$. Each client $k \in C_t$ receives the current global model parameters θ_t from the server and computes the

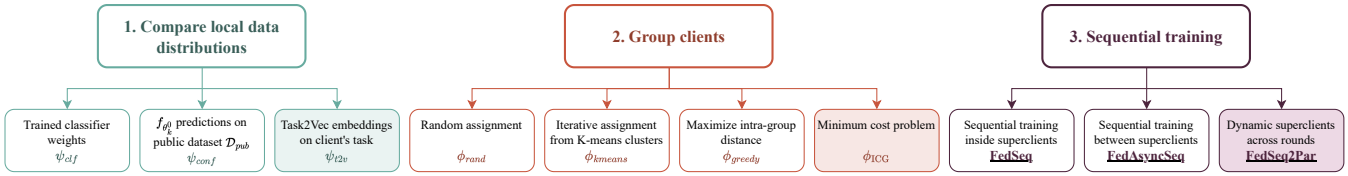

FIGURE 2. **Summary of the method components. Extensions with respect to [28] are highlighted with colored boxes. Best seen in colors.**

update θ_{t+1}^k by minimizing the local empirical risk $L_k(\theta) = \mathbb{E}_{(x,y)\sim\mathcal{D}_k}[\ell_k(f_{\theta_t};(x,y))]$ where ℓ_k is the loss function of the k-th client, *e.g.* the cross-entropy loss. The updates $\{\theta_{t+1}^k\}_{k\in C_t}$ are then sent to the server to be aggregated into the global model $f_{\theta_{t+1}}$. The global training objective is

$$\arg\min_{\theta\in\mathbb{R}^d} \sum_{k\in C_t} \frac{n_k}{n} L_k(\theta),\ d\in\mathbb{N}^+ \tag{1}$$

where $n = \sum_{k\in C_t} n_k$. The de-facto standard algorithm for solving the FL objective in Eq. 1 is FedAvg [1], which computes a weighted average of the clients' updates as $\theta_{t+1} \leftarrow \sum_{k\in C_t} {}^{n_k}/_{n}\,\theta_{t+1}^k$. As noted by [49], this is equivalent to performing one step of stochastic gradient descent (SGD) with unitary learning rate as $\theta_{t+1} \leftarrow \theta_t - \sum_{k\in C_t} {}^{n_k}/_{n}\,(\theta_t - \theta_{t+1}^k)$, where the difference between the local model and the round initialization acts as pseudo-gradient for the client's direction.

As shown in [83], in realistic scenarios, clients likely do not draw data from the same underlying distribution, namely $\mathcal{P}(D_i) \neq \mathcal{P}(D_j)\ \forall i \neq j \in C$, resulting in slower and unstable convergence [39]. More in general, $f_{\theta^k} \neq f_{\theta_.}\ \forall k \in C$ [23]. To address the challenges arising from data heterogeneity and speed up convergence, FedSeq introduces modifications to the training orchestration process. Specifically, clients with diverse data distributions are grouped into *superclients* $\{S_i\}_{i=1}^{N_S}$, aiming to minimize the divergence in distribution among superclients. Sequential training is then performed by clients within the same group. Intuitively, this approach allows local models to accumulate knowledge from the overall data distribution, even when client datasets exhibit significant heterogeneity.

B. BUILDING SUPERCLIENTS

This section details how to create a superclient S from users with diverse local distributions while respecting the privacy constraints, *i.e.* without accessing clients' data directly. Assigning clients to equally-sized groups while minimizing distribution distance is a challenging problem similar to the bin packing problem [84], and is NP-hard in nature. Thus, this work proposes using multiple greedy strategies to estimate the local distributions in a privacy-preserving way and solve the clients' clustering problem, being flexible towards dynamic and constantly evolving FL environments. In this Section, we introduce different grouping criteria G_S which are based on *i*) a *client distribution estimator* $\psi_{(.)}$, providing privacy-preserving statistics on the local data distribution, *ii*) a *metric* τ, for evaluating the distance between the estimated data distributions, and *iii*) a *grouping method* $\phi_{(.)}$, to assemble

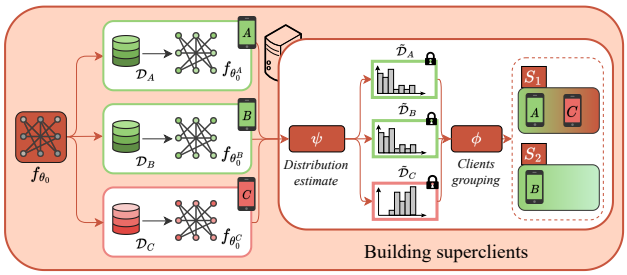

FIGURE 3. **FedSeq pre-training phase to build superclients. a) The initial random global model f_{θ_0} is sent to all the clients, which train it using their local data $\mathcal{D}_k\ \forall k \in C$. b) The local data distributions are estimated (ψ) using the clients' updates while preserving their privacy. c) Based on the grouping strategy ϕ, clients are assigned to N_S superclients. Best seen in colors.**

dissimilar clients, *i.e.* $G_S := \{\psi_{(.)};\tau;\phi_{(.)}\}$. The approach is depicted in Fig. 3.

1) Clients distribution comparison

The model f_θ can be defined as a combination of a deep feature extractor $h_{\theta_{\text{feat}}} : \mathcal{X} \to \mathcal{Z}$ and a classifier $g_{\theta_{\text{clf}}} : \mathcal{Z} \to \mathcal{Y}$, where $\theta = \{\theta_{\text{feat}}, \theta_{\text{clf}}\}$ is the entire set of model parameters and $\mathcal{Z}$ the output feature space. The classification output is given by $g \circ h : \mathcal{X} \to \mathcal{Y}$, where we drop the subscripts to ease the notation. FedSeq exploits a *pre-training phase* to estimate the users' data distribution. Each client $k \in C$ trains a common random model θ_0 on its dataset $\mathcal{D}_k$ for e epochs, resulting in $f_{\theta_0^k}$, which serves as a starting point for the following distribution estimation approaches. The proposed client distribution estimators are

- ψ_{clf}: as the **model classifier** is biased towards the training data [85], its parameters $\theta_{0,\text{clf}}^k$ serve as a proxy for the client's local data distribution.
- ψ_{conf} relies on the **predictions of the local models on a server-side public dataset**, *i.e.*, $\{f_{\theta_0^k}(z) =: f_0^k,\ z \in \mathcal{D}_{pub},\ k \in C\}$. $\mathcal{D}_{pub}$ contains J samples for each class $c \in [N_C]$. The predictions are averaged by class as $p_{k,c} = \frac{1}{J}\sum_{x\in\mathcal{D}_c} f_0^k(x)$, where $\mathcal{D}_c \subset \mathcal{D}_{pub}$ contains only samples of class c. The k-th client's *confidence vector* is defined as:

$$p_k := \text{softmax}(\{p_{k,1},\ldots,p_{k,N_C}\}) \in [0,1]^{N_C} \tag{2}$$

  Since the k-th model's predictions are favorable towards the majority of the classes seen in $\mathcal{D}_k$ [86], p_k is an acceptable privacy-preserving representation of $\mathcal{D}_k$.

Although ψ_{conf} has proven to be the most effective approach [28], it relies on the availability of a public dataset that accu-

rately reflects the overall data distribution — a requirement that can be challenging to meet. Therefore, extending [28], this study introduces an alternative method to overcome this limitation, without introducing any privacy liabilities.

- ψ_{t2v}: based on *Task2Vec* [29], which extracts **vectorial representations of given tasks** based on an approximation of the Fisher Information Matrix (FIM), defined as

$$F := \mathbb{E}_{(x,y)\sim f_\theta(x,y)}[\nabla_\theta \log f_\theta(y|x) \nabla_\theta \log f_\theta(y|x)^T]. \quad (3)$$

  The FIM serves as a metric for the information content of a parameter regarding the joint distribution $f_\theta(x, y)$. If it has limited influence on the classification performance for a specific task, its corresponding entry in the FIM will be low. Thus, the FIM represents the task itself, here corresponding to each client's local dataset. Starting from a pre-trained set of weights $\tilde{\theta}_0$, the classifier is fine-tuned on $\mathcal{D}_k$ and the FIM is computed on the feature extractor parameters. The resulting representations are demonstrated to capture taxonomic and semantic similarities between tasks.

In the following sections, we indicate as $\tilde{\mathcal{D}}_k$ the estimate provided by $\psi_{(.)}$ for the k-th device's data distribution.

2) Grouping clients

$\mathcal{D}_S = \bigcup_{k\in C_S} \mathcal{D}_k$ is defined as the union of the data from the clients $C_S \subset C$ belonging to a superclient S. The aim is to find the maximum amount of superclients N_S satisfying the following constraints: *i*) minimum number of samples $|\mathcal{D}_S|_{min}$, and *ii*) maximum number of clients $K_{S,max}$ per superclient. Given $\psi_{(.)}$ and τ, FedSeq [28] approximates the solution of the problem using:

- ϕ_{rand}, a naïve yet practical method that **randomly** assigns clients to superclients until the stopping criterion is met.
- ϕ_{kmeans}: **K-means** [87] is first applied to obtain N_S homogeneous clusters. Each superclient is formed by iteratively extracting one client at a time from each cluster, until $|\mathcal{D}_S| \geq |\mathcal{D}_S|_{min}$ and $K_S \leq K_{S,max}\ \forall S$.
- ϕ_{greedy} initially assigns one random client $k_i \in C$ to the superclient S. The next $k_j \in C \setminus \{k_i\}$ is chosen so as **the distance between k_i and k_j is maximized**, *i.e.* $\max_{j\in[K]} \tau(\tilde{\mathcal{D}}_{k_i}, \tilde{\mathcal{D}}_{k_j})$. The process is repeated by iteratively maximizing $\tau(\tilde{\mathcal{D}}_j, \frac{1}{|S|}\sum_{i\in|S|}\tilde{\mathcal{D}}_i)$, with $|S|$ being the cardinality of S, until the defined constraints are met.
- While highly effective [28], the best-performing ϕ_{greedy} is hindered by its iterative nature, leading to slower execution. In response, this work adapts the faster ***Inter-Cluster Grouping*** (ICG) algorithm from [30] to our approach (ϕ_{ICG}). Differently from ICG that requires superclients of equal size, FedSeq relaxes this constraint by redistributing the unassigned clients.

At the end of this procedure, we obtain a set $\mathcal{S}$ of N_S superclients, where each superclient S_i includes $K_{S_i} := |C_{S_i}| \leq K_{S,max}$ clients and $|\mathcal{D}_{S_i}| \geq |\mathcal{D}_S|_{min}$ data points, with $\sum_{S_i\in\mathcal{S}} K_{S_i} = K$.

Algorithm 1 FedSeq, FedAsyncSeq and FedSeq2Par

Require: $f_{\theta_0}, G_S, K_{S,max}, |\mathcal{D}_S|_{min}$. Epochs e, E_k, E_S. T rounds. Clients C. Fraction C of superclients selected at each round. Growth function and parameters f_{gr}, α_{gr} and β_{gr}.

$\mathcal{S} \leftarrow$ CREATE_SUPERCLIENTS$(f_{\theta_0}, G_S, e, K_{S,max}, |\mathcal{D}_S|_{min}, K, f_{gr}(\alpha_{gr}, \beta_{gr}; t))$
$N_S \leftarrow |\mathcal{S}|$
$\Theta \leftarrow [\theta_0, \ldots, \theta_0]_{1\times CN_S}$, $w \leftarrow [0, \ldots, 0]_{1\times CN_S}$
for $t = 1$ to T **do**
  if $f_{gr}(t, \alpha_{gr}, \beta_{gr}) > |N_S|$ **then**
    $\mathcal{S} \leftarrow$ CREATE_SUPERCLIENTS$(f_{\theta_0}, G_S, e, K, f_{gr}(\alpha_{gr}, \beta_{gr}; t))$
    $N_S \leftarrow |\mathcal{S}|$
  end if
  $\mathcal{S}_t \leftarrow$ Subsample fraction C of N_S superclients
  for $S_i \in \mathcal{S}_t$ **in parallel do**
    Shuffle clients in S_i
    $\theta_t^{S_i,0} \leftarrow \theta_t$ {FedSeq and FedSeq2Par}
    $\theta_t^{S_i,0} \leftarrow \Theta[i]$
    for $e_S = 1$ to E_S **do**
      $\theta_{t+1}^{S_i} \leftarrow$ SEQUENTIAL_TRAINING$(\theta_t^{S_i,0}, E_k)$
    end for
    $\Theta[i] \leftarrow \theta_{t+1}^{S_i}$, $w_i \leftarrow w_i + |\mathcal{D}_{S_i}|$
  end for
  $\theta_{t+1} \leftarrow$ FedAvg$(\{\theta_{t+1}^{S_i}, \forall S_i \in \mathcal{S}_t\})$ {FedSeq and FedSeq2Par}
  if $t \mod N_S = 0$ **then**
    $\theta_{t+1} \leftarrow \sum_i \frac{w_i}{w}\Theta[i]$, $w = \sum_i w_i$
    $\Theta \leftarrow [\theta_{t+1}, \ldots, \theta_{t+1}]_{1\times CN_S}$, $w \leftarrow [0, \ldots, 0]_{1\times CN_S}$
  end if
end for

C. SEQUENTIAL TRAINING

This section introduces three alternatives to leverage *sequential training* within the created superclients, namely FedSeq, FedAsyncSeq and FedSeq2Par. The approaches are summarized in Algorithm 1.

1) FedSeq

As showed in Fig. 4, the clients $\{k_{i,1}, \ldots, k_{i,|S_i|}\} \in C$ belonging to the superclient $S_i\ \forall i \in [N_S]$ form a chain performing sequential training. At each round t, the server selects a subset of $S_t \in \mathcal{S}$ superclients. Within each superclient S_i, the first device $k_{i,1}$ receives the global model f_{θ_t} from the server and locally trains it for E_k epochs on $\mathcal{D}_{k_{i,1}}$. The updated parameters $\theta_{t+1}^{k_{i,1}}$ are sent to the next client of the sequence $k_{i,2}$. Such training procedure continues until the last client $k_{|S_i|}$ updates the received model and sends it back to the server. The passage over all the clients of the chain can be repeated E_S times allowing a ring communication strategy. However, $E_S \neq 1$ leads to an increase in communication as multiple messages have to be exchanged between clients, which is why we discourage this approach and set $E_S = 1$ in our experiments. On the server-side, the superclients updates are averaged following Eq. 1. Intuitively, by training the model over multiple clients' data before the averaging step, we simulate the existence of a larger, homogeneous dataset.

2) FedAsyncSeq

In realistic scenarios, synchronous federated training can become impractical, especially when considering factors such

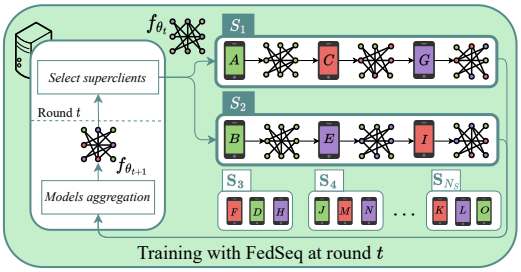

FIGURE 4. Sequential training with FedSeq. At each round t, a subset of superclients (here S_1 and S_2) is selected and receives θ_t, which is trained sequentially by the clients. Final updates are sent back to the server, where they are aggregated with FedAvg. Best seen in colors.

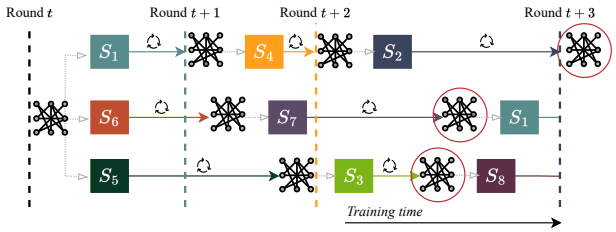

FIGURE 5. Training with FedAsyncSeq with $R = 3$. At round t, the global model is sent to the selected superclients $\{S_1, S_5, S_6\}$, having varying latency (*full arrows*). The first superclient to complete training (S_1) marks the start of the new round $t+1$. As soon as the server receives the updated model, it sends it to another superclient (*e.g.*, from S_1 to S_4). Within the time it takes for S_5 to finish training, the faster S_1 and S_4 can also complete theirs. After R rounds, the latest updates from each chain of superclients (*circled in red*) are combined, and the process begins anew. Best seen in colors.

as the latency of slower devices. The delays can be further exacerbated by FedSeq since there is no control over the capabilities of clients' systems, and multiple slow clients may be grouped together within the same superclient, leading to a significant increase in server-side waiting time. To mitigate the challenges posed by latency and ensure efficient training, we propose **FedAsyncSeq**, which **leverages sequentiality at the superclient level** while allowing asynchronous updates to be merged (Fig. 5). Instead of aggregating the models after *every* round, the server does so every R rounds. During this time frame, the model received by each superclient S_i is instantly sent to another random superclient S_j (where $i \neq j$). This approach allows the fastest chains to continue with additional training iterations instead of waiting for slower ones. After every R rounds, the most recent updates from each chain of superclients, which may originate from distinct rounds, are combined. This approach shares similarities with asynchronous settings, where models are aggregated as soon as they become available. It is worth noting that R can potentially equal T, meaning that the updated models are only averaged at the end of training. This strategy reduces the number of aggregation and synchronization steps with the server while enabling the model to potentially observe the entire dataset before averaging. This brings us closer to the centralized setting while still leveraging FL's parallelism.

3) FedSeq2Par

Sequential training of the model through a long chain of clients with highly diverse data distributions may lead to catastrophic forgetting [88]. This refers to the model forgetting the knowledge acquired from the initial users while becoming overly specialized to the most recently seen datasets [30]. Additionally, maintaining a static sequence of clients within each superclient may result in the model learning information based on the order of the clients, introducing biases or unintended patterns. Building upon [30], this extension introduces the *Sequential-to-Parallel* (STP) approach in FedSeq2Par to overcome these issues. Sequential training is exploited during the initial stages, facilitating information exchange among heterogeneous clients. **Parallelism is gradually introduced by incrementally increasing the number of superclients** and reducing their sizes, promoting faster convergence. The dynamic creation of superclients allows accounting for new clients, while eliminating potential biases related to the static nature of superclients. Formally, in each training round t, STP dynamically constructs an increasing number of superclients $\{S_i\}_{i=1}^{f_{gr}(\alpha_{gr},\beta_{gr};t)}$, where $f_{gr}(\cdot,\cdot;\cdot)$ is a non-decreasing growth function dependent on the training round t, and the hyperparameters $\alpha_{gr}, \beta_{gr} \in \mathbb{R}^+$ control the growth rate and the initial number of superclients respectively. The choice of f_{gr} is critical for the behavior of STP. As in [30], we consider three possible growth functions: *linear*, allowing for a smooth growth; *logarithmic*, promoting an initial faster dynamic; *exponential*, with slower changes at first,

$$f_{linear}(\alpha_{gr},\beta_{gr},t) = \beta_{gr}[\alpha_{gr}(t-1)+1], \tag{4}$$

$$f_{log}(\alpha_{gr},\beta_{gr},t) = \beta_{gr}[\alpha_{gr}\ln t+1], \tag{5}$$

$$f_{exp}(\alpha_{gr},\beta_{gr},t) = \beta_{gr}(1+\alpha_{gr})^{t-1}. \tag{6}$$

IV. EXPERIMENTS

This section introduces the empirical evaluations of the proposed approaches, comprising ablation studies (Sec. IV-C) and comparison with the state of the art (Sec. IV-B).

A. DATASETS

FedSeq, FedAsyncSeq and FedSeq2Par are evaluated on both vision and NLP datasets. Additionally to the CIFAR10 and CIFAR100 [89] datasets proposed in our previous [28], this extension aims to enlarge the paper's scope and introduces FEMNIST [90] for image classification, SHAKESPEARE [90] for next character prediction and STACKOVERFLOW [91] for next word prediction, all widely used as FL benchmarks [10], [15], [49]. It is important to note that *all* datasets exhibit heterogeneous distributions concerning label skew. Differently from the CIFAR datasets, FEMNIST also presents notable feature shifts attributed to diverse calligraphy styles depicting the same letter or number. Moreover, the local dataset cardinality significantly varies across clients in both FEMNIST and the NLP datasets.

TABLE 2. Comparison with state-of-the-art FL algorithms. Color coding: first, second and third best results.

Algorithm	CIFAR10			CIFAR100			FEMNIST		SHAKESPEARE		STACKOVERFLOW
	$\alpha = 0$	$\alpha = 0.2$	$\alpha = 0.5$	$\alpha = 0$	$\alpha = 0.2$	$\alpha = 0.5$	NIID	IID	NIID	IID	-
Centralized	85.64±0.07			54.97±0.19			87.52±0.27		52.00±0.02		28.50±0.20
FedAvg	71.27±0.29	76.32±0.36	77.39±0.43	42.68±0.22	48.79±0.55	49.51±0.61	81.55±0.11	83.06±0.12	48.68±0.12	48.50±0.07	24.68±0.15
FedProx	71.52±0.08	76.21±0.50	77.38±0.57	42.83±0.18	48.84±0.65	49.44±0.49	81.55±0.05	83.07±0.05	48.73±0.12	48.51±0.15	24.59±0.19
SCAFFOLD	78.82±0.15	78.02±1.13	78.51±0.24	42.17±0.10	51.06±0.03	51.03±0.12	82.56±0.07	82.70±0.01	50.91±0.19	50.91±0.12	26.10±0.15
FedDyn	**83.31**±0.15	82.31±0.41	82.97±0.40	**50.35**±0.27	**53.50**±0.76	**54.32**±0.63	81.95±0.42	82.00±0.13	**51.77**±0.03	**51.94**±0.09	25.51±0.02
FedCyclic	82.45±0.18	**82.61**±0.27	**83.49**±0.08	47.46±0.42	49.93±0.16	50.47±0.27	85.46±0.05	**87.47**±0.09	50.25±0.03	50.68±0.03	26.44±1.68
FedSeq (ours)	81.89±0.28	82.19±0.26	82.77±0.12	45.87±0.45	49.26±0.40	49.63±0.42	**87.11**±0.03	**87.48**±0.03	51.70±0.13	51.82±0.00	**28.91**±0.21
FedAsyncSeq (ours)	**83.03**±0.31	**83.17**±0.27	**83.57**±0.22	**50.23**±0.11	**51.39**±0.18	**51.27**±0.24	**86.96**±0.07	87.20±0.01	**51.82**±0.13	**51.88**±0.04	**27.44**±0.35
FedSeq2Par (ours)	**83.68**±0.14	**84.21**±0.40	**84.26**±0.06	**51.46**±0.23	**51.44**±0.02	**51.72**±0.10	**87.58**±0.15	**87.95**±0.05	**52.75**±0.20	**52.71**±0.05	**29.79**±0.33

In order to set up a heterogeneous scenario for CIFAR10 and CIFAR100, the local class distribution is sampled from a Dirichlet distribution with $\alpha \in \{0, 0.2, 0.5\}$ [40]. Both CIFAR datasets are divided into 500 clients with 100 images each. The IID and non-IID ("NIID" for short) data distributions of FEMNIST introduced in [90] follow the writers' ownership, *i.e.* each client is a distinct writer. The NIID split accounts for both label skew and feature shift. The IID and NIID distributions of SHAKESPEARE [90] reflect some Shakespearean characters, with 100 clients owning around 3,743 samples each, while the implementation of STACKOVERFLOW follows [49]. FEMNIST and STACKOVERFLOW better represent realistic cross-device settings thanks to a larger number of clients: $K = 3,500$ and $40k$ respectively. Additional information can be found in Appendix A. As proposed by previous works [3], [15] accounting for the learning trends instability, the results are averaged over the last 25 rounds on SHAKESPEARE and 100 rounds on CIFAR10/100 and FEMNIST. As done in [49], due to its prohibitively large number of clients and examples, testing on STACKOVERFLOW full dataset is performed only at the end and a subsample is used during training.

B. COMPARISON WITH STATE-OF-THE-ART FL ALGORITHMS

1) Algorithms

To validate the effectiveness of the proposed approaches, this study conducts a comparison with state-of-the-art (SOTA) algorithms for heterogeneous FL. In addition to the standard FedAvg [1] described in Sec. III, FedSeq and its variants are tested against FedProx [10], which adds a regularization term in the local loss function to encourage proximity between clients' and server's model parameters, SCAFFOLD [13], addressing the client drift via stochastic variance reduction, FedDyn [15], that aligns local and global stationary points, and FedCyclic [67], which cyclically exchanges models across clients without relying on a central server. FedCyclic can be seen as the extreme case of FedSeq with $N_S = 1$ and no server-side aggregation. FedCyclic was chosen over FedStar (from the same paper [67]) due to the latter's impractical communication overhead in realistic scenarios. FedStar indeed requires each client to send its updates to *all* the other participants, leading to an exponential increase in communication costs. To ensure a fair comparison, FedCyclic is not trained on all clients at each round but randomly selects a fraction C of the available ones. The same applies to all the other approaches and, most importantly, we ensure that the number of model updates within rounds is the same for all the compared methods. FedSeq and FedSeq2Par build superclients using ψ_{t2v} and the best corresponding $\phi_{(\cdot)}$ (see Sec. IV-C). Local pre-training runs for 10 epochs chosen from $\{1, 5, 10, 20, 30, 40\}$ (see Appendix B for details). We select $R = N_S$ in FedAsyncSeq.

2) Final performance

Table 2 shows that FedSeq and its extensions are either competitive or outperform other SOTA algorithms on all tasks and datasets, especially on severe heterogeneous data distributions. We point out that both SCAFFOLD and FedDyn require stateful clients, and SCAFFOLD doubles the size of the communicated message, differently from this paper's approaches. Being the extreme case of FedSeq with $N_S = 1$, FedCyclic reaches similar performances on most of the datasets, implying that having one single superclient containing all clients does not dramatically increase performances and instead hugely increments the amount of training time, as each client needs to wait for the previous one's update. Focusing on the extensions of FedSeq, both FedAsyncSeq and FedSeq2Par improve the baseline's performances on all tasks. We note that aggregating superclients updates every N_S rounds not only requires less frequent synchronization between clients and server, but also improves the reached accuracy. FedSeq2Par achieves the best results in most cases, exploiting the benefits of sequential training and parallelism, being the second best to FedDyn only in the case of $\alpha = 0.2$ and $\alpha = 0.5$ in CIFAR100. However, differently from FedSeq2Par, FedDyn relies on stateful clients, posing a significant challenge for real-world deployments with billions of edge devices [3], [92]. In such large-scale settings, individual devices are unlikely to be called upon for multiple training rounds. This transience renders their local states obsolete quickly, compromising their effectiveness in subsequent training iterations. The results obtained by FedDyn on the more realistic FEMNIST (3,500 clients) and STACKOVERFLOW ($40k$ clients) datasets un-

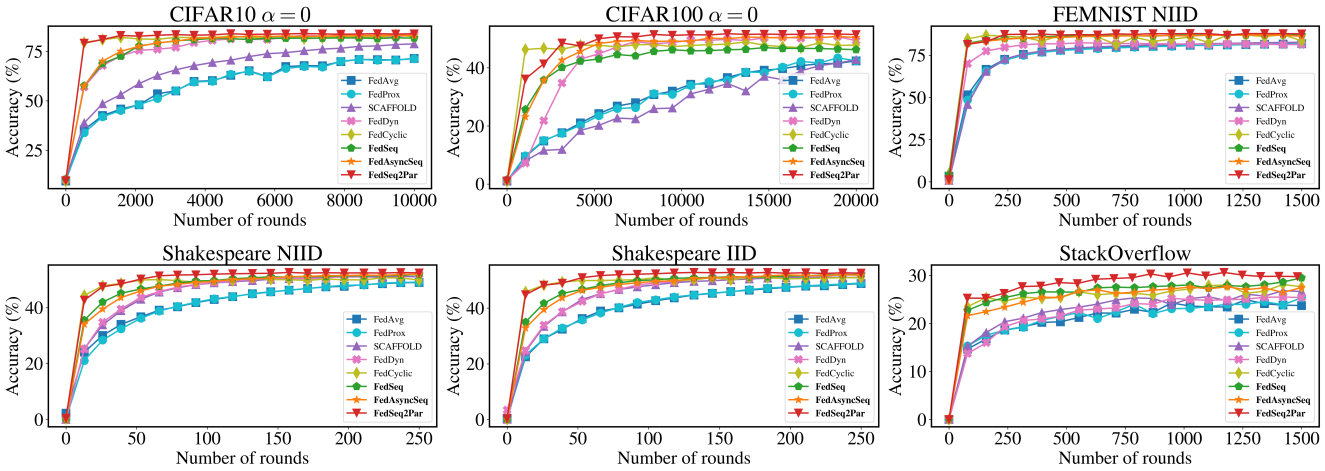

FIGURE 6. Accuracy convergence plots of FedSeq, FedSeq2Par, FedAsyncSeq (in bold) and SOTA algorithms on vision and NLP datasets. On average, FedSeq2Par is the best-performing algorithm. All the proposed approaches can be distinguished for their improved speed. Best seen in colors. Full results are reported in the Appendix E.

derscore this point. On FEMNIST, FedDyn outperforms the baseline FedAvg by only ≈ 0.4 points in accuracy. In contrast, FedSeq's variants, especially FedSeq2Par, achieve a significant improvement of +6 points over FedAvg. On the even larger STACKOVERFLOW dataset, FedDyn shows a loss of ≈ 0.4 points compared to the baseline, while FedSeq2Par exhibits a gain of ≈ 1.3 points. These results confirm the limitations of FedDyn in real-world cross-device scenarios, where its reliance on stateful clients becomes a significant disadvantage.

3) Convergence speed

Fig. 6 evidently shows that FedSeq and its extensions not only achieve superior results but also exhibit accelerated performance. Fig 7 compares the rounds necessary to each algorithm to reach 70% and 90% of the centralized accuracy on CIFAR10/100 and FEMNIST. Table E.1 in Appendix E integrates the results for the other datasets. FedSeq consistently demonstrates significant improvements in convergence speed across all tasks, achieving a speed-up factor of over $18x$ on FEMNIST and $10x$ on STACKOVERFLOW, presenting high cross-device variability. Achieving superior overall accuracy, asynchronous training, and reduced latency compared to FedSeq does not compromise the convergence speed of FedAsyncSeq. FedSeq2Par notably improves convergence speed on all tasks and datasets through its STP approach. FedDyn enhances the convergence rates of FedAvg but experiences parameter explosion in highly imbalanced settings [92], requiring gradient clipping techniques. Among the considered methods, FedCyclic achieves comparable or better results than FedSeq2Par. However, it is important to note that FedCyclic, as an extreme case of FedSeq with $N_S = 1$, eliminates any form of parallelism inherent in distributed and FL settings. The results highlight the significance of sequential training for rapid convergence in initial rounds, while the superior performance of FedSeq2Par in later stages underlines the role of parallelism in achieving both improved final performance and convergence speed up.

4) Communication cost

Communication is the main bottleneck of federated training [93], due to the overload of the networks and message size. Thus, when comparing the performance of FL algorithms, their impact on the communication cost is of the utmost importance. As already shown in Sec. IV-B3, our methods speed up the convergence, implying that fewer communication rounds are needed to reach a target performance. This study additionally compares the number of client-server exchanges required by the proposed method (FedSeq) with leading SOTA algorithms. Table 3 demonstrates that FedSeq achieves **less network communication** thanks to its client-to-client approach. Similar analyses can be easily extended to FedSeq2Par and FedAsyncSeq. Given the total number of clients K, the fraction selected at each round C, the total number of superclients N_S, and the rounds T, we first analyze the case in which all superclients are equally sized, *i.e.*, $K_{S_i} = K_{S_j} = {}^{K}/_{N_S} =: K_S \, \forall i \neq j$. In FedAvg, the server sends the global model to the $C \cdot K$ selected clients, which then send back the updated version. As summarized in Table 3, this process accounts for $2C \cdot K$ exchanges over the network. The same goes for FedProx and FedDyn. SCAFFOLD requires double the communication. In FedSeq with equal K_S instead, the server-to-client (S2C) and client-to-server (C2S) communication only happens between the first and last clients of the chain of each superclient respectively. Since the server selects $C \cdot N_S$ superclients, the process sums up to $2C \cdot N_S$ exchanges. Moreover, within each superclient, the clients

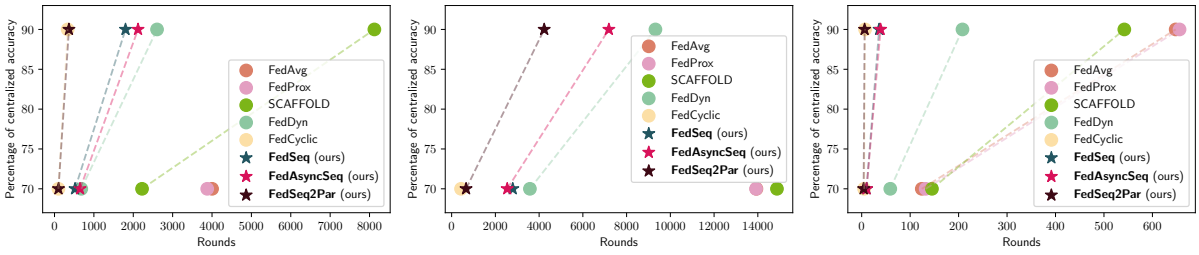

CIFAR10 $\alpha = 0$ **CIFAR100** $\alpha = 0$ **FEMNIST NIID**

FIGURE 7. Convergence rates in non-i.i.d. scenarios. Each plot shows the rounds necessary for each method to reach 70% and 90% of the centralized accuracy. Not all the algorithms reach the 90% target (missing line). Our methods (*in bold, stars*) outperform the others in all settings. Best seen in colors.

TABLE 3. Number of communication exchanges from server to client (C2S), client to server (S2C) and client to client (C2C) at each round t and across all rounds T.

Method	S2C	C2S	C2C	Total @ round	Total T rounds
FedAvg	CK	CK	0	$2CK$	$2TCK$
FedProx	CK	CK	0	$2CK$	$2TCK$
FedDyn	CK	CK	0	$2CK$	$2TCK$
SCAFFOLD	$2CK$	$2CK$	0	$4CK$	$4TCK$
FedCyclic	1	1	$CK-1$	$CK+1$	$T(CK+1)$
FedSeq $K_{S_i} = K_{S_j}$	CN_S	CN_S	$\left(\frac{K}{N_S}-1\right)CN_S$	$C(N_S+K)$	$TC(N_S+K)$
FedSeq $K_{S_i} \neq K_{S_j}$	CN_S	CN_S	$\sum_{S_i \in \mathcal{S}_t} K_{S_i} - CN_S$	CN_S+ $+\sum_{S_i \in \mathcal{S}_t} K_{S_i}$	TCN_S+ $+\sum_{t \in [T]} \sum_{S_i \in \mathcal{S}_t} K_{S_i}$

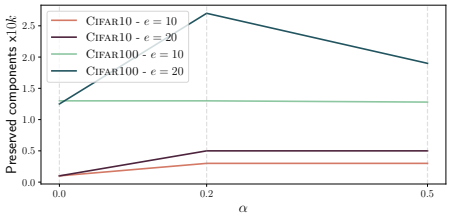

FIGURE 8. CIFAR datasets. Ratio of the preserved components after applying PCA with 90% of explained variance when varying the number of local epochs e.

exchange messages following the chain, for a total of $K_S - 1$ transmissions $\forall S$. If we consider all $C \cdot N_S$ groups involved, this is equivalent to

$$(K_S - 1)\, C \cdot N_S = \left(\frac{K}{N_S} - 1\right) C \cdot N_S = C \cdot K - C \cdot N_S. \qquad (7)$$

By summing everything up, the total is $2C \cdot N_S + C \cdot K - C \cdot N_S = C(N_S + K)$. Since $N_S < K$, **the overall communication cost of FedSeq is smaller than FedAvg, FedProx, FedDyn and SCAFFOLD**. If superclients are not equally sized, the client-to-client (C2C) cost is $\sum_{S_i \in \mathcal{S}_t}(K_{S_i} - 1) = \sum_{S_i \in \mathcal{S}_t} K_{S_i} - C \cdot N_S$, where $|\mathcal{S}_t| = C \cdot N_S$, and the total becomes $C \cdot N_S + \sum_{S_i \in \mathcal{S}_t} K_{S_i}$, *i.e.*, depends on the size of the selected superclients. However, to ensure a fair comparison, we select $K_{S,max}$ s.t. $K/K_{S,max} \approx N_S$, *i.e.*, most of the superclients are of the same size, falling back to the first scenario. This implies that **FedSeq always has the lowest cumulative cost**. Lastly, FedCyclic is a limit case of FedSeq with one superclient made of $C \cdot K$ clients, with a total cost of $T \cdot C(K + 1/C)$.

C. ABLATION STUDIES AND ANALYSES

This section discusses the impact of each method component introduced in Sec. III.

1) Estimating clients' data distribution

The proposed method utilizes the parameters (ψ_{clf}) or pre-trained model predictions ($\psi_{\text{conf}}, \psi_{\text{t2v}}$) to compute a privacy-preserving estimate of the clients' dataset distribution. To mitigate the *curse of dimensionality* [94] on the classifier parameters in ψ_{clf}, PCA [95] is applied, keeping 90% of the explained variance. Fig. 8 shows that the percentage of preserved components decreases with the complexity of the dataset, *e.g.* fewer components are needed for CIFAR10, and increases directly proportional to e. As for ψ_{conf}, not to severely impact the original dataset, $\mathcal{D}_{pub}$ is built using 10 images per class from the test set for computing the *confidence vectors* (Eq. 2). Once $\mathcal{D}_{pub}$ has served its purpose, it is not used again.

Since a public dataset capturing the overall global distribution may not be available in realistic settings, this paper introduces ψ_{t2v}, based on Task2Vec [29], which presents two main advantages: *i*) no external dataset is required, and *ii*) clients only fine-tune the classifier, reducing the latency. Following [29], we use a pre-trained ResNet18 for image classification tasks, and a GPT-2-like [96] language modeling transformer for NLP tasks. To better understand the difference in their behavior, the embeddings of ψ_{t2v} (*right*) and ψ_{conf} (*left*) are compared in Fig. 9a. Specifically, we illustrate the distance between their embeddings computed over the first 75 clients of CIFAR100 with $\alpha = 0$. The first 25 clients exclusively have images of aquatic mammals (beavers, dolphins, otters, seals, and whales), the next 25 clients have images of fishes (aquarium fishes, flatfishes, rays, sharks, and trouts), and the last 25 clients have images of various flowers. Vectors from ψ_{conf} lack class similarity representation, while the ψ_{t2v} distance matrix reveals that clients with fish and aquatic mammal images (*red square*) cluster together more closely than those with flower images. In Fig. 9b, one client

TABLE 4. FedSeq baselines: comparison of grouping criteria by varying ϕ, ψ and τ. Results in terms of accuracy (%).

Method	ψ	ϕ	τ	CIFAR10			CIFAR100			FEMNIST		SHAKESPEARE	
				$\alpha=0$	$\alpha=0.2$	$\alpha=0.5$	$\alpha=0$	$\alpha=0.2$	$\alpha=0.5$	NIID	IID	NIID	IID
FedSeq	-	Random	-	81.90	82.09	82.12	46.39	48.62	49.44	87.07	87.49	51.55	51.84
	clf	K-means	Euclidean	**82.30**	81.78	82.48	44.91	48.74	49.60	87.07	87.45	51.78	51.70
		Greedy	Cosine	79.95	82.06	83.32	45.22	48.92	49.62	87.12	87.42	51.72	51.88
	conf	K-means	Euclidean	82.04	81.99	82.37	43.55	49.43	49.79	87.10	**87.50**	51.79	51.87
		Greedy	KL	82.21	82.20	82.22	45.97	49.56	49.82	87.01	87.46	51.70	51.65
		Greedy	Gini Index	82.09	81.85	82.71	45.79	48.98	49.61	87.05	87.42	51.59	**51.98**
	t2v	Greedy	Norm-Cosine	82.28	82.48	**82.86**	**46.51**	**50.06**	**50.31**	87.11	87.48	51.65	51.82
		ICG	Euclidean	82.05	**82.51**	82.54	46.33	49.13	49.86	**87.18**	87.45	**51.84**	51.77
FedAsyncSeq	*t2v*	*	*	83.40	83.37	83.85	50.38	51.64	51.58	86.96	87.20	51.93	52.02
FedSeq2Par	*t2v*	ICG	Euclidean	**83.86**	**84.66**	**84.35**	**51.23**	**51.41**	**51.78**	**87.58**	**87.96**	**52.84**	**52.54**

(*): (Greedy, Norm-Cosine) for CIFAR10/100; (K-means, Euclidean) for FEMNIST and SHAKESPEARE.

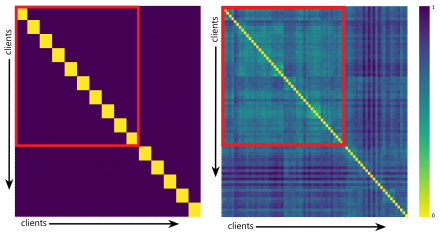

(a) Distance matrices on ψ_{conf} (*left*) and ψ_{t2v} (*right*) tasks embeddings.

(b) ψ_{conf} and ψ_{t2v}'s tasks embeddings similarity.

FIGURE 9. CIFAR100, $\alpha = 0$. *(a)* Focus on 75 clients. Each group of 25 clients has access to either images of aquatic mammals, fishes or flowers. *(b)* Focus on client with images of *whales*. Comparison of embedding distances with clients containing images of progressively different entities. ψ_{t2v} accurately recognizes the similarities between animals, in contrast to ψ_{conf}.

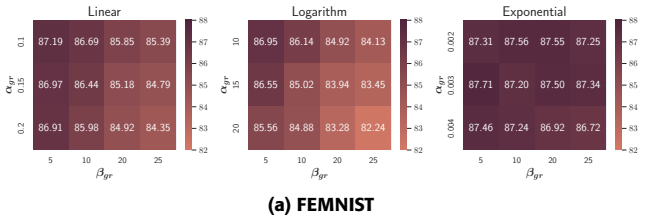

(a) FEMNIST

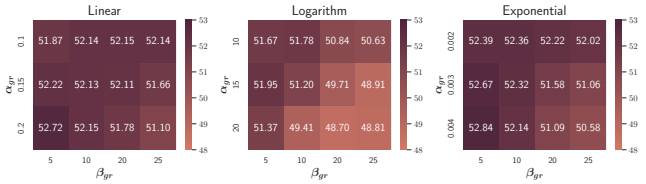

(b) SHAKESPEARE

FIGURE 10. Sensitivity of FedSeq2Par to f_{gr} and the growth parameters α_{gr} and β_{gr}. Results in test accuracy (%) on the NIID splits.

with only whale images is compared in terms of distance with other clients having progressively dissimilar images to whales. Once again, ψ_{t2v} accurately recognizes the similarities between animals, in contrast to ψ_{conf}. To understand this behavior, we note that the embedding of the k-th client with samples belonging to class $c \in [N_c]$ given by ψ_{conf} is $p_k[i] \approx \begin{cases} 1 & if\, i = c, \\ 0 & otherwise \end{cases}$. This aligns with our expectations, as $f_{\theta_0^k}$ is trained to classify observations with label c. As a result, the embeddings of clients seeing different classes (regardless of the similarity of the depicted subjects) are equally distant. However, this contradicts our intuitive understanding, as we would expect that similarities in data distributions would manifest as closeness in the vector space. In contrast, the distance between ψ_{t2v} embeddings aligns with our intuition on semantic and taxonomic relations among entities. This behavior is evidently reflected in its performance in Table 4, where ψ_{t2v} consistently outperforms the previous best approach, ψ_{conf} [28], for both vision and NLP tasks.

2) Grouping criteria

Table 4 compares the different combinations of grouping criteria G_S. As for ψ_{kmeans}, a reasonable value for the number of clusters is N_c, and the Euclidean distance is used to compare the resulting superclients. The confidence vectors extracted by ψ_{conf} have the form of a probability distribution (Sec. III), additionally comparable via disomogeneity measures such as the KL divergence and the Gini Index. The normalized embeddings obtained with Task2Vec are compared with the cosine distance ("Norm-Cosine" in the Table) [29]. Notably, ϕ_{rand} returns groups obtaining competitive results with the other grouping methods. The reason lies in statistical considerations on the cross-device setting: with the number of clients being large in all datasets, a randomly created group is unlikely to contain clients belonging all to the same data distribution. ψ_{t2v} achieves the best performance across all settings, demonstrating effective capture of task similarities. We select ϕ_{icg} as grouping method due to its satisfying results and efficiency, useful in the dynamic creation of superclients especially with several groups.

3) FedSeq2Par

As described in Sec. III-C3, the Sequential-to-Parallel approach used in FedSeq2Par is based on the function $f_{gr}(\alpha_{gr}, \beta_{gr}, t)$ (Eq. 4-6), that defines the number of super-

TABLE 5. Parallelism and test accuracy: FedSeq *vs* FedSeq2Par.

Dataset	N_S	$\overline{N}_S$	Parallelism	Accuracy ($\alpha = 0$ / NIID)	
	FedSeq	FedSeq2Par	↑	FedSeq	FedSeq2Par
CIFAR10	50	**104**	2.09x	81.89	**83.68**
CIFAR100	50	**113**	2.26x	45.87	**51.46**
FEMNIST	175	**383**	2.19x	87.11	**87.58**
SHAKESPEARE	25	**52**	2.09x	51.70	**52.75**
STACKOVERFLOW	1900	**5966**	3.14x	28.91	**29.79**

clients at each round t. This section aims to understand which growth function better suits the analyzed settings (*linear*, *logarithmic*, or *exponential*) and the effect of the parameters α_{gr} (growth rate) and β_{gr} (initial number of superclients). α_{gr} is chosen so that a fully parallel scenario is reached in the last rounds, while favoring sequential training at the beginning. We test $\beta_{gr} \in \{5, 10, 20, 25\}$ for all datasets except for the larger STACKOVERFLOW, for which we select $\beta_{gr} \in \{50, 100, 200, 250\}$. Fig. 10 analyzes the impact of these parameters on the NIID splits of FEMNIST and SHAKESPEARE. Notably, starting with the smallest number of superclients β_{gr} consistently yields superior performance, as it exploits sequentiality more. f_{exp} has the best and most consistent results. Due to the uniformity of these results, the same configuration is maintained across all datasets. We further compare FedSeq2Par with FedSeq in terms of number of superclients and final performance in Table 5. For FedSeq2Par, we compute the average number of superclients $\overline{N}_S$ across rounds. We note that $\overline{N}_S$ is constantly larger than N_S, implying a more parallelized scenario *on average* with FedSeq2Par w.r.t. FedSeq even if $\beta_{gr} \ll N_S$. This behavior positively reflects on the final performance, confirming the efficacy of the STP approach.

V. PRIVACY

Recent FL literature has highlighted the potential for attackers to reconstruct sensitive information through the clients' updates [36]. Thus, concerns on the potential privacy implications of the *client-to-client* sequential training approach introduced by FedSeq arise. Specifically, this extension poses the question: *does FedSeq's client-to-client sequential training facilitate the retrieve of previous users' personal information by a malicious client?* To answer, FedSeq is evaluated against two famous attacks, namely the **label flipping** [35] (LFA) and the **GAN recovery** attacks (GRA) [75], and study potential private information leakages. The well-known gradient inversion attack [36] is not considered here, as its assumptions do not align with our approach (*e.g.*, access of the attacker to both initial and updated models, knowledge of private labels). Differently, in this case, clients only receive the updated parameters from the previous user and potentially malicious clients are not aware of other users' private labels.

A. LABEL FLIPPING ATTACK

LFA is an *active* privacy attack aiming at deteriorating the global model performances by switching labels at training time. Here, the focus is on models solving the classification task. To mislead the global model classification ability, the set of malicious clients $\mathcal{A} := \{a_i\}_{i=1}^{L \cdot K} \subseteq C$ with $L \in [0, 1]$ willingly swaps the labels of their local data following a set of criteria $\{\gamma_i\}_{i=1}^{L \cdot K}$. The criterion γ_i defines the labels to be swapped during the attack for each attacker a_i. For instance, $\gamma_i = \gamma_j$ implies that the attackers a_i and a_j will swap the same classes. This work tests two possible situations:

1) Different attackers swap distinct classes, *i.e.* each attacker a_i chooses its γ_i independently (γ_{random}),
2) All the attackers swap the same classes, *i.e.* $\gamma_i = \gamma_j \, \forall i, j \in \mathcal{A}$ (γ_{fixed}).

B. GAN RECOVERY ATTACK

GRA is a *passive* privacy attack that aims at reconstructing other clients' private information using GAN architectures [97]. It is important to highlight that the primary objective of GANs is to *generate* samples that closely resemble those found in the training set without direct access to the original ones. GANs rely on interactions with a *discriminative* deep neural network to learn and capture the underlying data distribution [75]. They are trained to mimic the images encountered by the discriminative network, starting from random initialization. However, a potential concern arises when the discriminator is trained on private data, as it can potentially be exploited to train a generator network capable of reconstructing the sensitive data. This poses significant privacy and security concerns. Formally, the GANs' optimization problem [75] is

$$\min_{\theta_G} \max_{\theta_D} \sum_{i=1}^{n} \log f(x_i, \theta_D) + \sum_{j=1}^{n} \log(1 - f(g(z_j, \theta_G), \theta_D)), \quad (8)$$

where $f(x, \theta_D) : \mathcal{X} \to \mathcal{Y}$ is a discriminative network parametrized by θ_D that, given an image, outputs a class label. The generative network $g(z, \theta_G) : \mathcal{X} \to \mathcal{X}$ receives random noise as input and outputs an image. x_i is the original image and $g(z_j)$ is a randomly generated one. In GRA, at round t, an attacker $a_i \in C$ disguised as a client exploits the incoming trained model θ_t as the discriminator of a *GAN*, *i.e.* $\theta_D \leftarrow \theta_t$. The generator g is then trained for E_a epochs to reconstruct inputs similar to the ones previously accessed by θ_t, thus breaking the clients' privacy.

C. TESTING FEDSEQ PRIVACY RESILIENCE

This section provides quantitative results of FedSeq's resilience against the *LFA* and *GRA* attacks. We show that FedSeq not only does not introduce additional privacy liabilities w.r.t. FedAvg, but it learns more robust models.

1) FedSeq against LFA

Table 6 summarizes the results on the different setups proposed to evaluate the robustness of FedSeq to the LFA attack. We distinguish between the fraction of malicious *superclients* L_S and the fraction of malicious clients *within* each malevolent superclient L. The corresponding fraction of attackers in FedAvg becomes $L_S \cdot L$. We test L_S in $\{0.1, 0.3, 0.5\}$ and L in

TABLE 6. Label Flipping Attack experiments after 1k rounds. Results in accuracy (%) and drop in accuracy (↓) w.r.t. to the reference. In bold smaller drops in each attack. Symbols: "◦" (negligible or non-existing drops), "Fixed" (γ_{fixed}) and "Random" (γ_{random}).

		CIFAR10						CIFAR100					
L_S	L	Swapped Labels	γ	FedSeq Accuracy ↑	FedSeq Drop ↓	FedAvg Accuracy ↑	FedAvg Drop ↓	Swapped Labels	γ	FedSeq Accuracy ↑	FedSeq Drop ↓	FedAvg Accuracy ↑	FedAvg Drop ↓
0.1	0.1	0↔2	Fixed	76.06	◦	48.72	◦	$INTRA_{SC}$	Fixed	38.81	−0.89	12.57	**−0.28**
	0.5			75.92	**−0.20**	47.55	−1.14			38.87	−0.83	12.35	**−0.50**
0.3	0.1			76.25	**−0.15**	47.74	−0.95			39.66	◦	12.45	−0.40
	0.5			74.75	−1.35	48.30	**−0.39**			39.04	**−0.66**	11.55	−1.30
0.5	0.1			75.65	**−0.45**	47.55	−1.14			39.23	**−0.47**	12.35	−0.50
	0.5			75.50	**−0.60**	47.44	−1.25			38.40	**−1.30**	11.02	−1.83
0.1	0.1	3↔5	Fixed	76.05	◦	47.92	−0.77		Random	39.36	−0.34	13.23	◦
	0.5			75.50	**−0.60**	47.93	−0.76			39.59	◦	13.05	◦
0.3	0.1			75.75	**−0.35**	47.80	−0.89			39.67	◦	13.23	◦
	0.5			75.28	**−0.82**	47.74	−0.95			38.88	−0.82	12.47	**−0.38**
0.5	0.1			76.15	◦	47.93	−0.76			39.68	◦	13.05	◦
	0.5			75.09	**−1.01**	46.94	−1.75			38.98	**−0.72**	11.27	−1.58
0.1	0.1	0↔2 ↔3↔5	Fixed	76.65	◦	47.36	−1.33	$EXTRA_{SC}$	Fixed	40.11	◦	12.37	−0.48
	0.5			76.01	◦	48.12	−0.58			40.03	◦	12.20	−0.65
0.3	0.1			75.79	**−0.31**	48.29	−0.40			39.29	−0.41	13.08	◦
	0.5			75.19	**−0.91**	46.42	−2.27			38.73	−0.97	12.34	**−0.51**
0.5	0.1			75.35	−0.75	48.12	**−0.58**			38.95	−0.75	12.20	**−0.65**
	0.5			75.09	**−1.01**	46.94	−1.75			38.46	−1.24	12.30	**−0.55**
0.1	0.1		Random	76.52	◦	48.65	◦		Random	39.17	−0.53	13.12	◦
	0.5			75.71	**−0.39**	47.52	−1.17			39.12	−0.58	12.83	◦
0.3	0.1			75.98	**−0.12**	47.01	−1.68			39.54	◦	13.06	◦
	0.5			74.75	**−1.35**	46.70	−1.99			37.90	−1.80	11.81	**−1.04**
0.5	0.1			75.93	**−0.17**	47.52	−1.17			39.03	−0.67	13.12	◦
	0.5			73.46	−2.64	46.49	**−2.20**			37.76	−1.94	11.94	**−0.91**
		Reference accuracy: **FedSeq** 76.10% - **FedAvg** 48.69%						Reference accuracy: **FedSeq** 39.70% - **FedAvg** 12.85%					

$\{0.1, 0.5\}$. For example, $L_S = 0.1$ and $L = 0.5$ implies that 10% of the superclients are malevolent, and 50% of their clients are attackers. Following [35], the four swapped classes in CIFAR10 are: *airplanes* (label 0) exchanged with *birds* (label 2), and *dogs* (label 5) with *cats* (label 3). We additionally swap all the aforementioned classes (0,2,3,5) at the same time. When using CIFAR100 instead, the concept of "superclass" proper of the dataset is exploited (*e.g.*, *aquatic mammals*, *flowers*). We either swap 20 classes that do not belong to the same superclass, *i.e.* one class for each superclass (*e.g.*, *dolphins* and *roses*), or exchange pair of 20 labels belonging to the same superclass (*e.g.*, *dolphins* with *whales*). We refer to the former as $EXTRA_{SC}$, and to the latter as $INTRA_{SC}$. To evaluate FedSeq against the easiest scenario for the attacker, all the experiments are run with $\alpha = 100$ on both CIFAR10 and CIFAR100, meaning that all K clients see all the classes, and the LFA is always feasible. T is set equal to $1k$. Table 6 shows the results of the attack on each proposed configuration, analyzing both the accuracy of the model on the overall test set and the drop w.r.t. the reference experiment without attackers. On average, the *fixed* attacks are more effective than the *random* ones. The reason behind this behavior is intuitive: when using γ_{fixed}, the attackers never let the model learn the correct patterns for classifying the swapped labels, differently from the random acting. Swapping 4 labels in CIFAR10 rather than 2 brings on average more damage. For CIFAR100, the $EXTRA_{SC}$ attacks are significantly more effective: the model likely learns some common features for images belonging to the same superclass, leading to a reduced efficacy of the $INTRA_{SC}$ attack. Importantly, FedSeq outperforms FedAvg on most scenarios both in terms of accuracy and drop w.r.t. to the reference: this means FedSeq is still able to achieve faster convergence if under attack, and is more robust than FedAvg.

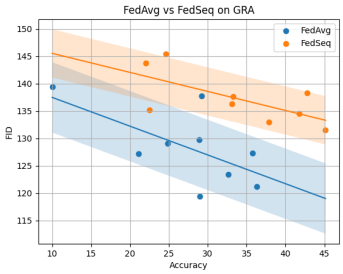

(a) FID scores after GRA attack on FedSeq and FedAvg

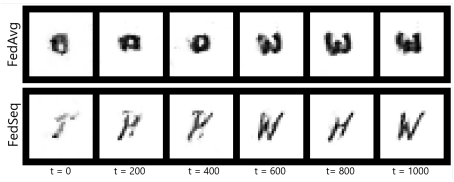

(b) GRA attacker images reconstruction

FIGURE 11. (a) GRA attack on global model with different accuracy. The resulting FID scores on FedSeq are consistently higher, implying a less effective attack. (b) Examples of images reconstructed by the GRA attacker at distinct rounds.

2) FedSeq against GRA

The *Fréchet inception distance* (FID) [98] assesses the quality of the images created by a generative model. Given a dataset $\mathcal{D}$ and its reconstruction $\hat{\mathcal{D}}$, the FID measures the distribution of their features, extracted using an InceptionV3 network [99], using the *Fréchet distance* [98]. A lower score indicates better-quality images. Within the context of an attack, the

FID has to be as large as possible, signifying the attacker's inability to reconstruct private data effectively. Unlike the approach in [75], we refrain from incorporating a "fake" class in the classifier, deeming it unrealistic. Instead, we allow the attacker to utilize an additional binary dense layer on top of the model to distinguish between "fake" and "real" data. Fig. 11a resulting from attacks on models with varying levels of accuracy. It is clear that the attack conducted on FedSeq consistently yields higher FID scores in comparison to FedAvg, underscoring its enhanced privacy characteristics. Fig. 11b shows some examples of images reconstruction at different rounds.

VI. CONCLUSION

This work addresses the issues arising from the inherent statistical heterogeneity in Federated Learning (FL) introducing FedSeq. FedSeq leverages sequential training among groups of heterogeneous clients (*superclients*) to obtain more robust models before the server-side averaging step. Various strategies are proposed to effectively group clients according to their data distribution. To reduce the waiting time due to the latency of the slowest superclients, we develop FedAsyncSeq, which allows asynchronous communication between clients and server. Lastly, to exploit sequentiality and parallelism at their best, FedSeq2Par dynamically changes the number of superclients at each round. The extensive experiments on multiple FL benchmarks prove the efficacy of our approaches, in terms of final performances, convergence speed and privacy resilience.

REFERENCES

[1] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017.

[2] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020.

[3] Debora Caldarola, Barbara Caputo, and Marco Ciccone. Improving generalization in federated learning by seeking flat minima. In *European Conference on Computer Vision*, pages 654–672. Springer, 2022.

[4] Andrew Hard, Kanishka Rao, Rajiv Mathews, Swaroop Ramaswamy, Françoise Beaufays, Sean Augenstein, Hubert Eichner, Chloé Kiddon, and Daniel Ramage. Federated learning for mobile keyboard prediction. *arXiv:1811.03604*, 2018.

[5] Liu Yang, Ben Tan, Vincent W Zheng, Kai Chen, and Qiang Yang. Federated recommendation systems. *Federated Learning: Privacy and Incentive*, pages 225–239, 2020.

[6] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Federated visual classification with real-world data distribution. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X 16*, pages 76–92. Springer, 2020.

[7] Lidia Fantauzzo, Eros Fanì, Debora Caldarola, Antonio Tavera, Fabio Cermelli, Marco Ciccone, and Barbara Caputo. Feddrive: Generalizing federated learning to semantic segmentation in autonomous driving. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 11504–11511. IEEE, 2022.

[8] Donald Shenaj, Eros Fanì, Marco Toldo, Debora Caldarola, Antonio Tavera, Umberto Michieli, Marco Ciccone, Pietro Zanuttigh, and Barbara Caputo. Learning across domains and devices: Style-driven source-free domain adaptation in clustered federated learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 444–454, 2023.

[9] Felix Sattler, Simon Wiedemann, Klaus-Robert Müller, and Wojciech Samek. Robust and communication-efficient federated learning from non-iid data. *IEEE TNNLS*, 31(9):3400–3413, 2019.

[10] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2:429–450, 2020.

[11] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data. *arXiv:1806.00582*, 2018.

[12] Qinbin Li, Yiqun Diao, Quan Chen, and Bingsheng He. Federated learning on non-iid data silos: An experimental study. *CoRR*, abs/2102.02079, 2021.

[13] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *ICML*, pages 5132–5143. PMLR, 2020.

[14] Tailin Zhou, Zehong Lin, Jun Zhang, and Danny HK Tsang. Understanding model averaging in federated learning on heterogeneous data. *arXiv:2305.07845*, 2023.

[15] Durmus Alp Emre Acar, Yue Zhao, Ramon Matas Navarro, Matthew Mattina, Paul N Whatmough, and Venkatesh Saligrama. Federated learning based on dynamic regularization. *ICLR*, 2021.

[16] Zhe Qu, Xingyu Li, Rui Duan, Yao Liu, Bo Tang, and Zhuo Lu. Generalized federated learning via sharpness aware minimization. In *ICML*, pages 18250–18280. PMLR, 2022.

[17] Yan Sun, Li Shen, Tiansheng Huang, Liang Ding, and Dacheng Tao. Fedspeed: Larger local interval, less communication round, and higher generalization accuracy. *ICLR*, 2023.

[18] Virginia Smith, Chao-Kai Chiang, Maziar Sanjabi, and Ameet Talwalkar. Federated multi-task learning. *NeurIPS*, 2017.

[19] Hadi Jamali-Rad, Mohammad Abdizadeh, and Anuj Singh. Federated learning with taskonomy for non-iid data. *IEEE TNNLS*, pages 1–12, 2022.

[20] Canh T. Dinh, Tung T. Vu, Nguyen H. Tran, Minh N. Dao, and Hongyu Zhang. A new look and convergence rate of federated multitask learning with laplacian regularization. *IEEE TNNLS*, pages 1–11, 2022.

[21] Felix Sattler, Klaus-Robert Müller, and Wojciech Samek. Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints. *IEEE TNNLS*, 2020.

[22] Kavya Kopparapu and Eric Lin. Fedfmc: Sequential efficient federated learning on non-iid data. *arXiv:2006.10937*, 2020.

[23] Christopher Briggs, Zhong Fan, and Peter Andras. Federated learning with hierarchical clustering of local updates to improve training on non-iid data. In *International Joint Conference on Neural Networks*, pages 1–9. IEEE, 2020.

[24] Ming Xie, Guodong Long, Tao Shen, Tianyi Zhou, Xianzhi Wang, Jing Jiang, and Chengqi Zhang. Multi-center federated learning. *arXiv:2108.08647*, 2021.

[25] Debora Caldarola, Massimiliano Mancini, Fabio Galasso, Marco Ciccone, Emanuele Rodolà, and Barbara Caputo. Cluster-driven graph federated learning over multiple domains. In *Proceedings of the IEEE/CVF CVPR*, pages 2749–2758, 2021.

[26] Yihan Yan, Xiaojun Tong, and Shen Wang. Clustered federated learning in heterogeneous environment. *IEEE TNNLS*, 2023.

[27] 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. *arXiv:1912.04977*, 2019.

[28] Riccardo Zaccone, Andrea Rizzardi, Debora Caldarola, Marco Ciccone, and Barbara Caputo. Speeding up heterogeneous federated learning with sequentially trained superclients. In *26th International Conference on Pattern Recognition*, pages 3376–3382. IEEE, 2022.

[29] Alessandro Achille, Michael Lam, Rahul Tewari, Avinash Ravichandran, Subhransu Maji, Charless C Fowlkes, Stefano Soatto, and Pietro Perona. Task2vec: Task embedding for meta-learning. In *Proceedings of the IEEE/CVF ICCV*, pages 6430–6439, 2019.

[30] Shenglai Zeng, Zonghang Li, Hongfang Yu, Yihong He, Zenglin Xu, Dusit Niyato, and Han Yu. Heterogeneous federated learning via grouped sequential-to-parallel training. In *International Conference on Database Systems for Advanced Applications*, pages 455–471. Springer, 2022.

[31] Lingjuan Lyu, Han Yu, and Qiang Yang. Threats to federated learning: A survey. *arXiv:2003.02133*, 2020.

[32] Lingjuan Lyu, Han Yu, Xingjun Ma, Chen Chen, Lichao Sun, Jun Zhao, Qiang Yang, and S Yu Philip. Privacy and robustness in federated learning: Attacks and defenses. *IEEE TNNLS*, 2022.

[33] Junbo Wang, Amitangshu Pal, Qinglin Yang, Krishna Kant, Kaiming Zhu, and Song Guo. Collaborative machine learning: Schemes, robustness, and privacy. *IEEE TNNLS*, pages 1–18, 2022.

[34] Luis Muñoz-González, Battista Biggio, Ambra Demontis, Andrea Paudice, Vasin Wongrassamee, Emil C Lupu, and Fabio Roli. Towards poisoning of deep learning algorithms with back-gradient optimization. In *Proceedings of the 10th ACM workshop on artificial intelligence and security*, pages 27–38, 2017.

[35] Vale Tolpegin, Stacey Truex, Mehmet Emre Gursoy, and Ling Liu. Data poisoning attacks against federated learning systems. In *European Symposium on Research in Computer Security*, pages 480–501. Springer, 2020.

[36] Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. Inverting gradients - how easy is it to break privacy in federated learning? *CoRR*, abs/2003.14053, 2020.

[37] Jinwoo Jeon, Jaechang Kim, Kangwook Lee, Sewoong Oh, and Jungseul Ok. Gradient inversion with generative image prior. *CoRR*, abs/2110.14962, 2021.

[38] Jie Xu, Benjamin S Glicksberg, Chang Su, Peter Walker, Jiang Bian, and Fei Wang. Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 5(1):1–19, 2021.

[39] Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zhihua Zhang. On the convergence of fedavg on non-iid data. *ICLR*, 2020.

[40] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Measuring the effects of non-identical data distribution for federated visual classification. *NeurIPS Workshop*, 2019.

[41] Shiqiang Wang, Tiffany Tuor, Theodoros Salonidis, Kin K Leung, Christian Makaya, Ting He, and Kevin Chan. Adaptive federated learning in resource constrained edge computing systems. *IEEE Journal on Selected Areas in Communications*, 37(6):1205–1221, 2019.

[42] Ahmed Khaled, Konstantin Mishchenko, and Peter Richtárik. First analysis of local gd on heterogeneous data. *arXiv:1909.04715*, 2019.

[43] Liang Gao, Huazhu Fu, Li Li, Yingwen Chen, Ming Xu, and Cheng-Zhong Xu. Feddc: Federated learning with non-iid data via local drift decoupling and correction. In *Proceedings of the IEEE/CVF CVPR*, pages 10112–10121, 2022.

[44] Matias Mendieta, Taojiannan Yang, Pu Wang, Minwoo Lee, Zhengming Ding, and Chen Chen. Local learning matters: Rethinking data heterogeneity in federated learning. In *Proceedings of the IEEE/CVF CVPR*, pages 8397–8406, June 2022.

[45] Yicheng Chen, Rick S Blum, and Brian M Sadler. Communication efficient federated learning via ordered admm in a fully decentralized setting. In *2022 56th Annual Conference on Information Sciences and Systems*, pages 96–100. IEEE, 2022.

[46] Yonghai Gong, Yichuan Li, and Nikolaos M Freris. Fedadmm: A robust federated deep learning framework with adaptivity to system heterogeneity. In *2022 IEEE 38th International Conference on Data Engineering*, pages 2575–2587. IEEE, 2022.

[47] Han Wang, Siddartha Marella, and James Anderson. Fedadmm: A federated primal-dual algorithm allowing partial participation. In *2022 IEEE 61st Conference on Decision and Control*, pages 287–294. IEEE, 2022.

[48] Yan Sun, Li Shen, Shixiang Chen, Liang Ding, and Dacheng Tao. Dynamic regularized sharpness aware minimization in federated learning: Approaching global consistency and smooth landscape. *arXiv:2305.11584*, 2023.

[49] Sashank Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečnỳ, Sanjiv Kumar, and H Brendan McMahan. Adaptive federated optimization. *ICLR*, 2021.

[50] Yujia Wang, Lu Lin, and Jinghui Chen. Communication-efficient adaptive federated learning. *arXiv:2205.02719*, 2022.

[51] Sai Praneeth Karimireddy, Martin Jaggi, Satyen Kale, Mehryar Mohri, Sashank J Reddi, Sebastian U Stich, and Ananda Theertha Suresh. Mime: Mimicking centralized stochastic algorithms in federated learning. *Advances in Neural Information Processing Systems*, 2021.

[52] Jianyu Wang, Vinayak Tantia, Nicolas Ballas, and Michael Rabbat. Slowmo: Improving communication-efficient distributed sgd with slow momentum. *ICLR*, 2020.

[53] Emre Ozfatura, Kerem Ozfatura, and Deniz Gündüz. Fedadc: Accelerated federated learning with drift control. In *IEEE International Symposium on Information Theory*, pages 467–472. IEEE, 2021.

[54] Jing Xu, Sen Wang, Liwei Wang, and Andrew Chi-Chih Yao. Fedcm: Federated learning with client-level momentum. *arXiv:2106.10874*, 2021.

[55] Rudrajit Das, Anish Acharya, Abolfazl Hashemi, Sujay Sanghavi, Inderjit S Dhillon, and Ufuk Topcu. Faster non-convex federated learning via global and local momentum. In *Uncertainty in Artificial Intelligence*, pages 496–506. PMLR, 2022.

[56] Geeho Kim, Jinkyu Kim, and Bohyung Han. Communication-efficient federated learning with acceleration of global momentum. *arXiv:2201.03172*, 2022.

[57] Daliang Li and Junpu Wang. Fedmd: Heterogenous federated learning via model distillation. *arXiv:1910.03581*, 2019.

[58] Wenke Huang, Mang Ye, and Bo Du. Learn from others and be yourself in heterogeneous federated learning. In *Proceedings of the IEEE/CVF CVPR*, pages 10143–10153, June 2022.

[59] Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.

[60] Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Personalized federated learning: A meta-learning approach. *NeurIPS*, 2020.

[61] Xiaochen Zhou and Xudong Wang. Memory and communication efficient federated kernel k -means. *IEEE TNNLS*, pages 1–12, 2022.

[62] Martin Papenberg and Gunnar W Klau. Using anticlustering to partition data sets into equivalent parts. *Psychological Methods*, 26(2):161, 2021.

[63] Mohammed Fayaz A., Neethimani S. M., Sai Lokesh Reddy Y., Srijayanthi Subramanian, and Sakthivel Ravichandran. Comparative analysis of anticlusters formed using various distance metrics and k-medoids algorithm. *International Journal of Advanced Science and Technology*, 29:7705–7717, 2020.

[64] Ventzeslav Valev. Set partition principles revisited. In Adnan Amin, Dov Dori, Pavel Pudil, and Herbert Freeman, editors, *Advances in Pattern Recognition*, pages 875–881, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg.

[65] Abhijit Guha Roy, Shayan Siddiqui, Sebastian Pölsterl, Nassir Navab, and Christian Wachinger. Braintorrent: A peer-to-peer environment for decentralized federated learning. *arXiv:1905.06731*, 2019.

[66] Chenghao Hu, Jingyan Jiang, and Zhi Wang. Decentralized federated learning: A segmented gossip approach. *arXiv:1908.07782*, 2019.

[67] Shreyansh Jain and Koteswar Rao Jerripothula. Federated learning for commercial image sources. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6534–6543, 2023.

[68] Cong Xie, Sanmi Koyejo, and Indranil Gupta. Asynchronous federated optimization. *arXiv:1903.03934*, 2019.

[69] Chenhao Xu, Youyang Qu, Yong Xiang, and Longxiang Gao. Asynchronous federated learning on heterogeneous devices: A survey. *arXiv:2109.04269*, 2021.

[70] Yang Chen, Xiaoyan Sun, and Yaochu Jin. Communication-efficient federated deep learning with layerwise asynchronous model update and temporally weighted aggregation. *IEEE TNNLS*, 31(10):4229–4238, 2019.

[71] Xiaofeng Lu, Yuying Liao, Pietro Lio, and Pan Hui. Privacy-preserving asynchronous federated learning mechanism for edge network computing. *IEEE Access*, 8:48970–48981, 2020.

[72] Jiangshan Hao, Yanchao Zhao, and Jiale Zhang. Time efficient federated learning with semi-asynchronous communication. In *IEEE 26th International Conference on Parallel and Distributed Systems*, pages 156–163. IEEE, 2020.

[73] Franziska Boenisch, Adam Dziedzic, Roei Schuster, Ali Shahin Shamsabadi, Ilia Shumailov, and Nicolas Papernot. When the curious abandon honesty: Federated learning is not private. In *IEEE 8th European Symposium on Security and Privacy*, pages 175–199. IEEE, 2023.

[74] Franziska Boenisch, Adam Dziedzic, Roei Schuster, Ali Shahin Shamsabadi, Ilia Shumailov, and Nicolas Papernot. Is federated learning a practical pet yet? *arXiv:2301.04017*, 2023.

[75] Briland Hitaj, Giuseppe Ateniese, and Fernando Perez-Cruz. Deep models under the gan: information leakage from collaborative deep learning. In *Proceedings of the 2017 ACM SIGSAC conference on computer and communications security*, pages 603–618, 2017.

[76] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

[77] Zhuohang Li, Jiaxin Zhang, Luyang Liu, and Jian Liu. Auditing privacy defenses in federated learning via generative gradient leakage. In *Proceedings of the IEEE/CVF CVPR*, pages 10132–10142, 2022.

[78] 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*, pages 691–706. IEEE, 2019.

[79] Yuxin Wen, Jonas Geiping, Liam Fowl, Micah Goldblum, and Tom Goldstein. Fishing for user data in large-batch federated learning via gradient magnification. *arXiv:2202.00580*, 2022.

[80] Cynthia Dwork. Differential privacy. In *International Colloquium on Automata, Languages, and Programming*, pages 1–12. Springer, 2006.
[81] Kang Wei, Jun Li, Ming Ding, Chuan Ma, Howard H Yang, Farhad Farokhi, Shi Jin, Tony QS Quek, and H Vincent Poor. Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security*, 15:3454–3469, 2020.
[82] Najeeb Moharram Jebreel, Josep Domingo-Ferrer, Alberto Blanco-Justicia, and David Sánchez. Enhanced security and privacy via fragmented federated learning. *IEEE TNNLS*, pages 1–15, 2022.
[83] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Federated visual classification with real-world data distribution. *CoRR*, abs/2003.08082, 2020.
[84] Michael R Garey and David S Johnson. "strong"np-completeness results: Motivation, examples, and implications. *Journal of the ACM*, 25(3):499–508, 1978.
[85] Mi Luo, Fei Chen, Dapeng Hu, Yifan Zhang, Jian Liang, and Jiashi Feng. No fear of heterogeneity: Classifier calibration for federated learning with non-iid data. *Advances in Neural Information Processing Systems*, 34:5972–5984, 2021.
[86] Haibo He and Edwardo A Garcia. Learning from imbalanced data. *IEEE Transactions on knowledge and data engineering*, 21(9):1263–1284, 2009.
[87] Douglas Steinley. K-means clustering: a half-century synthesis. *British Journal of Mathematical and Statistical Psychology*, 59(1):1–34, 2006.
[88] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
[89] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
[90] Sebastian Caldas, Sai Meher Karthik Duddu, Peter Wu, Tian Li, Jakub Konečnỳ, H Brendan McMahan, Virginia Smith, and Ameet Talwalkar. Leaf: A benchmark for federated settings. *Workshop on Federated Learning for Data Privacy and Confidentiality*, 2019.
[91] The TensorFlow Federated Authors. Tensorflow federated stack overflow dataset, 2019.
[92] Farshid Varno, Marzie Saghayi, Laya Rafiee, Sharut Gupta, Stan Matwin, and Mohammad Havaei. Minimizing client drift in federated learning via adaptive bias estimation. *arXiv:2204.13170*, 2022.
[93] Qinbin Li, Zeyi Wen, Zhaomin Wu, Sixu Hu, Naibo Wang, Yuan Li, Xu Liu, and Bingsheng He. A survey on federated learning systems: vision, hype and reality for data privacy and protection. *IEEE Transactions on Knowledge and Data Engineering*, 2021.
[94] Richard Bellman. Dynamic programming. *Science*, 153:34–37, 1966.
[95] Karl Pearson F.R.S. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901.
[96] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 2019.
[97] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63:139–144, 2020.
[98] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 2017.
[99] Xiaoling Xia, Cui Xu, and Bing Nan. Inception-v3 for flower classification. In *2nd international conference on image, vision and computing*, pages 783–787. IEEE, 2017.

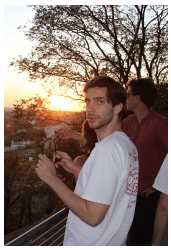

ANDREA SILVI received his Bachelor's in Computer Engineering from Politecnico di Torino in October 2019 and also completed there his Master's in Data Science and Engineering in December 2022. He joined the VANDAL (Visual and Applied Learning) Laboratory to work on Federated Learning during his Master's thesis, under the supervision of Prof. Barbara Caputo, Dr. Marco Ciccone and Debora Caldarola. In September 2023, Andrea started his Ph.D. at Chalmers University of Technology in Göteborg, Sweden, supervised by Professor Moa Johansson. His current focus lies in emergent communications in multi-agent systems and neuro-symbolic AI.

ANDREA RIZZARDI received his Master's degree in Data Science and Engineering at Polytechnic of Turin in 2022, and his Bachelor's degree in 2018. His background includes statistics and coding, while his main research areas are Federated Learning and Tiny ML.

DEBORA CALDAROLA is a 4th year PhD student from Polytechnic of Turin (Italy), supervised by Barbara Caputo and co-advised by Marco Ciccone. She is currently visiting Stanford University advised by Sanmi Koyejo. Her research focuses on trustworthy machine learning, with a specific interest in federated learning and fairness studied through the lens of the loss landscape (e.g., sharpness-aware training). She recently organized the Women in Computer Vision Workshop (WiCV) in conjunction with ICCV, and is a member of Eta Kappa Nu, the IEEE honor society. She received both her Master's (2020) and Bachelor's (2018) degrees in Computer Engineering from Polytechnic of Turin with maximum grade.

BARBARA CAPUTO received the Ph.D. degree in computer science from the KTH Royal Institute of Technology, Stockholm, Sweden, in 2005. From 2007 to 2013, she was a Senior Researcher at Idiap-EPFL. Then, she moved to Sapienza Rome University, thanks to a MUR professorship, and joined the Politecnico di Torino, in 2018. She is currently a Full Professor with the Politecnico di Torino, where she leads the Artificial Intelligence (AI) Hub, Politecnico di Torino. Since 2017, she has been a double affiliation with the Italian Institute of Technology (IIT). She is one of the 30 experts who contributed to write the Italian strategy on AI and coordinator of the Italian National Ph.D. on AI and industry 4.0, sponsored by MUR. She is an ERC Laureate, ELLIS Fellow, and since 2019, she serves on the ELLIS Board.

MARCO CICCONE is an ELLIS Postdoctoral Researcher in the VANDAL group at Polytechnic of Turin. He received a Ph.D. in Computer Science and Engineering cum laude at Polytechnic of Milan, working on iterative and conditional models for visual representation learning. His research interests are in the intersection of meta, continual, and federated learning to scale the training of agents with heterogeneous data and mitigate the effect of catastrophic forgetting and heterogeneity across tasks, domains, and devices.