

AFL: A Single-Round Analytic Approach for Federated Learning with Pre-trained Models

Run He¹, Kai Tong¹, Di Fang¹, Han Sun^{2,3}, Ziqian Zeng¹, Haoran Li⁴,
Tianyi Chen⁵, Huiping Zhuang^{1*}

¹ South China University of Technology ² Tsinghua University

³ Beijing National Research Center for Information Science and Technology

⁴ The Hong Kong University of Science and Technology ⁵ Microsoft

*corresponding: hpzhuang@scut.edu.cn

Abstract

In this paper, we introduce analytic federated learning (AFL), a new training paradigm that brings analytical (i.e., closed-form) solutions to the federated learning (FL) with pre-trained models. Our AFL draws inspiration from analytic learning—a gradient-free technique that trains neural networks with analytical solutions in one epoch. In the local client training stage, the AFL facilitates a one-epoch training, eliminating the necessity for multi-epoch updates. In the aggregation stage, we derive an absolute aggregation (AA) law. This AA law allows a single-round aggregation, reducing heavy communication overhead and achieving fast convergence by removing the need for multiple aggregation rounds. More importantly, the AFL exhibits a property that invariance to data partitioning, meaning that regardless of how the full dataset is distributed among clients, the aggregated result remains identical. This could spawn various potentials, such as data heterogeneity invariance and client-number invariance. We conduct experiments across various FL settings including extremely non-IID ones, and scenarios with a large number of clients (e.g., ≥ 1000). In all these settings, our AFL constantly performs competitively while existing FL techniques encounter various obstacles. Our codes are available at <https://github.com/ZHUANGHP/Analytic-federated-learning>.

1. Introduction

Federated learning (FL) [24] aims to collectively train a machine learning model over data silos by aggregating their individual trained weights, while preserving the privacy of their source data. This training paradigm has received high popularity, particularly in sensitive domains where data privacy is crucial, such as in banks [27, 41] and hospitals [11, 13].

Conventional FL techniques rely on weight aggregation

among clients over multiple rounds of training. The objective is to achieve convergence and approximate its joint-training counterpart, where all clients' data are accessible in a single location. To accomplish this, many contributions have been made. One widely recognized method is FedAvg [24]. Relying on a large number of aggregation rounds, the FedAvg employs a simple yet effective weight averaging technique across local clients. Building upon this, various methods have been proposed (e.g., the FedProx [17] and the FedNova [35]), each with its own specific focus within the field of FL.

However, training a model from scratch via FL can be computationally intensive and demanding in terms of communication bandwidth, especially with large models and numerous participating clients. Several efforts have explored utilizing pre-trained models to mitigate these challenges [23, 45]. Typically, this involves freezing the backbone and only updating and sharing lightweight parameters, such as prototypes [29, 30] or prompts [7, 40], to reduce the substantial training costs.

Although leveraging pre-trained models can circumvent the high costs associated with training backbones from scratch, existing FL techniques with pre-trained models are primarily based on a gradient-based iterative approach, necessitating iterative optimization on each client and multi-round aggregation across clients. The gradient-based optimization used in the existing FL faces various challenges and imposes several constraints. The faced challenges include, but are not limited to: 1) *Data heterogeneity*, where the data distribution in each client is not independently identical (non-IID), even with mutually exclusive data categories across different clients (i.e., pathological distribution), 2) *Large client number*, where the aggregation involving a significant number of clients (i.e., ≥ 1000) can lead to substantial performance degradation in FL systems as the client count increases [16], 3) *Slow convergence*: where FL methods may struggle to converge within limited communication rounds, especially

in severe non-IID scenarios, and 4) *High communication cost*, where multi-round aggregation in existing FL methods escalates the communication costs associated with parameter sharing between clients and servers.

In this paper, we propose a new FL training framework named analytic federated learning (AFL), which provides a single-round aggregation for federated learning with pre-trained models. The AFL draws inspiration from analytic learning [8, 32, 48]—a gradient-free technique with a closed-form solution obtained from reshaping the network training into linearized formulation. The AL paradigm receives several benefits over gradient-based techniques. First, it is gradient-free, thereby avoiding gradient-related issues, such as vanishing and exploding gradients. Second, the analytical solution frees AL from convergence issues during training. Also, the AL requires only one visit to the dataset while gradient-based mechanism usually needs hundreds of epochs or beyond. These properties are attractive in FL to accomplish fast convergence and low communication cost. Here, we are able to incorporate this mechanism into the FL domain to overcome limitations inherent in gradient-based techniques. Our contributions are summarized as follows:

- We propose the AFL, a gradient-free FL framework with analytical (closed-form) solutions. These analytical solutions apply both in the local client training stage and the aggregation stage.
- In the local stage, we adopt a pre-trained network to harness input embeddings, and formulate the training in each client into a localized linear regression problem. This leads to a least squares (LS) based *one-epoch client training*, eliminating the need for multi-epoch training and enabling fast convergence in local clients.
- In the aggregation stage, we derive an absolute aggregation (AA) law in analytical form, optimally establishing a *single-round aggregation*. That is, the aggregation happens only once, avoiding multiple FL rounds that bring high communication costs. Additionally, in scenarios where the AA law becomes suboptimal due to a large number of clients, we introduce a regularization intermediary (RI) process to restore its optimality.
- Owing to analytical solutions, the AFL exhibits a property that *invariance to data partitioning*. This means that regardless of how the full dataset is distributed (e.g., non-IID) among local clients, the result remains identical. This property spawns several appealing characteristics: i) *Data heterogeneity invariance* where the result is invariant to arbitrary heterogeneous data partition scenarios. ii) *Client-number invariance*, which produces identical results regardless of the number of clients involved.
- We conduct extensive experiments spanning diverse scenarios, including a wide variety of non-IID partitions and large client number (up to 1000) settings. Our AFL consistently showcases competitive performance throughout all

these settings when compared with other methods.

2. Related Works

In this section, we review existing related FL literature. Additionally, we explore various AL techniques and their variants to reveal their underlying mechanisms.

2.1. Federated Learning Methods

Following the FedAvg [24], to address non-IID issues in FL, various methods have been proposed. One common approach involves assessing the significance of parameters during aggregation to ensure that local updates do not diverge substantially from the global model. For instance, the FedProx [17] restricts the size of local updates, while the FedNova [35] employs a normalized averaging method to eliminate target inconsistency while maintaining fast error convergence. These methods are frequently used as baselines, and we compare our results against them in our experiments. Another set of methods focuses on determining adaptive aggregation weights obtained from multiple clients. The FedLAW [19] learns these weights to achieve a global model with state-of-the-art performance, though it requires a proxy dataset to learn the weights, making the results sensitive to the selection of the proxy dataset. To address this sensitivity, the FedCDA [34] proposes a proxy-free method that reduces each client’s deviation from the local models of other participants and selects a local model from its multiple recent models acquired over several rounds.

Some methods address the parameter order mismatch issue across clients, which can occur during global aggregation. The Fed2 [43] designs a model structure adaptation method to ensure explicit feature allocation across different network structures. Similarly, the method in [18] seeks a position-aware neuron to fuse position-related values (i.e., position encodings) into neuron outputs. Distillation methods [2, 9, 33, 38] represent another branch, where the average of logits from client models is used for the local model aggregation, thereby enhancing generalization. [21] pioneers to apply knowledge distillation on the server side, transferring knowledge from multiple local models to the global model using an unlabeled proxy dataset. To overcome the limitation of using a proxy dataset, recent studies such as [47] and [44] suggest substituting the proxy dataset with generated data.

Existing FL techniques have several significant drawbacks, including challenges with data heterogeneity, large client numbers, convergence issues and high communication costs. Our AFL framework addresses these issues by utilizing a gradient-free, closed-form analytic learning approach, avoiding gradient-related problems (e.g., multi-epoch training, convergence issues and multi-round aggregation).

2.2. Analytic Learning

The AL has been developed as a strategy to address issues associated with gradient-based update, such as gradient vanishing/exploding, divergence during iteration, and long training time due to multi-epoch training. The AL is also referred to as pseudoinverse learning [8] owing to its utilization of matrix inversion. The AL starts from shallow learning, which is investigated prior to the advent of deep networks in the realm of research. For instance, the radial basis network [26] trains parameters using an LS estimation after performing a kernel transformation in the first layer. The multilayer AL [31, 37] comes up with a one-epoch training style, using LS techniques to resolve linear segments transformed by nonlinear network. One instance of this method is the dense pseudoinverse autoencoder [36], which uses LS solutions to combine shallow and deep features to train a stacked autoencoder layer-by-layer.

Nonetheless, earlier AL techniques train their weights by processing the entire dataset simultaneously, therefore facing memory challenge. This memory concern is alleviated by the block-wise recursive Moore-Penrose inverse [48], which equivalently replaces the joint learning with a recursive approach. This recursive equivalent property echoes well with the continual learning community. Naturally, analytic continual learning techniques [50, 51, 53] adopt this equivalent characteristic, thrive in handling the catastrophic forgetting problem, and are invariant to the sequential data partition in continual learning. Our AFL draws inspiration from these adaptations, aiming to introduce similar equivalent patterns (e.g., invariant to heterogeneous data) to the FL community.

3. Analytic Federated Learning

In this section, we provide a detailed exposition of AFL derivations, organized into a local training stage and a centralized aggregation stage. In the local stage, a pre-trained backbone serves as a feature extractor, facilitating an AL network learning that allows the training to be completed in one epoch. In the aggregation stage, we introduce the AA law, establishing a single-round aggregation. We elaborate on AFL’s invariance to data partitioning here, bringing benefits such as data heterogeneity invariance, client-number invariance and fast convergence in a single round. An Overview of the proposed AFL paradigm is depicted in Figure 1.

Prior to further developments, here let $\mathcal{D} = \{\mathcal{D}_k\}_{k=1}^K$ be the complete training data, where $\mathcal{D}_k \sim \{\mathcal{X}_{k,i}, y_{k,i}\}_{i=1}^{N_k}$ suggests an N_k -sample sub-dataset accessible to the k -th client with $\mathcal{X}_{k,i}$ and $y_{k,i}$ representing the i -th input-label pair. In this paper, all these K clients share the same backbone network f_{backbone} parameterized by Θ to map their inputs (e.g., \mathcal{X}) to embedding vectors.

3.1. Local Stage: Localized Analytic Learning

In this stage, each local client’s network is trained using the AL technique. This involves transforming the neural network’s classification head into a linear regression problem, thereby enabling the derivation of a closed-form LS solution.

At the initial step, client k extracts its embedding vector $\mathbf{x}_{k,i}$ by passing the i -th data $\mathcal{X}_{k,i}$ from \mathcal{D}_k through the frozen backbone network f_{backbone} , i.e.,

$$\mathbf{x}_{k,j} = f_{\text{backbone}}(\mathcal{X}_{k,j}, \Theta) \quad (1)$$

where $\mathbf{x}_{k,j} \in \mathbb{R}^{1 \times y_e}$, with y_e indicating the embedding length.

For the k -th client (with N_k samples in \mathcal{D}_k), we can stack the extracted embeddings and their corresponding one-hot labels via mapping $\mathcal{D}_k \sim \{\mathcal{X}_{k,i}, y_{k,i}\}_{i=1}^{N_k}$ to $\mathcal{D}_k \sim \{\mathbf{X}_k, \mathbf{Y}_k\}$, i.e.,

$$\mathbf{X}_k = \begin{bmatrix} \mathbf{x}_{k,1} \\ \mathbf{x}_{k,2} \\ \vdots \\ \mathbf{x}_{k,N_k} \end{bmatrix} = \begin{bmatrix} f_{\text{backbone}}(\mathcal{X}_{k,1}, \Theta) \\ f_{\text{backbone}}(\mathcal{X}_{k,2}, \Theta) \\ \vdots \\ f_{\text{backbone}}(\mathcal{X}_{k,N_k}, \Theta) \end{bmatrix} \quad \mathbf{Y}_k = \begin{bmatrix} \text{onehot}(y_{k,1}) \\ \text{onehot}(y_{k,2}) \\ \vdots \\ \text{onehot}(y_{k,N_k}) \end{bmatrix}, \quad (2)$$

where the embedding matrix $\mathbf{X}_k \in \mathbb{R}^{N_k \times y_e}$, and the label matrix $\mathbf{Y}_k \in \mathbb{R}^{N_k \times C}$ has C classes. The $\text{onehot}(\ast)$ operator converts the index label $y_{k,j}$ into a C -dimension one-hot row vector.

Subsequently, we approach the local client training with AL technique [8]. Specially, the target of the k -th client is to linearly map the extracted embeddings onto the one-hot labels by minimizing the mean square error (MSE) loss function as follows.

$$\mathcal{L}(\mathbf{W}_k) = \|\mathbf{Y}_k - \mathbf{X}_k \mathbf{W}_k\|_{\text{F}}^2, \quad (3)$$

where $\|\ast\|_{\text{F}}$ indicates the Frobenius norm. This leads to an optimal weight estimation $\hat{\mathbf{W}}_k$, i.e.,

$$\hat{\mathbf{W}}_k = \underset{\mathbf{W}_k}{\text{argmin}} \mathcal{L}(\mathbf{W}_k) = \mathbf{X}_k^\dagger \mathbf{Y}_k, \quad (4)$$

where \dagger denotes the Moore-Penrose (MP) inverse (also referred as generalized inverse or pseudoinverse) [8, 48].

The solution presented in (4) optimally addresses the MSE loss function described in (3), effectively establishing an LS-based AL solution for localized network learning.

Why One-epoch Analytic Learning Works. AL methods are generally effective for training shallow networks but face challenges when applied to deeper ones. This can be attributed to the fact that AL techniques are often designed as classifiers rather than end-to-end learning approaches. Despite this limitation, recent research has demonstrated that with a well-trained backbone, the AL performs adequately in various complex scenarios [52]. The practice of using a “pre-trained backbone + downstream tasks” has become increasingly common. This has allowed the one-epoch AL




Figure 1. An overview of the AFL. During the local stage, each client calculates C_k^r and \hat{W}_k^r based on the same pre-trained backbone and its own dataset. The server obtained the $C_{agg,K}^r$ and $\hat{W}_{agg,K}^r$ then get \hat{W} in the aggregation stage.

to thrive in various areas such as continual learning [50] and reinforcement learning [22]. Hence, it could also be well incorporated in the individual client training.

Adopting AL is the key to enforcing the upcoming single-round aggregation (by deriving the AA law). The affine characteristic of linear regression in each client opens up new possibilities for exploration in FL. We provide a comprehensive explanation of such an exploration in later sections.

3.2. Aggregation Stage: Absolute Aggregation Law

In the aggregation stage, we introduce the Absolute Aggregation (AA) law, a key contribution in AFL. The AA law facilitates a single-round aggregation, i.e., the aggregation happens only once. Additionally, in scenarios where the AA law becomes suboptimal due to a large number of clients, we introduce a regularization intermediary (RI) process to restore its optimality.

The MP inverse partition [4] inspires our derivation, which is reformulated into Lemma 1.

Lemma 1. Let $X = \begin{bmatrix} X_u \\ X_v \end{bmatrix}$ with X_u and X_v having full column ranks, and X follows a partition

$$X^\dagger = [\bar{U} \quad \bar{V}], \quad (5)$$

where

$$\begin{cases} \bar{U} = X_u^\dagger - R_u C_v X_u^\dagger - R_u C_v (C_u + C_v)^{-1} C_v X_u^\dagger \\ \bar{V} = X_v^\dagger - R_v C_u X_v^\dagger - R_v C_u (C_u + C_v)^{-1} C_u X_v^\dagger \end{cases}, \quad (6)$$

$$\begin{cases} C_u = X_u^\top X_u \\ C_v = X_v^\top X_v \end{cases}, \quad \begin{cases} R_u = C_u^{-1} \\ R_v = C_v^{-1} \end{cases}.$$

Proof. See Supplementary Materials A. \square

Lemma 1 points out that, a matrix's MP inverse (e.g., X^\dagger) can be computed using the inverse matrices of its block

components (e.g., X_u^\dagger and X_v^\dagger). This introduces possibilities for aggregating a weight $W = X^\dagger Y$ equally from manipulating constituent counterparts $W_u = X_u^\dagger Y_u$ and $W_v = X_v^\dagger Y_v$. That is, $W = f_{agg}(W_u, W_v)$, i.e., a single-aggregation strategy.

Bearing the above intuition in mind, we are able to derive such a single-aggregation strategy in action. This is delivered in Theorem 1.

Theorem 1. Absolute Aggregation Law: Let $\hat{W} = X^\dagger Y$, where $X = \begin{bmatrix} X_u \\ X_v \end{bmatrix}$ and $Y = \begin{bmatrix} Y_u \\ Y_v \end{bmatrix}$ with X_u and X_v having full column ranks. Let $\hat{W}_u = X_u^\dagger Y_u$, $\hat{W}_v = X_v^\dagger Y_v$, and we have

$$W = W_u W_u + W_u W_v, \quad (7)$$

where

$$\begin{cases} W_u = I - R_u C_v - R_u C_v (C_u + C_v)^{-1} C_v \\ W_v = I - R_v C_u - R_v C_u (C_u + C_v)^{-1} C_u \end{cases}$$

$$\begin{cases} C_u = X_u^\top X_u \\ C_v = X_v^\top X_v \end{cases} \quad \begin{cases} R_u = C_u^{-1} \\ R_v = C_v^{-1} \end{cases}. \quad (8)$$

Proof. See Supplementary Materials B. \square

The AA law, as stated in Theorem 1, provides a powerful insight. It establishes an intuition that we can aggregate two independently trained weights, such as W_u and W_v , into their jointly trained counterpart W . This is achieved in an optimal way without any approximation or parameter tuning.

Invariance to data partitioning. To a certain extent, the achievement in Theorem 1 attains the ultimate goal of FL, i.e., the *equivalence* between weights trained in FL fashion and that trained on a centralized joint dataset. Traditionally, the FL aims to approximate or converge to the performance of the joint-trained model through multiple rounds of aggregation in a central server. However, the AA law provides a

more direct path to this goal. It allows for an equivalence (not approximation or convergence) to manifest in a linear regression standpoint.

Supported by the AA law, the AFL achieves a level of performance that is on par with the joint-trained model, without the need for multiple rounds of aggregation. This direct equivalence could establish significant advancement in FL, as it simplifies the process and reduces the heavy computational overhead associated with multiple aggregation rounds.

Although the AA law in Theorem 1 admits the absolute aggregation between two clients (i.e., $\hat{\mathbf{W}}_u$ and $\hat{\mathbf{W}}_v$), this pattern can be trivially broadcast to multi-client scenario. To elaborate, without loss of generality, we denote $\hat{\mathbf{W}}_{\text{agg},k-1}$ as the accumulated aggregation (AcAg) weight that has aggregated $k-1$ clients. By rewriting (1), the next aggregation with $\hat{\mathbf{W}}_k$ ($i = k, \dots, K$) reads

$$\hat{\mathbf{W}}_{\text{agg},k} = \mathcal{W}_{\text{agg}} \hat{\mathbf{W}}_{\text{agg},k-1} + \mathcal{W}_k \hat{\mathbf{W}}_k. \quad (9)$$

According to (1), let $\mathbf{C}_u \rightarrow \mathbf{C}_{\text{agg},k-1}$, $\mathbf{C}_v \rightarrow \mathbf{C}_k$, and we have $\mathbf{C}_{\text{agg},k} = \mathbf{C}_{\text{agg},k-1} + \mathbf{C}_k$. Hence,

$$\begin{cases} \mathcal{W}_{\text{agg}} = \mathbf{I} - \mathbf{C}_{\text{agg},k-1}^{-1} \mathbf{C}_k (\mathbf{I} + \mathbf{C}_{\text{agg},k}^{-1} \mathbf{C}_k), \\ \mathcal{W}_k = \mathbf{I} - \mathbf{C}_k^{-1} \mathbf{C}_{\text{agg},k-1} (\mathbf{I} + \mathbf{C}_{\text{agg},k}^{-1} \mathbf{C}_{\text{agg},k-1}), \end{cases} \quad (10)$$

where

$$\begin{cases} \mathbf{C}_{\text{agg},k} = \mathbf{C}_{\text{agg},k-1} + \mathbf{C}_k = \sum_i^k \mathbf{C}_i, \\ \mathbf{C}_i = \mathbf{X}_i^\top \mathbf{X}_i. \end{cases} \quad (11)$$

As such, the joint-trained weight $\hat{\mathbf{W}} = \hat{\mathbf{W}}_{\text{agg},k}$ is produced by aggregating among individual clients in a pair-wise manner. It is interesting to find that the optimal aggregation is in fact a linear combination between two matrices (e.g., $\hat{\mathbf{W}}_{\text{agg},k-1}$ and $\hat{\mathbf{W}}_k$) weighted by \mathcal{W}_{agg} and \mathcal{W}_k respectively.

Note that the aggregation does NOT necessarily follow a sequential index from 1 to K . We can randomly sample an available client to aggregate with the AcAg weight. This is revealed by the fact that elements in the weighting matrices are somewhat interchangeable (e.g., see (10)).

3.3. RI Process: AA Law in Rank-deficient Scenario

As indicated in Theorem 1, the equivalence in AA law relies on an assumption of a full-column rank in each client, e.g., \mathbf{X}_k having full-column rank. This may not hold in the large client number scenario where each client has limited data (e.g., $N_k < y_e$), rendering the full-column rank assumption invalid. To address this, we implement the AA law with an RI process. Specially, we include a regularization term as an intermediary during the local stage, and remove it after the aggregation stage.

To this end, we include a regularization term controlled by γ in the objective function, i.e.,

$$\mathcal{L}(\mathbf{W}_k^r) = \|\mathbf{Y}_k - \mathbf{X}_k \mathbf{W}_k^r\|_{\mathbb{F}}^2 + \gamma \|\mathbf{W}_k^r\|_{\mathbb{F}}^2, \quad (12)$$

which rewrites the MP inverse based solution in (4) into

$$\hat{\mathbf{W}}_k^r = \underset{\mathbf{W}_k^r}{\text{argmin}} \mathcal{L}(\mathbf{W}_k^r) = (\mathbf{X}_k^\top \mathbf{X}_k + \gamma \mathbf{I})^{-1} \mathbf{X}_k^\top \mathbf{Y}_k. \quad (13)$$

Such a solution does not suffer from rank-deficiency issues, as $\mathbf{X}_k^\top \mathbf{X}_k + \gamma \mathbf{I}$ is positive-definite thereby a full-rank matrix.

During aggregation, we substitute $\hat{\mathbf{W}}_k$ in (4) with $\hat{\mathbf{W}}_k^r$ using (13). This substitution would clearly result in deviations (i.e., $\hat{\mathbf{W}}_{\text{agg},k}^r \neq \hat{\mathbf{W}}_{\text{agg},k}$), which is depicted in Theorem 2.

Theorem 2. RI-AA Law: *The relation between $\hat{\mathbf{W}}_{\text{agg},k}^r$ and $\hat{\mathbf{W}}_{\text{agg},k}$ follows*

$$\hat{\mathbf{W}}_{\text{agg},k}^r = (\mathbf{C}_{\text{agg},k}^r)^{-1} \mathbf{C}_{\text{agg},k} \hat{\mathbf{W}}_{\text{agg},k}, \quad (14)$$

where

$$\mathbf{C}_{\text{agg},k}^r = \mathbf{C}_{\text{agg},k} + k\gamma \mathbf{I} = \sum_i^k \mathbf{C}_i^r, \quad \mathbf{C}_i^r = \mathbf{X}_i^\top \mathbf{X}_i + \gamma \mathbf{I}. \quad (15)$$

Proof. See Supplementary Materials C. \square

Theorem 2 establishes the relation between $\hat{\mathbf{W}}_{\text{agg},k}^r$ and $\hat{\mathbf{W}}_{\text{agg},k}$, which is a one-to-one mapping, such that $\hat{\mathbf{W}}_{\text{agg},k}$ can be restored by manipulating $\hat{\mathbf{W}}_{\text{agg},k}^r$, i.e.,

$$\begin{aligned} \hat{\mathbf{W}}_{\text{agg},k} &= (\mathbf{C}_{\text{agg},k})^{-1} \mathbf{C}_{\text{agg},k}^r \hat{\mathbf{W}}_{\text{agg},k}^r \\ &= (\mathbf{C}_{\text{agg},k}^r - k\gamma \mathbf{I})^{-1} \mathbf{C}_{\text{agg},k}^r \hat{\mathbf{W}}_{\text{agg},k}^r. \end{aligned} \quad (16)$$

That is, we are able to attain $\hat{\mathbf{W}}_{\text{agg},k}$ by removing the impact of the regularization term γ to counter the ill-conditioned constraint in the large client number scenario. The implementation of AFL is summarized in Algorithm 1.

Benefits of Adopting AL in AFL. Inheriting from the AL technique, the AFL admits several merits over its gradient-based counterparts as follows. i) *Fast training and convergence:* the analytical solutions allow AFL to finish the training and aggregation in one shot, exhibiting fast training and convergence. Also, the analytical solutions free the AFL from any convergence issue as no iterative-search based action is executed. ii) *Low communication cost:* the single-round aggregation only requires a single communication between the clients and the server, which significantly reduces the communication cost. iii) *Data heterogeneity invariance:* the invariance to data partitioning does not pose any constraint on data partition strategy. That is, the equivalence is hold across all possible data heterogeneous scenarios (e.g., see Section 4.2). iv) *Client-number invariance:* for a complete dataset \mathcal{D} partitioned among K clients (i.e., $\{\mathcal{D}_k\}_{k=1}^K$), according to Theorem 1 and (9), when the weights from all K clients are aggregated, the resulting weight is identical to that trained on the full dataset \mathcal{D} . To validate AA law with RI-process, we conduct an experiment on a dummy dataset and show the invariance (see Supplementary Materials D).

Algorithm 1 Analytic Federated Learning

Input: $\mathcal{D}_k, k = 0, \dots, K, \gamma$, and pre-trained backbone Θ .

Server Executes:

1. **for** each client k **in parallel do**
2. $\hat{\mathbf{W}}_k^r, \mathbf{C}_k^r \leftarrow \text{Local Stage}(k, \mathcal{D}_k, \gamma)$.
3. **end for**
4. $\hat{\mathbf{W}} \leftarrow \text{Aggregation Stage}(\{\hat{\mathbf{W}}_k^r, \mathbf{C}_k^r, \gamma\}_{k=1}^K)$.

Local Stage: client k with \mathcal{D}_k and γ .

1. Get embedding and label matrices using (2).
2. Obtain weight matrix $\hat{\mathbf{W}}_k^r$ by (13).
3. Get $\mathbf{C}_k^r = \mathbf{X}_k^\top \mathbf{X}_k + \gamma \mathbf{I}$.
4. Return $\hat{\mathbf{W}}_k^r, \mathbf{C}_k^r$.

Aggregation Stage: with $\{\hat{\mathbf{W}}_k^r, \mathbf{C}_k^r, \gamma\}_{k=1}^K$.

1. Initialize $\hat{\mathbf{W}}_{\text{agg},0}^r = \mathbf{0}, \mathbf{C}_{\text{agg},0}^r = \mathbf{0}$.
 2. **for** k in $\text{range}(K)$:
 - i) Aggregate $\hat{\mathbf{W}}_{\text{agg},k}^r$ with $\hat{\mathbf{W}}_k^r$ using (9).
 - ii) Update $\mathbf{C}_{\text{agg},k}^r = \mathbf{C}_{\text{agg},k-1}^r + \mathbf{C}_k^r$.
 3. **end for**.
 4. Restore $\hat{\mathbf{W}} = \hat{\mathbf{W}}_{\text{agg},K}^r$ with $\hat{\mathbf{W}}_{\text{agg},K}^r$ in (16).
-

An AL Branch of Federated Learning. The AFL incorporates the AL technique and can be considered as an AL branch within the FL context. The AL and its recursive formulation have demonstrated remarkable adaptability in continual learning utilizing a well-trained backbone [52, 54]. In this case, this intuition has been extended to the FL field through non-trivial derivations.

4. Experiments

In this section, we provide extensive experiments to validate the proposed AFL, including comparison with FL state-of-the-arts and analysis under various settings. The training time and ablation study of regularization are also investigated.

4.1. Comparison with FL Techniques

We conduct comparison with FL state-of-the-arts, including FedAvg [24], FedProx [17], MOON [15], FedGen [46], FedDyn [1], FedNTD [14] and FedDisco [42] under various non-IID settings.

Dataset and Model. We validate the baselines and our proposed AFL in 3 popular benchmark datasets in FL: CIFAR-10 [12], CIFAR-100 [12] and Tiny-ImageNet [25]. For all datasets, we use a ResNet-18 [10] pretrained on ImageNet-1k [5] as backbone. We freeze the backbones in all FL methods.

Data Partition. For simulating Non-IID scenarios in FL, we specify two Non-IID data partition methods including Latent Dirichlet Allocation [20] (LDA, denoted as NIID-1) and Sharding [20] (denoted as NIID-2). In the LDA setting, data assigned to each clients is forced to satisfy the Dirichlet distribution, and degree of the data heterogeneity is controlled

by parameter α . Smaller α leads to a more heterogeneous data distribution. In the Sharding strategy, the data is sorted by labels and divided into same-sized shards, and s control the heterogeneity, i.e. the number of shards per client. When s takes a smaller value, the data is more heterogeneous. We choose $\alpha = 0.1, 0.01$ and $s = 10, 4$ for CIFAR-100 and Tiny-ImageNet. For CIFAR-10, α is set to 0.1 and 0.01, and s is set to 4 and 2. Most existing methods are validated on data partition of $\alpha = 0.3$ to 1.0 and $s = 10$ [14, 19]. Here we provide more challenging settings to validate the robustness under extremely heterogeneous cases.

Implementation Details. In all the experiments, we use 100 clients for each method and use the same partitioned dataset within experiments of the same data setting. We implement the AFL with a $\gamma = 1$ RI process (any γ would suffice, see the ablation study). Each experiment setting is run 3 times and the mean and standard deviation of the best top-1 classification accuracy during training are reported. The implementation details of gradient-based compared methods can be found in Supplementary Materials E.

Experimental Results. We report the results of the compared methods under the setting of NIID-1 and NIID-2 in Table 1. As shown in the table, except for slightly weaker results than those of FedDyn in the NIID-2 setting, the AFL obtains very competitive performance compared with other methods across various settings. The degree of data heterogeneity does not at all affect the AFL. For instance, the accuracy remains 80.75% on CIFAR-10 for various NIID-1 and NIID-2 settings. Although slight differences could occur among various settings, it barely impacts the classification accuracy (an AL property indicated in [48]). The same pattern repeats on CIFAR-100 and Tiny-ImageNet. Uniquely, the AFL obtains identical results for all 3 repeated runs, i.e., the standard deviations are zeros! This is because the AFL does not introduce any stochastic element, so the repeated computation in each run is naturally equivalent to one another, hence the zero standard deviation.

Notably, when introducing a pre-trained backbone, the compared methods yield similar results and even FedAvg can become a competitive baseline. The reason could be that, when incorporating a pre-trained backbone, the FL process can be more stable with a better start point and methods stabilizing the FL process could be less effective. This phenomena has also been witnessed in other FL studies [3, 7]. However, other FL methods still experience performance reductions under severe non-IID scenarios. For example, FedDyn performs relatively well (e.g., 57.55%) under NIID-1 ($\alpha = 0.1$) on CIFAR-100 but undergoes a performance degradation (to 36.12%) when $\alpha = 0.01$. This pattern is rather consistent in other compared methods, such as FedAvg (56.62% \rightarrow 32.99%), FedProx (56.45% \rightarrow 33.37%), and MOON (56.58% \rightarrow 33.34%), and is also true across all datasets. The performance distributions regarding NIID-2 for

Table 1. The top-1 accuracy (%) of compared methods under two non-IID settings. Settings controlled by α and s are NIID-1 and NIID-2 respectively. The data is reported as average and standard deviation after 3 runs. Results in **bold** are the best within the compared methods in the same setting.

Dataset	Setting	FedAvg	FedProx	MOON	FedGen	FedDyn	FedNTD	FedDisco	AFL
CIFAR-10	$\alpha = 0.1$	64.02 \pm 0.18	64.07 \pm 0.08	63.84 \pm 0.03	64.14 \pm 0.24	64.77 \pm 0.11	64.64 \pm 0.02	63.83 \pm 0.08	80.75 \pm 0.00
	$\alpha = 0.05$	60.52 \pm 0.39	60.39 \pm 0.09	60.28 \pm 0.17	60.65 \pm 0.19	60.35 \pm 0.54	61.16 \pm 0.33	59.90 \pm 0.05	80.75 \pm 0.00
	$s = 4$	68.47 \pm 0.13	68.46 \pm 0.08	68.47 \pm 0.15	68.24 \pm 0.28	73.50 \pm 0.11	70.24 \pm 0.11	65.04 \pm 0.11	80.75 \pm 0.00
	$s = 2$	57.81 \pm 0.03	57.61 \pm 0.12	57.72 \pm 0.15	57.02 \pm 0.18	64.07 \pm 0.09	58.77 \pm 0.18	58.78 \pm 0.02	80.75 \pm 0.00
CIFAR-100	$\alpha = 0.1$	56.62 \pm 0.12	56.45 \pm 0.22	56.58 \pm 0.02	56.48 \pm 0.17	57.55 \pm 0.08	56.60 \pm 0.14	55.79 \pm 0.04	58.56 \pm 0.00
	$\alpha = 0.01$	32.99 \pm 0.20	33.37 \pm 0.09	33.34 \pm 0.11	33.09 \pm 0.09	36.12 \pm 0.08	32.59 \pm 0.21	25.72 \pm 0.08	58.56 \pm 0.00
	$s = 10$	55.76 \pm 0.13	55.80 \pm 0.16	55.70 \pm 0.25	60.93 \pm 0.17	61.09 \pm 0.09	54.69 \pm 0.15	54.65 \pm 0.09	58.56 \pm 0.00
	$s = 5$	48.33 \pm 0.15	48.29 \pm 0.14	48.34 \pm 0.19	48.12 \pm 0.06	59.34 \pm 0.11	47.00 \pm 0.19	45.86 \pm 0.18	58.56 \pm 0.00
Tiny-ImageNet	$\alpha = 0.1$	46.04 \pm 0.27	46.47 \pm 0.23	46.21 \pm 0.14	46.27 \pm 0.14	47.72 \pm 0.22	46.17 \pm 0.16	47.48 \pm 0.06	54.67 \pm 0.00
	$\alpha = 0.01$	32.63 \pm 0.19	32.26 \pm 0.14	32.38 \pm 0.20	32.33 \pm 0.14	35.19 \pm 0.06	31.86 \pm 0.44	27.15 \pm 0.10	54.67 \pm 0.00
	$s = 10$	39.06 \pm 0.26	38.97 \pm 0.23	38.79 \pm 0.14	38.82 \pm 0.16	41.36 \pm 0.06	37.55 \pm 0.09	38.86 \pm 0.12	54.67 \pm 0.00
	$s = 5$	29.66 \pm 0.19	29.17 \pm 0.16	29.24 \pm 0.30	29.37 \pm 0.25	35.18 \pm 0.18	29.01 \pm 0.14	27.72 \pm 0.18	54.67 \pm 0.00

these compared methods resemble those in NIID-1, where smaller s values invite performance degradation among existing FL counterparts. For instance, the FedDyn exhibits 73.50% \rightarrow 64.07% for $s = 4 \rightarrow 2$ on CIFAR-10 while the AFL obtains competitive and identical results (e.g., 80.75%).

4.2. Analysis on Data Partition

Here we provide broaden non-IID partitions to demonstrate AFL’s invariance to data partitioning. This includes varying the client number and the non-IID degree. We also provide the IID partition results.

Client-number Invariance. We compare our AFL and the FedAvg under NIID-1 setting on CIFAR-100 and Tiny-ImageNet with $\alpha = 0.1$, and vary the number of clients from 100 to 500 and 1000. The results are shown in Figure 2. We observe that the AFL keeps an identical performance when scaling the number of clients, while the FedAvg experiences a performance decline along the increasing number (e.g., 56.57% \rightarrow 41.01% for $K = 100 \rightarrow 1000$ on CIFAR-100). This provides a strong evidence to support the invariance to data partitioning in our AFL. It also showcases the capability of pushing the AFL to large-scale client training scenario without any performance compromise.

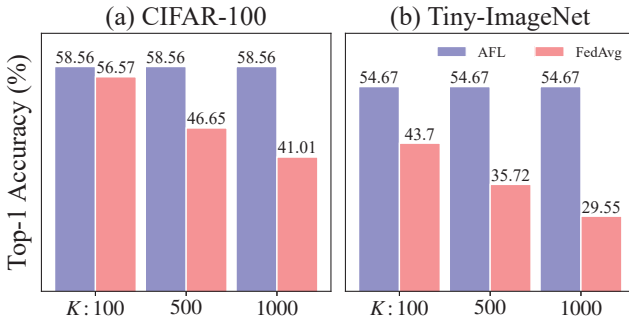


Figure 2. Accuracy over various number of clients.

Data Heterogeneity Invariance. Here, we fix the client number to 100 and partition the CIFAR-100 under the setting of NIID-1 with $\alpha = 0.005, 0.01, 0.1, 1$, including the IID setting as well. We report the results of AFL and FedAvg in Table 2. The FedAvg suffers from more accuracy losses (e.g., 57.72% \rightarrow 24.74% for $\alpha = 0.1 \rightarrow 0.005$) as the data heterogeneity grows higher. Under the IID partition, the FedAvg receives its best performance (i.e., 57.89%), which is still less competitive than our AFL (i.e., 58.56%). On the other hand, AFL obtains identical results (i.e., 58.56%) across various settings, including non-IID and IID ones. This is another strong proof of the AA law indicating the weight-invariant property of AFL. Our AFL is invariant to any degree of data heterogeneity, leading to unchanged performance in all possible data heterogeneous partition scenarios, even in extreme data heterogeneous cases (e.g., $\alpha = 0.005$).

Table 2. The top-1 classification accuracy (%) of AFL and FedAvg under different data heterogeneity.

Acc. (%)	$\alpha = 0.005$	$\alpha = 0.01$	$\alpha = 0.1$	$\alpha = 1$	IID
FedAvg	24.74	33.09	56.57	57.72	57.89
AFL	58.56	58.56	58.56	58.56	58.56

4.3. Training Efficiency

Fast Training with Single-round Aggregation. We plot the training evolution curves of accuracy on CIFAR-100 and Tiny-ImageNet in Figure 3 and report the execution time for each method in the legend bars. Compared FL methods take 60s to 100s on CIFAR-100 (100s to 160s on Tiny-ImageNet) to complete an aggregation round, leading to a total training time of 30,000s to 50,000 (50,000s to 80,000s). AFL, however, spends 236.61s on CIFAR-100 and 349.50s on Tiny-ImageNet, achieving approximately 150 \times -200 \times speedups over its FL counterparts due to only one aggregation.




Figure 3. Accuracy curves with communication rounds. Average training time is reported in the legends.

4.4. Ablation Study of RI Process

Here, we conduct an ablation study regarding the RI process by reporting accuracies of AFL with $\alpha = 0.1$ and $K = 100, 500, 1000$ under different values of γ . The results without and with the RI process are provided in Table 3. When $\gamma = 0$ (i.e., no regularization involved), the AFL stops working with $K = 500$ and 1000 due to the ill-conditioned matrix scenario (e.g., $N_k < y_c$). Such an ill-conditioned case is avoided by introducing γ . However, the lack of RI process (see left columns in Table 3) could lead to accuracy loss. For instance, for $\gamma = 100$, the AFL could suffer a loss of 9% (i.e., $58.56\% \rightarrow 49.62\%$). This is the result of regularization accumulation (see (15)). With the RI process, the AFL obtains an identical result across various γ values. More importantly, this demonstrates that adopting the RI avoids the need to find proper γ values. That is, the regularization is a removable intermediary, not a hyperparameter that requires tuning.

Table 3. Ablation study of RI under various γ and K . The left/right results are performance w/o and w/ the RI process in (16).

Acc.(%)	$\gamma = 0$	$\gamma = 0.1$	$\gamma = 1$	$\gamma = 10$	$\gamma = 100$
K=100	58.56 N/A	58.54 58.56	58.51 58.56	58.15 58.56	55.77 58.56
K=500	1.11 N/A	58.52 58.56	58.30 58.56	56.72 58.56	51.77 58.56
K=1000	0.75 N/A	58.51 58.56	58.15 58.56	55.77 58.56	49.62 58.56

4.5. Validation with Different Backbones

To explore the effect of different backbones used in the AFL, we extend the experiments with VGG11 [28] and ViT-B-16 [6]. All these backbone are pre-trained in ImageNet-1k and we conduct the experiments under the same setting in Section 4.2. Due to the invariance to data partitioning, we only report one result in single dataset. As shown in the Table 4, with different pre-trained backbones, the AFL can all obtain competitive results.

5. Limitations and Future Work

Utilizing Pre-trained Backbone. The AFL approach is both facilitated and constrained by the requirement of having a well-trained feature extractor. However, this limitation has

Table 4. The results of top-1 accuracy in % of the AFL with different backbones including ResNet-18, VGG11 and ViT-B-16.

Acc.(%)	CIFAR-10	CIFAR-100	Tiny-ImageNet
ResNet-18	80.75	58.56	54.67
VGG11	82.72	60.43	54.73
ViT-B-16	93.92	75.45	82.02

been significantly mitigated by the emergence of reusing pre-trained models for new tasks. This “pre-trained backbone + downstream task” paradigm has become a standard practice in numerous deep learning domains, offering improved generalization and reduced computational costs. FL can further enhance this paradigm and we validate that collaboration can still be beneficial with pre-trained backbones in Supplementary Materials F. The proposal of the AFL aligns with these recent research trends, making it a sensible FL advancement.

Partially Participating and Stragglers. The AFL formulates a single-round aggregation for FL systems, promoting rapid convergence and reducing communication overhead. However, challenges arise when clients engage partially or when stragglers impede progress. Since clients can only contribute to the aggregation after finishing local computations, the AFL needs to wait for all the clients. This potentially hampers the AFL’s overall efficiency and inspires us to further refine the AFL to address these issues.

Linear Assumptions of AFL. The AFL is established upon linear classifiers and may be less effective with non-linear data distribution. To address this, AFL can incorporate non-linear projections including non-linear activations or kernel functions. Also, for multi-layer model, AFL can formulate local least-square problem at each layer by label projection [49]. These techniques have been utilized in various AL-based work [39] and the AA law holds theoretically. We will conduct a further exploration in future.

6. Conclusion

In this paper, we introduce a gradient-free FL framework named analytic federated learning (AFL). The AFL unveils analytical solutions both in the local client training stage and the aggregation stage. This leads to one-epoch local training, single-round aggregation, and fast convergence. In particular, the single-round aggregation property is theoretically supported and proved by the well-formulated AA law. Additionally, by introducing the RI process, we re-establish the AFL’s optimality which could be compromised in the scenario of rank-deficient with normally a large number of clients. The AFL demonstrates its invariance to data partitioning, a property that allows several appealing FL characteristics such as data heterogeneity invariance and client-number invariance. These characteristics are empirically validated through experiments across various settings, where the AFL achieves a consistent and competitive performance.

Acknowledgment

This research was supported by the National Natural Science Foundation of China (62306117, 62406114), the Guangzhou Basic and Applied Basic Research Foundation (2024A04J3681, 2023A04J1687), GJYC program of Guangzhou (2024D03J0005), the National Key R & D Project from Minister of Science and Technology (2024YFA1211500), and the Fundamental Research Funds for the Central Universities (2024ZYGXZR074).

References

- [1] Durmus Alp Emre Acar, Yue Zhao, Ramon Matas, Matthew Mattina, Paul Whatmough, and Venkatesh Saligrama. Federated learning based on dynamic regularization. In *International Conference on Learning Representations*, 2021. 6
- [2] Ilai Bistriz, Ariana Mann, and Nicholas Bambos. Distributed distillation for on-device learning. In *Advances in Neural Information Processing Systems*, pages 22593–22604. Curran Associates, Inc., 2020. 2
- [3] Hong-You Chen, Cheng-Hao Tu, Ziwei Li, Han Wei Shen, and Wei-Lun Chao. On the importance and applicability of pre-training for federated learning. In *The Eleventh International Conference on Learning Representations*, 2023. 6
- [4] Randall E. Cline. Representations for the generalized inverse of a partitioned matrix. *Journal of the Society for Industrial and Applied Mathematics*, 12(3):588–600, 1964. 4
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. 6
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021. 8
- [7] Chun-Mei Feng, Bangjun Li, Xinxing Xu, Yong Liu, Huazhu Fu, and Wangmeng Zuo. Learning federated visual prompt in null space for mri reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8064–8073, 2023. 1, 6
- [8] Ping Guo, Michael R Lyu, and NE Mastorakis. Pseudoinverse learning algorithm for feedforward neural networks. *Advances in Neural Networks and Applications*, pages 321–326, 2001. 2, 3
- [9] Qiushan Guo, Xinjiang Wang, Yichao Wu, Zhipeng Yu, Ding Liang, Xiaolin Hu, and Ping Luo. Online knowledge distillation via collaborative learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [11] Georgios A Kaissis, Marcus R Makowski, Daniel Rückert, and Rickmer F Braren. Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence*, 2(6):305–311, 2020. 1
- [12] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. 6
- [13] Rajesh Kumar, Abdullah Aman Khan, Jay Kumar, Zakria, Noorbakhsh Amiri Golilarz, Simin Zhang, Yang Ting, Chengyu Zheng, and Wenyong Wang. Blockchain-federated-learning and deep learning models for covid-19 detection using ct imaging. *IEEE Sensors Journal*, 21(14):16301–16314, 2021. 1
- [14] Gihun Lee, Minchan Jeong, Yongjin Shin, Sangmin Bae, and Se-Young Yun. Preservation of the global knowledge by not-true distillation in federated learning. In *Advances in Neural Information Processing Systems*, pages 38461–38474. Curran Associates, Inc., 2022. 6
- [15] Qinbin Li, Bingsheng He, and Dawn Song. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10713–10722, 2021. 6
- [16] Qinbin Li, Yiqun Diao, Quan Chen, and Bingsheng He. Federated learning on non-iid data silos: An experimental study. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*, pages 965–978, 2022. 1
- [17] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. In *Proceedings of Machine Learning and Systems*, pages 429–450, 2020. 1, 2, 6
- [18] Xin-Chun Li, Yi-Chu Xu, Shaoming Song, Bingshuai Li, Yinchuan Li, Yunfeng Shao, and De-Chuan Zhan. Federated learning with position-aware neurons. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10082–10091, 2022. 2
- [19] Zexi Li, Tao Lin, Xinyi Shang, and Chao Wu. Revisiting weighted aggregation in federated learning with neural networks. In *Proceedings of the 40th International Conference on Machine Learning*, pages 19767–19788. PMLR, 2023. 2, 6
- [20] Tao Lin, Lingjing Kong, Sebastian U Stich, and Martin Jaggi. Ensemble distillation for robust model fusion in federated learning. In *Advances in Neural Information Processing Systems*, pages 2351–2363. Curran Associates, Inc., 2020. 6
- [21] Tao Lin, Lingjing Kong, Sebastian U Stich, and Martin Jaggi. Ensemble distillation for robust model fusion in federated learning. In *Advances in Neural Information Processing Systems*, pages 2351–2363. Curran Associates, Inc., 2020. 2
- [22] Zichen Liu, Chao Du, Wee Sun Lee, and Min Lin. Locality sensitive sparse encoding for learning world models online. In *The Twelfth International Conference on Learning Representations*, 2024. 4
- [23] Shubham Malaviya, Manish Shukla, and Sachin Lodha. Reducing communication overhead in federated learning for pre-trained language models using parameter-efficient fine-tuning. In *Proceedings of The 2nd Conference on Lifelong Learning Agents*, pages 456–469. PMLR, 2023. 1
- [24] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-Efficient Learning of Deep Networks from Decentralized

- Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pages 1273–1282. PMLR, 2017. [1](#), [2](#), [6](#)
- [25] Mohammed Ali mnmoustafa. Tiny imagenet, 2017. [6](#)
- [26] J. Park and I. W. Sandberg. Universal approximation using radial-basis-function networks. *Neural Computation*, 3(2): 246–257, 1991. [3](#)
- [27] Geet Shingi. A federated learning based approach for loan defaults prediction. In *2020 International Conference on Data Mining Workshops (ICDMW)*, pages 362–368, 2020. [1](#)
- [28] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition, 2015. [8](#)
- [29] Yue Tan, Guodong Long, LU LIU, Tianyi Zhou, Qinghua Lu, Jing Jiang, and Chengqi Zhang. Fedproto: Federated prototype learning across heterogeneous clients. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(8):8432–8440, 2022. [1](#)
- [30] Yue Tan, Guodong Long, Jie Ma, LU LIU, Tianyi Zhou, and Jing Jiang. Federated learning from pre-trained models: A contrastive learning approach. In *Advances in Neural Information Processing Systems*, pages 19332–19344. Curran Associates, Inc., 2022. [1](#)
- [31] Kar-Ann Toh. Learning from the kernel and the range space. In *2018 IEEE/ACIS 17th International Conference on Computer and Information Science (ICIS)*, pages 1–6, 2018. [3](#)
- [32] Kar-Ann Toh. Learning from the kernel and the range space. In *the Proceedings of the 17th 2018 IEEE Conference on Computer and Information Science*, pages 417–422. IEEE, 2018. [2](#)
- [33] Haozhao Wang, Yichen Li, Wenchao Xu, Ruixuan Li, Yufeng Zhan, and Zhigang Zeng. DaFKD: Domain-aware federated knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20412–20421, 2023. [2](#)
- [34] Haozhao Wang, Haoran Xu, Yichen Li, Yuan Xu, Ruixuan Li, and Tianwei Zhang. FedCDA: Federated learning with cross-rounds divergence-aware aggregation. In *The Twelfth International Conference on Learning Representations*, 2024. [2](#)
- [35] Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H. Vincent Poor. Tackling the objective inconsistency problem in heterogeneous federated optimization. In *Advances in Neural Information Processing Systems*, pages 7611–7623. Curran Associates, Inc., 2020. [1](#), [2](#)
- [36] Jue Wang, Ping Guo, and Yanjun Li. DensePILAE: a feature reuse pseudoinverse learning algorithm for deep stacked autoencoder. *Complex & Intelligent Systems*, pages 1–11, 2021. [3](#)
- [37] X. Wang, T. Zhang, and R. Wang. Noniterative deep learning: Incorporating restricted boltzmann machine into multilayer random weight neural networks. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 49(7):1299–1308, 2019. [3](#)
- [38] Guile Wu and Shaogang Gong. Peer collaborative learning for online knowledge distillation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(12):10302–10310, 2021. [2](#)
- [39] Yicheng Xu, Yuxin Chen, Jiahao Nie, Yusong Wang, Huiping Zhuang, and Manabu Okumura. Advancing cross-domain discriminability in continual learning of vision-language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. [8](#)
- [40] Fu-En Yang, Chien-Yi Wang, and Yu-Chiang Frank Wang. Efficient model personalization in federated learning via client-specific prompt generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 19159–19168, 2023. [1](#)
- [41] Wensi Yang, Yuhang Zhang, Kejiang Ye, Li Li, and Cheng-Zhong Xu. Ffd: A federated learning based method for credit card fraud detection. In *Big Data – BigData 2019*, pages 18–32, Cham, 2019. Springer International Publishing. [1](#)
- [42] Rui Ye, Mingkai Xu, Jianyu Wang, Chenxin Xu, Siheng Chen, and Yanfeng Wang. FedDisco: Federated learning with discrepancy-aware collaboration. In *Proceedings of the 40th International Conference on Machine Learning*, pages 39879–39902. PMLR, 2023. [6](#)
- [43] Fuxun Yu, Weishan Zhang, Zhuwei Qin, Zirui Xu, Di Wang, Chenchen Liu, Zhi Tian, and Xiang Chen. Fed2: Feature-aligned federated learning. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pages 2066–2074, 2021. [2](#)
- [44] Lin Zhang, Li Shen, Liang Ding, Dacheng Tao, and Ling-Yu Duan. Fine-tuning global model via data-free knowledge distillation for non-iid federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10174–10183, 2022. [2](#)
- [45] Zhuo Zhang, Yuanhang Yang, Yong Dai, Qifan Wang, Yue Yu, Lizhen Qu, and Zenglin Xu. Fedpetuning: When federated learning meets the parameter-efficient tuning methods of pre-trained language models. In *Findings of the Association for Computational Linguistics: ACL 2023*, page 9963–9977. Association for Computational Linguistics (ACL), 2023. Annual Meeting of the Association of Computational Linguistics 2023, ACL 2023 ; Conference date: 09-07-2023 Through 14-07-2023. [1](#)
- [46] Zhuangdi Zhu, Junyuan Hong, and Jiayu Zhou. Data-free knowledge distillation for heterogeneous federated learning. In *Proceedings of the 38th International Conference on Machine Learning*, pages 12878–12889. PMLR, 2021. [6](#)
- [47] Zhuangdi Zhu, Junyuan Hong, and Jiayu Zhou. Data-free knowledge distillation for heterogeneous federated learning. In *Proceedings of the 38th International Conference on Machine Learning*, pages 12878–12889. PMLR, 2021. [2](#)
- [48] Huiping Zhuang, Zhiping Lin, and Kar-Ann Toh. Blockwise recursive Moore-Penrose inverse for network learning. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, pages 1–14, 2021. [2](#), [3](#), [6](#)
- [49] Huiping Zhuang, Zhiping Lin, and Kar-Ann Toh. Correlation projection for analytic learning of a classification network. *Neural Processing Letters*, pages 1–22, 2021. [8](#)
- [50] Huiping Zhuang, Zhenyu Weng, Hongxin Wei, Renchunzi Xie, Kar-Ann Toh, and Zhiping Lin. ACIL: Analytic class-incremental learning with absolute memorization and privacy protection. In *Advances in Neural Information Processing*

Systems, pages 11602–11614. Curran Associates, Inc., 2022. [3](#), [4](#)

- [51] Huiping Zhuang, Zhenyu Weng, Run He, Zhiping Lin, and Ziqian Zeng. GKEAL: Gaussian kernel embedded analytic learning for few-shot class incremental task. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7746–7755, 2023. [3](#)
- [52] Huiping Zhuang, Yizhu Chen, Di Fang, Run He, Kai Tong, Hongxin Wei, Ziqian Zeng, and Cen Chen. GACL: Exemplar-free generalized analytic continual learning. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2024. [3](#), [6](#)
- [53] Huiping Zhuang, Run He, Kai Tong, Ziqian Zeng, Cen Chen, and Zhiping Lin. DS-AL: A dual-stream analytic learning for exemplar-free class-incremental learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(15):17237–17244, 2024. [3](#)
- [54] Huiping Zhuang, Yuchen Liu, Run He, Kai Tong, Ziqian Zeng, Cen Chen, Yi Wang, and Lap-Pui Chau. F-OAL: Forward-only online analytic learning with fast training and low memory footprint in class incremental learning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. [6](#)