

000
001
002
003
004
005
006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053

DEEPAFL: DEEP ANALYTIC FEDERATED LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Federated Learning (FL) is a popular distributed learning paradigm to break down data silo. Traditional FL approaches largely rely on gradient-based updates, facing significant issues about heterogeneity, scalability, convergence, and overhead, etc. Recently, some analytic-learning-based work has attempted to handle these issues by eliminating gradient-based updates via analytical (i.e., closed-form) solutions. Despite achieving superior invariance to data heterogeneity, these approaches are fundamentally limited by their single-layer linear model with a frozen pre-trained backbone. As a result, they can only achieve suboptimal performance due to their lack of representation learning capabilities. In this paper, to enable representable analytic models while preserving the ideal invariance to data heterogeneity for FL, we propose our *Deep* Analytic *Federated* Learning approach, named **DeepAFL**. Drawing inspiration from the great success of ResNet in gradient-based learning, we design gradient-free residual blocks in our DeepAFL with analytical solutions. We further introduce an efficient layer-wise protocol for training our deep analytic models layer by layer in FL through least squares. Both theoretical analyses and empirical evaluations validate our DeepAFL’s superior performance with its dual advantages in heterogeneity invariance and representation learning, outperforming state-of-the-art baselines by up to 5.68%–8.42% across three benchmark datasets. The related codes will be made open-sourced upon the acceptance of this paper.

1 INTRODUCTION

Federated Learning (FL) has emerged as a prominent paradigm that enables distributed machine learning to break down data silos (Fan et al., 2025c; Yang et al., 2023). The objective of FL is to allow a group of clients to collaboratively train a powerful and robust global model, while preserving their data privacy (Ren et al., 2025; Liu et al., 2024b). Currently, the field of FL has seen substantial growth across a wide variety of applications (Zhou et al., 2024; Wu et al., 2024; Rong et al., 2025).

Traditional FL methods rely on a gradient-based optimization paradigm, exemplified by the classic FedAvg (McMahan et al., 2017) and its following variants (Li et al., 2021b; Yang et al., 2024). These methods typically necessitate iterative optimization processes to achieve convergence (Wang et al., 2024; Tan et al., 2022a). Yet, these gradient-based techniques are widely acknowledged to suffer from several major challenges (Ye et al., 2023a; Chai et al., 2024; He et al., 2025b), as follows.

- (1) **Heterogeneity Issues:** The data across clients are often Not Independently and Identically Distributed (Non-IID), which can severely impact model performance and convergence.
- (2) **Scalability Issues:** As the number of clients increases, especially to a large scale (e.g., thousands of clients), the FL systems can experience substantial performance degradation.
- (3) **Convergence Issues:** The FL methods may struggle to converge within limited aggregation rounds, particularly in challenging scenarios of non-IID data or large-scale clients.
- (4) **Overhead Issues:** The overall FL process incurs significant overhead from multi-epoch training on each client and multi-round model aggregation across clients for convergence.

Many researchers have come to realize that the aforementioned challenges in FL are fundamentally rooted in the long-standing reliance on gradient-based updates, which are inherently sensitive and costly in the distributed FL scenario (Ye et al., 2023a; Fanì et al., 2024; He et al., 2025b). From this perspective, existing gradient-based methods can only superficially alleviate these issues rather than fundamentally address them. Therefore, a natural and promising avenue to fundamentally address these gradient-related issues is to eliminate gradient-based updates entirely (He et al., 2025b).

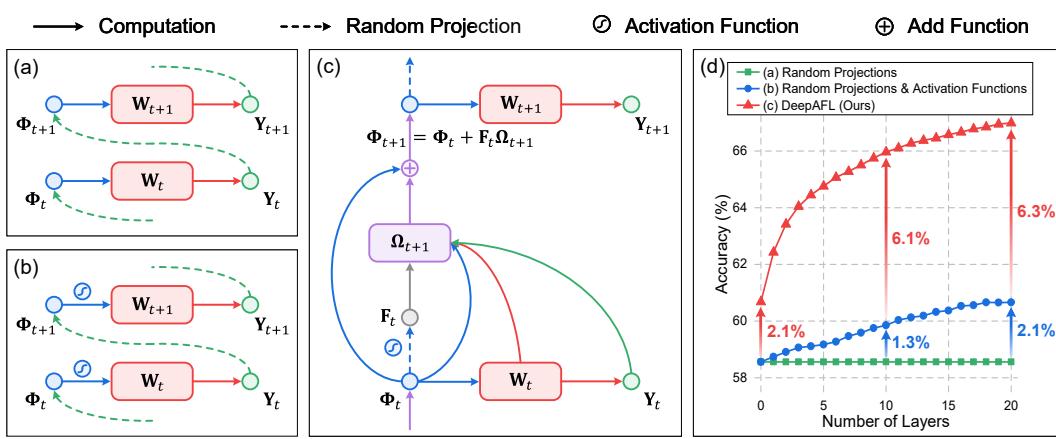


Figure 1: Comparing our proposed DeepAFL with the naive approaches for representation learning. (a) Illustration of multi-analytic layers with random projections. (b) Illustration of multi-analytic layers with random projections & activation functions. (c) Illustration of our proposed DeepAFL. (d) Among these, our DeepAFL exhibits the best performance improvements with increasing layers. For all these approaches, the activation function is GELU, the random projection dimension is 1024, the pre-trained model is the ResNet-18 used in AFL, and the evaluation dataset is the CIFAR-100.

As a prominent gradient-free technique, analytic learning has exhibited great promise by achieving analytical (closed-form) solutions through least squares (Zhuang et al., 2024a; Fan et al., 2025a). To introduce this technique into the FL, Analytic Federated Learning (AFL) has been proposed with the help of pre-trained models (He et al., 2025b). The core idea of AFL is to leverage frozen pre-trained models (e.g., foundation models) for feature extraction on input data. Based on the extracted features, AFL can then build a single-layer linear model, which has an analytical solution. Through its specific protocols on local training and global aggregation, AFL achieves ideal invariance to data heterogeneity, as its final aggregated global model is equivalent to the centralized analytical solution.

Despite achieving State-Of-The-Art (SOTA) performance with its unique invariance to data heterogeneity, AFL is significantly limited by its single-layer linear model. On the one hand, this analytic model is foundational to AFL, as its simplicity with a convex optimization objective is a prerequisite for deriving an analytical solution. On the other hand, it also severely limits AFL’s application, as it can only learn a **linear mapping** from the frozen backbone’s features to the final output, thus **fundamentally failing to perform representation learning within the FL systems**. Consequently, due to the limited learning capacity of the linear analytic model, AFL is prone to underfitting, especially when the backbone itself is lightweight. Additionally, even if the backbone itself possesses sufficient feature extraction capabilities, the linear separability of its output features may still be insufficient.

Thus, an interesting and challenging problem arises: *can we deepen AFL’s analytic model to enable its representation learning capabilities while simultaneously preserving its analytical solutions for invariance to data heterogeneity?* A quite naive approach is using a random projection after each analytic model to construct the input features for the next layer, as displayed in Figure 1(a). Yet, this approach brings almost no performance improvement, as all its layers are simply linear mappings. Then, we attempt to introduce an activation function after each feature to provide non-linearity, as illustrated in Figure 1(b). Based on the evaluation results presented in Figure 1(d), the deep activated random projections can, to some extent, enrich the feature representations. However, as the number of layers increases, the deepening approach (b) struggles to improve the performance very little from representation learning. Therefore, these naive approaches fall far short of our requirements.

Drawing inspiration from the great success of ResNet in gradient-based learning (He et al., 2016), in this paper, we adopt the similar skip connections to boost the representation of the analytic layers. Specifically, we model this representation learning as $\Phi_t = \Phi_{t-1} + g_t(\Phi_{t-1})$, where $g_t(\cdot)$ represents a nonlinear feature transformation. In gradient-based learning, the residual blocks $g_t(\Phi_{t-1})$ can be easily learned, as Stochastic Gradient Descent (SGD) with backpropagation can automatically adjust the network weights to learn appropriate representations. Given that analytic models preclude gradient-based updates via backpropagation, a key technical challenge lies in how to effectively learn the residual blocks $g_t(\Phi_{t-1})$ for meaningful boosting within the framework of analytic learning.

To enable representable analytic models while preserving their ideal invariance to data heterogeneity for FL, in this paper, we propose our *Deep* Analytic Federated Learning approach, named *DeepAFL*. Inherited from AFL, our DeepAFL also employs a pre-trained backbone for initial feature extraction. Then, we propose to randomly project and activate these features to form the zero-layer features Φ_0 , which can yield an immediate performance gain of about 2.1%, as shown in Figure 1(d). After this setup, our DeepAFL continuously refines the features Φ_t layer by layer. Specifically, to obtain the residual block $g_t(\Phi_{t-1})$ for layer t , we first build a nonlinear representation \mathbf{F}_{t-1} from Φ_{t-1} using a random projection layer with an activation function, as illustrated in Figure 1(c). Subsequently, we can obtain $g_t(\Phi_{t-1}) = \mathbf{F}_{t-1}\Omega_t$ by introducing a learnable transformation Ω_t to adjust and scale the representation \mathbf{F}_{t-1} . We derive the optimal analytical solution for Ω_t via *sandwiched least squares*. As shown in Figure 1(d), our DeepAFL exhibits a desired capability of deep representation learning, as its performance continuously and markedly improves with an increasing number of layers.

Our primary contributions are summarized as follows:

- Conceptually, we propose our DeepAFL, a novel approach that can achieve gradient-free representation learning while preserving ideal invariance to data heterogeneity in FL.
- Technically, we develop an efficient layer-wise protocol for learning deep analytic models via least squares. In our DeepAFL, the clients only need to conduct lightweight forward-propagation computations, so that the server can aggregate the global models layer by layer.
- Theoretically, we demonstrate two ideal properties of our DeepAFL: its invariance to data heterogeneity and its capability of representation learning. To the best of our knowledge, our DeepAFL represents the first to achieve both of these ideal properties simultaneously.
- Experimentally, we provide extensive evaluations on three benchmark datasets to show the superiority of our DeepAFL, which outperforms SOTA baselines by up to 5.68%–8.42%, thanks to its dual advantages in heterogeneity invariance and representation learning.

2 RELATED WORK

2.1 FEDERATED LEARNING

As a prominent distributed learning paradigm, FL allows multiple clients to collaboratively train a global model to break down data silos (Fan et al., 2025c; Yang et al., 2023). Existing FL techniques are largely derived from FedAvg (McMahan et al., 2017) and rely on gradient-based optimization. Despite advancements in FL, their reliance on gradients also causes inherent issues about overhead, convergence, heterogeneity, and scalability (Ye et al., 2023a; Chai et al., 2024; He et al., 2025b). While several studies have explored representation learning in FL, such as FedRep (Collins et al., 2021) and FedU² (Liao et al., 2024), they still fail to solve the preceding gradient-based issues.

Recently, AFL has introduced a new and promising wave that fundamentally handles these issues by avoiding gradient-based updates via analytic learning (He et al., 2025b). Based on it, some subsequent research has extended this concept to personalized FL and federated continual learning (Fan et al., 2025b; Tang et al., 2025b), achieving excellent performance. Nevertheless, a key limitation of existing analytic-learning-based FL approaches is their lack of representation learning capability.

2.2 ANALYTIC LEARNING

Analytic learning (Zhuang et al., 2022; 2023; 2024b), also known as pseudoinverse learning (Cline, 1964; Guo et al., 2001), has emerged as a popular gradient-free technique to address gradient-related issues (Toh, 2018; Lanthaler & Nelsen, 2023; Prabhu et al., 2024; Zozoulenko et al., 2025; Fan et al., 2025a; Bolager et al., 2023). Its core idea is to directly derive analytical (i.e., closed-form) solutions using least squares, to eliminate gradient-based updates (Lai et al., 2025; Peng et al., 2025). Currently, analytic learning has been widely applied in various modern tasks and achieved promising results (Liu et al., 2024c; Li et al., 2025; He et al., 2025b; Tang et al., 2025a; Tran et al., 2025).

Despite the extensive superiority shown by analytic learning, current approaches are largely limited by their single-layer linear models, which sacrifice the capabilities of deep representation learning. On the other hand, achieving representable deep analytic models while preserving their closed-form solutions stands as a great challenge, especially in the distributed scenario of FL. Drawing inspiration from the success of ResNet (He et al., 2016) in gradient-based learning, we fundamentally solve this problem by carefully designing gradient-free residual blocks with closed-form solutions.

162 **3 OUR PROPOSED DEEPAFL**
 163

164 In this section, we elaborate in detail on our proposed DeepAFL. For clarity, we consider the typical
 165 FL setting involving one server and K clients. Moreover, to align with existing research, we use the
 166 most prevalent task in FL, image recognition, as the prime example to describe the workflow of our
 167 DeepAFL (Yu et al., 2024; 2025; Yang et al., 2024; He et al., 2025b). Each client’s local dataset is
 168 denoted as $\mathcal{D}^k = \{\mathbf{X}^k, \mathbf{Y}^k\}$, where $\mathbf{X}^k \in \mathbb{R}^{N_k \times d_H \times d_W \times d_C}$ and $\mathbf{Y}^k \in \mathbb{R}^{N_k \times C}$ represent the N_k local
 169 samples and their corresponding labels, respectively. Here, $d_H \times d_W \times d_C$ denotes the 3 dimensions
 170 (height, width, and channels) of the input images, and C denotes the number of output classes.

171 The objective of our proposed DeepAFL is to construct a deep residual analytic network comprising
 172 T layers for deriving the features Φ_T and the corresponding global classifier \mathbf{W}_T , in a distributed
 173 and gradient-free manner. In Section 3.1, we detail the motivation and insight of the proposed deep
 174 analytic learning in our DeepAFL, which is described from a centralized perspective for clarity. In
 175 Section 3.2, we elaborate on the specific workflow and implementation of our DeepAFL within the
 176 data-distributed FL scenario. In Section 3.3, we provide theoretical analyses of our DeepAFL for its
 177 validity, privacy, and efficiency, particularly showing its dual ideal properties.

178 **3.1 DEEP RESIDUAL ANALYTIC LEARNING OF OUR DEEPAFL**
 179

180 Here, we aim to introduce the key motivation and insights behind the proposed deep residual analytic
 181 model in our DeepAFL. For clarity and brevity, we first adopt a centralized perspective and consider
 182 a full dataset $\mathcal{D} = \{\mathbf{X}, \mathbf{Y}\}$ to construct the T -layer deep residual analytic network. Specifically, we
 183 progressively learn better representations and refine the features $\{\Phi_t\}_{t=0}^T$ layer by layer. Based on
 184 the features Φ_t for each layer t , we construct a corresponding analytic classifier \mathbf{W}_t .

185 First of all, inherited from AFL (He et al., 2025b), we employ a pre-trained backbone for our initial
 186 feature extraction from \mathbf{X} to obtain $\tilde{\mathbf{X}} \in \mathbb{R}^{N \times d_X}$ via (1). This has been widely adopted as a common
 187 practice in many recent studies of FL (Nguyen et al., 2023; Piao et al., 2024; Yu et al., 2024; 2025).

$$\tilde{\mathbf{X}} = \text{Backbone}(\mathbf{X}, \Theta), \quad (1)$$

188 where $\text{Backbone}(\cdot)$ represents the pre-trained backbone with frozen parameters Θ . Based on this
 189 initial feature extraction of AFL, we further incorporate the activated random projection to boost the
 190 features’ representation, thereby forming the zero-layer features $\Phi_0 \in \mathbb{R}^{N \times d_\Phi}$, as follows.

$$\Phi_0 = \sigma(\tilde{\mathbf{X}} \mathbf{A}), \quad (2)$$

191 where $\sigma(\cdot)$ represents the activation function, and $\mathbf{A} \in \mathbb{R}^{d_X \times d_\Phi}$ is the random projection matrix.
 192 This activated random projection can increase the feature dimension to a suitable size, boosting its
 193 linear separability, as widely validated in existing studies (Cover, 2006; Zhuang et al., 2022; 2024b;
 194 Zhang et al., 2025). To progressively learn richer representations, we attempt to deepen the network.
 195 Yet, naive approaches to deepening the network fail to achieve it effectively, as shown in Figure 1.

196 Drawing inspiration from the great success of ResNet in gradient-based learning (He et al., 2016),
 197 we adopt the similar skip connections to boost the representations of the analytic layers. Specifically,
 198 we model this representation learning process as the feature updating formula in (3):

$$\Phi_t = \Phi_{t-1} + g_t(\Phi_{t-1}), \forall t \in [1, T]. \quad (3)$$

199 Here, $g_t(\cdot)$ represents the residual block as a nonlinear feature transformation. We will specify the
 200 detailed gradient-free design of the residual block within our DeepAFL in (6) later. Built upon the
 201 obtained feature matrix Φ_t for each layer t , we can construct a corresponding analytic classifier \mathbf{W}_t .
 202 Specifically, consistent with other existing analytic-learning-based approaches (He et al., 2025b), the
 203 optimization objective for the analytic classifier \mathbf{W}_t can be formulated as follows.

$$\mathbf{W}_t = \arg \min_{\mathbf{W}} \|\mathbf{Y} - \Phi_t \mathbf{W}\|_F^2 + \lambda \|\mathbf{W}\|_F^2, \forall t \in [0, T], \quad (4)$$

204 where λ denotes the regularization parameter and $\|\cdot\|_F^2$ represents the Frobenius norm. Since cross-
 205 entropy loss does not admit a closed-form solution, we use the Mean Squared Error (MSE) loss here.
 206 In fact, the MSE loss is widely adopted in analytic learning (Zhuang et al., 2022; 2023; 2024b) and
 207 can achieve performance comparable to that of the cross-entropy loss (Hui & Belkin, 2021). Using
 208 the least squares method, we can derive the optimal analytical solution to the objective (4), as given
 209 by (5). The detailed proof of the analytical solution (5) is provided in **Lemma 1** of Section 3.3.

$$\mathbf{W}_t = (\Phi_t^\top \Phi_t + \lambda \mathbf{I})^{-1} \Phi_t^\top \mathbf{Y}, \forall t \in [0, T]. \quad (5)$$

216 After obtaining the analytic classifier \mathbf{W}_t , we then proceed with deep residual representation learning to update the $(t+1)$ -th-layer features. Based on (3), the key to this problem is the gradient-free
 217 design of the residual blocks, which possess analytical solutions with the properties of **stochasticity**,
 218 **nonlinearity**, and **learnability**. Thus, in our DeepAFL, the residual block is instantiated as follows:
 219

$$g_{t+1}(\Phi_t) = \sigma(\Phi_t \mathbf{B}_t) \Omega_{t+1} = \mathbf{F}_t \Omega_{t+1}, \forall t \in [0, T]. \quad (6)$$

220 Here, $\mathbf{B}_t \in \mathbb{R}^{d_\Phi \times d_F}$ represents the random projection matrix for **stochasticity**, akin to that provided
 221 by SGD in gradient-based learning. Meanwhile, $\sigma(\cdot)$ means the activation function for **nonlinearity**,
 222 a component that is widely demonstrated to be essential in deep representation learning. Moreover,
 223 $\Omega_{t+1} \in \mathbb{R}^{d_F \times d_\Phi}$ represents the trainable transformation matrix for providing **learnability**. In (6), we
 224 further define $\mathbf{F}_t = \sigma(\Phi_t \mathbf{B}_t)$ as the hidden random feature, thereby isolating the trainable Ω_{t+1} . In
 225 Section 4, we will provide extensive ablation studies to show the effectiveness of these components.
 226

227 Then, we focus on learning the optimal Ω_{t+1} for effective representation boosting within the residual
 228 block. Here, our objective can be written as, given a fixed classifier \mathbf{W}_t from the previous layer, to
 229 minimize the empirical risk by optimizing the new features $\Phi_{t+1} = \Phi_t + \mathbf{F}_t \Omega_{t+1}$, as follows:
 230

$$\begin{aligned} \Omega_{t+1} &= \arg \min_{\Omega} \|\mathbf{Y} - (\Phi_t + \mathbf{F}_t \Omega) \mathbf{W}_t\|_F^2 + \gamma \|\Omega\|_F^2 \\ &= \arg \min_{\Omega} \|\mathbf{R}_t - \mathbf{F}_t \Omega \mathbf{W}_t\|_F^2 + \gamma \|\Omega\|_F^2, \forall t \in [0, T]. \end{aligned} \quad (7)$$

231 Here, we define the residual of the current layer's classification as $\mathbf{R}_t = \mathbf{Y} - (\Phi_t \mathbf{W}_t)$ for notational
 232 convenience. The optimization objective (7) can be seen as a special case of generalized Sylvester
 233 matrix equations (Wu et al., 2008; Ding et al., 2008; Duan, 2015), with a specific structure where the
 234 unknown variable Ω is *sandwiched* between two known ones, \mathbf{F}_t and \mathbf{W}_t . This structure enables
 235 the derivation of an analytical solution, as shown in (8). We term it as the *sandwiched least squares*
 236 problem, and provide a detailed proof of the analytical solution (8) in **Lemma 2** of Section 3.3.
 237

$$\Omega_{t+1} = \mathbf{V}_t [(\mathbf{V}_t^\top \mathbf{F}_t^\top \mathbf{R}_t \mathbf{W}_t^\top \mathbf{U}_t) \oslash (\gamma \mathbf{1} + \text{diag}(\Lambda_t^F) \otimes \text{diag}(\Lambda_t^W))] \mathbf{U}_t^\top, \forall t \in [0, T], \quad (8)$$

238 where $\mathbf{F}_t^\top \mathbf{F}_t = \mathbf{V}_t \Lambda_t^F \mathbf{V}_t^\top$ and $\mathbf{W}_t^\top \mathbf{W}_t = \mathbf{U}_t \Lambda_t^W \mathbf{U}_t^\top$ are spectral decompositions, while \oslash and \otimes
 239 represent element-wise division and outer product, respectively. Once we obtain the transformation
 240 matrix Ω_{t+1} , we can compute the next-layer features via the following recursive formula.
 241

$$\Phi_{t+1} = \Phi_t + \mathbf{F}_t \Omega_{t+1}, \forall t \in [0, T]. \quad (9)$$

242 In summary, the construction of our deep residual analytic network proceeds in a layer-wise manner,
 243 involving the alternating derivation of \mathbf{W}_t and Ω_{t+1} via analytical solutions (5) and (8), respectively.
 244 Specifically, the solution procedure follows the sequence $\mathbf{W}_0 \mapsto \Omega_1 \mapsto \mathbf{W}_1 \mapsto \dots \mapsto \Omega_T \mapsto \mathbf{W}_T$,
 245 beginning with the zero-layer classifier \mathbf{W}_0 and concluding with the final-layer classifier \mathbf{W}_T . For
 246 further clarity, we provide detailed illustrations of our DeepAFL's formulations in Figures 14–16 of
 247 Appendix H. Based on the above centralized process, we will then elaborate on the specific workflow
 248 and implementation of our DeepAFL for the data-distributed FL scenario in Section 3.2.
 249

3.2 FEDERATED IMPLEMENTATION OF OUR DEEPAFL

250 In this subsection, we display the federated implementation of our DeepAFL for the data-distributed
 251 FL scenario, where each client owns its local dataset $\mathcal{D}^k = \{\mathbf{X}^k, \mathbf{Y}^k\}$. The overall implementation
 252 workflow of our DeepAFL still follows the layer-wise procedure that is described in Section 3.1. Yet,
 253 in the FL setting, the analytical computation of the global weights \mathbf{W}_t and Ω_{t+1} requires additional
 254 aggregation of all clients' local knowledge to update the global knowledge.
 255

256 First of all, each client k performs feature extraction on its local data \mathbf{X}^k to obtain its local zero-layer
 257 feature matrix $\Phi_0^k \in \mathbb{R}^{N_k \times d_\Phi}$ similar to (1) and (2), as follows.
 258

$$\Phi_0^k = \sigma(\text{Backbone}(\mathbf{X}^k, \Theta) \mathbf{A}). \quad (10)$$

259 Then, for each layer $t \in [0, T]$, each client utilizes its features Φ_t^k and labels \mathbf{Y}^k to compute its local
 260 *Feature Auto-Correlation Matrix* \mathbf{G}_t^k and *Label Cross-Correlation Matrix* \mathbf{H}_t^k , as follows.
 261

$$\mathbf{G}_t^k = (\Phi_t^k)^\top \Phi_t^k, \quad \mathbf{H}_t^k = (\Phi_t^k)^\top \mathbf{Y}^k. \quad (11)$$

262 Then, the server needs to aggregate $\mathbf{G}_t^{1:K}$ and $\mathbf{H}_t^{1:K}$ from the clients through (12). This process can
 263 be implemented via existing Secure Aggregation Protocols (Bonawitz et al., 2016; 2017; So et al.,
 264 2023). We provide related privacy analyses in Section 3.3.
 265

$$\mathbf{G}_t^{1:K} = \sum_{k=1}^K \mathbf{G}_t^k, \quad \mathbf{H}_t^{1:K} = \sum_{k=1}^K \mathbf{H}_t^k. \quad (12)$$

Once obtaining $\mathbf{G}_t^{1:K}$ and $\mathbf{H}_t^{1:K}$, the server can derive the global classifier \mathbf{W}_t via (13) and distribute it to all clients then. As shown in Section 3.3, regardless of how the data are distributed across clients, the global analytic classifier \mathbf{W}_t obtained through (13) is exactly identical to the optimal solution of centralized analytic learning on the full dataset, thus achieving invariance to data heterogeneity.

$$\begin{aligned}\mathbf{W}_t &= [(\Phi_t^{1:K})^\top \Phi_t^{1:K} + \lambda \mathbf{I}]^{-1} (\Phi_t^{1:K})^\top \mathbf{Y}^{1:K} \\ &= (\mathbf{G}_t^{1:K} + \lambda \mathbf{I})^{-1} \mathbf{H}_t^{1:K}.\end{aligned}\quad (13)$$

Then, each client is required to compute its hidden random features \mathbf{F}_t^k and local residual matrix \mathbf{R}_t^k based on the previously obtained features Φ_t^k and classifiers \mathbf{W}_t , as follows.

$$\mathbf{F}_t^k = \sigma(\Phi_t^k \mathbf{B}_t), \quad \mathbf{R}_t^k = \mathbf{Y}^k - (\Phi_t^k \mathbf{W}_t). \quad (14)$$

Subsequently, each client utilizes the hidden random features \mathbf{F}_t^k and local residual matrix \mathbf{R}_t^k to get its local *Hidden Auto-Correlation Matrix* Π_t^k and *Residual Cross-Correlation Matrix* Υ_t^k via (15).

$$\Pi_t^k = (\mathbf{F}_t^k)^\top \mathbf{F}_t^k, \quad \Upsilon_t^k = (\mathbf{F}_t^k)^\top \mathbf{R}_t^k. \quad (15)$$

Then, akin to \mathbf{G}_t^k and \mathbf{H}_t^k , the server aggregates $\Pi_t^{1:K}$ and $\Upsilon_t^{1:K}$ from the clients via (16).

$$\Pi_t^{1:K} = \sum_{k=1}^K \Pi_t^k, \quad \Upsilon_t^{1:K} = \sum_{k=1}^K \Upsilon_t^k. \quad (16)$$

Once obtaining $\Pi_t^{1:K}$ and $\Upsilon_t^{1:K}$, the server can derive Ω_{t+1} via (17) and distribute it to all clients.

$$\begin{aligned}\Omega_{t+1} &= \mathbf{V}_t [(\mathbf{V}_t^\top (\mathbf{F}_t^{1:K})^\top \mathbf{R}_t^{1:K} \mathbf{W}_t^\top \mathbf{U}_t) \oslash (\gamma \mathbf{1} + \text{diag}(\Lambda_t^F) \otimes \text{diag}(\Lambda_t^W))] \mathbf{U}_t^\top \\ &= \mathbf{V}_t [(\mathbf{V}_t^\top \Upsilon_t^{1:K} \mathbf{W}_t^\top \mathbf{U}_t) \oslash (\gamma \mathbf{1} + \text{diag}(\Lambda_t^F) \otimes \text{diag}(\Lambda_t^W))] \mathbf{U}_t^\top,\end{aligned}\quad (17)$$

where \mathbf{V}_t , Λ_t^F , \mathbf{U}_t , and Λ_t^W are obtained by spectral decompositions, as follows.

$$\Pi_t^{1:K} = (\mathbf{F}_t^{1:K})^\top \mathbf{F}_t^{1:K} = \mathbf{V}_t \Lambda_t^F \mathbf{V}_t^\top, \quad \mathbf{W}_t \mathbf{W}_t^\top = \mathbf{U}_t \Lambda_t^W \mathbf{U}_t^\top. \quad (18)$$

After receiving Ω_{t+1} from the server, each client can thus update the next-layer features Φ_{t+1}^k as:

$$\Phi_{t+1}^k = \Phi_t^k + \mathbf{F}_t^k \Omega_{t+1}. \quad (19)$$

The preceding process continues layer-by-layer until the final-layer classifier \mathbf{W}_T is completed. For clarity, we summarize the detailed training and inference procedures of our DeepAFL in Algorithm 1 and Algorithm 2, respectively, within Appendix A. In the next subsection, we will comprehensively provide theoretical analyses for our DeepAFL, including its validity, privacy, and efficiency.

3.3 THEORETICAL ANALYSES

Validity Analyses: Here, we provide detailed analyses of the validity of our DeepAFL. Specifically, we first derive the analytical solutions of our DeepAFL in Lemmas 1–2, and then demonstrate our DeepAFL’s dual properties of heterogeneity invariance and representation learning in Theorems 1–2.

Lemma 1: For any least squares problem with the following form:

$$\mathbf{W}^* = \arg \min_{\mathbf{W}} \|\mathbf{Y} - \Phi \mathbf{W}\|_F^2 + \lambda \|\mathbf{W}\|_F^2, \quad (20)$$

it yields a distinct analytical (i.e., closed-form) solution, which can be formulated as:

$$\mathbf{W}^* = (\Phi^\top \Phi + \lambda \mathbf{I})^{-1} \Phi^\top \mathbf{Y}. \quad (21)$$

Proof. See Appendix B.1 for details.

Lemma 2: For any sandwiched least squares problem with the following form:

$$\Omega^* = \arg \min_{\Omega} \|\mathbf{R} - \mathbf{F} \Omega \mathbf{W}\|_F^2 + \gamma \|\Omega\|_F^2, \quad (22)$$

it yields a distinct analytical (i.e., closed-form) solution, which can be formulated as:

$$\Omega^* = \mathbf{V} [(\mathbf{V}^\top \mathbf{F}^\top \mathbf{R} \mathbf{W}^\top \mathbf{U}) \oslash (\gamma \mathbf{1} + \text{diag}(\Lambda^F) \otimes \text{diag}(\Lambda^W))] \mathbf{U}^\top, \quad (23)$$

where $\mathbf{F}^\top \mathbf{F} = \mathbf{V} \Lambda^F \mathbf{V}^\top$ and $\mathbf{W} \mathbf{W}^\top = \mathbf{U} \Lambda^W \mathbf{U}^\top$ are spectral decompositions, while \oslash and \otimes represent element-wise division and outer product.

Proof. See Appendix B.1 for details.

324 **Theorem 1 (Invariance to Data Heterogeneity):** In the FL scenario, let \mathcal{D} be the full dataset, and
 325 $\mathcal{P} = \{\mathcal{D}^1, \mathcal{D}^2, \dots, \mathcal{D}^K\}$ be any heterogeneous partition of \mathcal{D} among K clients, where $\mathcal{D}_i \cap \mathcal{D}_j = \emptyset$
 326 for $i \neq j$ and $\bigcup_{i=1}^N \mathcal{D}_i = \mathcal{D}$. Given fixed seeds for the random projection matrices \mathbf{A} and $\{\mathbf{B}_t\}_{t=0}^{T-1}$,
 327 the global model weights $\{\mathbf{W}_t\}_{t=0}^T$ and $\{\Omega_t\}_{t=1}^T$ derived from DeepAFL are invariant to any data
 328 partition \mathcal{P} with different heterogeneity, being identical to the centralized analytical solutions on \mathcal{D} .
 329

330 *Proof.* See Appendix B.2 for details.

331 **Theorem 2 (Capability of Representation Learning):** Let the empirical risk $\mathcal{H}(\Phi, \mathbf{W})$, i.e., the
 332 loss function, for a given feature representation Φ and a classifier \mathbf{W} be defined as:

$$\mathcal{H}(\Phi, \mathbf{W}) = \|\mathbf{Y} - \Phi \mathbf{W}\|_F^2, \quad (24)$$

333 where \mathbf{Y} denotes the ground truth labels. Our DeepAFL yields a sequence of feature-classifier pairs
 334 (Φ_t, \mathbf{W}_t) at each layer $t \in [0, T]$. When setting regularization parameters γ and λ to 0, the sequence
 335 of the empirical risks $\{\mathcal{H}(\Phi_t, \mathbf{W}_t)\}_{t=0}^T$ in our DeepAFL keeps monotonically non-increasing, i.e.,
 336

$$\mathcal{H}(\Phi_t, \mathbf{W}_t) \geq \mathcal{H}(\Phi_{t+1}, \mathbf{W}_{t+1}), \forall t \in [0, T]. \quad (25)$$

337 Furthermore, as the number of layers T within our DeepAFL increases (i.e., $T \rightarrow \infty$), the sequence
 338 of empirical risks is guaranteed to converge to a limit $\mathcal{H}^* \leq \mathcal{H}(\Phi_0, \mathbf{W}_0)$.
 339

340 *Proof.* See Appendix B.3 for details.

341 Notably, Theorem 2 demonstrates our DeepAFL’s representation learning capability under the basic
 342 condition of no regularization ($\gamma = \lambda = 0$). In practice, we often employ regularization to enhance
 343 generalization performance and numerical stability. Thus, we introduce Theorem 3 in Appendix B.3
 344 to extend our analysis to cover the regularized settings. In addition, we will further provide extensive
 345 empirical evidence in Section 4 to support these dual advantageous properties of our DeepAFL.
 346

347 **Privacy Analyses:** For the privacy of our DeepAFL, several previous studies have provided similar
 348 analyses for *Auto-Correlation* and *Cross-Correlation Matrices* (Tan et al., 2022b; He et al., 2025b;
 349 Fan et al., 2025b). Particularly, inferring each client’s raw data is challenging as the size of its local
 350 dataset N_k remains private. Furthermore, this aggregation process is a form of *Secure Multi-Party
 351 Computation*, and many existing protocols can be readily integrated into our DeepAFL to further
 352 enhance privacy (Bonawitz et al., 2016; 2017; So et al., 2023). See Appendix C for more details.

353 **Efficiency Analyses:** In our DeepAFL, the overall complexities of computation and communication
 354 for each client are $\mathcal{O}(TN_k(d_\Phi^2 + d_\Phi d_F + d_F^2))$ and $\mathcal{O}(T(d_\Phi^2 + d_F^2))$, while those for the central
 355 server are $\mathcal{O}(T(d_\Phi^3 + d_F^3 + N_k d_F^2))$ and $\mathcal{O}(TKC(d_\Phi + d_F))$, respectively. See Appendix D for the
 356 detailed derivations and analyses. Notably, in Section 4, we will show the superior efficiency of our
 357 DeepAFL in comparison to existing gradient-based baselines. In our later experiments, we default
 358 to selecting projection dimensions $d_\Phi = d_F = 1024$ to balance the effectiveness and efficiency.
 359

360 4 EXPERIMENTAL EVALUATIONS

361 4.1 EXPERIMENTAL SETUP

362 **Datasets & Settings.** We conduct our experiments on three prominent benchmark datasets in FL:
 363 the CIFAR-10 (Krizhevsky & Hinton, 2009), CIFAR-100 (Krizhevsky & Hinton, 2009), and Tiny-
 364 ImageNet (Le & Yang, 2015). To simulate diverse data heterogeneity scenarios in FL, we use two
 365 common non-IID partitioning settings: Latent Dirichlet Allocation (Lin et al., 2020) (as Non-IID-1)
 366 and Sharding (Lin et al., 2020) (as Non-IID-2). We use the parameters α and s to control the level
 367 of heterogeneity in these two non-IID settings, respectively. In both cases, smaller parameter values
 368 indicate more heterogeneous data distributions. Specifically, we set $\alpha \in \{0.1, 0.05\}$, $s \in \{2, 4\}$ for
 369 CIFAR-10, while setting $\alpha \in \{0.1, 0.01\}$, $s \in \{5, 10\}$ for CIFAR-100 and Tiny-ImageNet.
 370

371 **Baselines & Metrics.** We compare our DeepAFL against 7 traditional gradient-based baselines,
 372 including FedAvg (McMahan et al., 2017), FedProx (Li et al., 2020) and MOON (Li et al., 2021a),
 373 FedGen (Zhu et al., 2021), FedDyn (Acar et al., 2021), FedNTD (Lee et al., 2022), and FedDisco (Ye
 374 et al., 2023b). Moreover, we also include the analytic learning-based method AFL (He et al., 2025b)
 375 as a baseline to further highlight our advantages. For fair comparisons, we align the experimental
 376 benchmark with those of AFL (He et al., 2025b) and use the same pre-trained ResNet-18 backbone
 377 for all methods. We employ the accuracy (%), computational cost (s), and communication cost (MB)
 378 as the main metrics. See Appendix E.2 for more details of our experimental implementation.

378
379
380
381
382
383
384 Table 1: Performance comparisons of the top-1 accuracy (%) among our DeepAFL and the baselines,
385 on the CIFAR-100 and Tiny-ImageNet. The best result is highlighted in **bold**, and the second-best
386 result is underlined. All the experiments were conducted three times, and the results are shown as
387 Mean \pm Standard Error. All the improvements of our DeepAFL were validated by Chi-squared tests.
388 The results of all baselines are directly obtained from the given benchmark in AFL (He et al., 2025b).

Baseline	CIFAR-100				Tiny-ImageNet			
	Non-IID-1		Non-IID-2		Non-IID-1		Non-IID-2	
	$\alpha = 0.1$	$\alpha = 0.01$	$s = 10$	$s = 5$	$\alpha = 0.1$	$\alpha = 0.01$	$s = 10$	$s = 5$
FedAvg (2017)	56.62 \pm 0.12	32.99 \pm 0.20	55.76 \pm 0.13	48.33 \pm 0.15	46.04 \pm 0.27	32.63 \pm 0.19	39.06 \pm 0.26	29.66 \pm 0.19
FedProx (2020)	56.45 \pm 0.22	33.37 \pm 0.09	55.80 \pm 0.16	48.29 \pm 0.14	46.47 \pm 0.23	32.26 \pm 0.14	38.97 \pm 0.23	29.17 \pm 0.16
MOON (2021a)	56.58 \pm 0.02	33.34 \pm 0.11	55.70 \pm 0.25	48.34 \pm 0.19	46.21 \pm 0.14	32.38 \pm 0.20	38.79 \pm 0.14	29.24 \pm 0.30
FedGen (2021)	56.48 \pm 0.17	33.09 \pm 0.09	60.93 \pm 0.17	48.12 \pm 0.06	46.27 \pm 0.14	32.33 \pm 0.14	38.82 \pm 0.16	29.37 \pm 0.25
FedDyn (2021)	57.55 \pm 0.08	36.12 \pm 0.08	<u>61.09\pm0.09</u>	<u>59.34\pm0.11</u>	47.72 \pm 0.22	35.19 \pm 0.06	41.36 \pm 0.06	35.18 \pm 0.18
FedNTD (2022)	56.60 \pm 0.14	32.59 \pm 0.21	54.69 \pm 0.15	47.00 \pm 0.19	46.17 \pm 0.16	31.86 \pm 0.44	37.55 \pm 0.09	29.01 \pm 0.14
FedDisco (2023b)	55.79 \pm 0.04	25.72 \pm 0.08	54.65 \pm 0.09	45.86 \pm 0.18	47.48 \pm 0.06	27.15 \pm 0.10	38.86 \pm 0.12	27.72 \pm 0.18
AFL (2025b)	<u>58.56\pm0.00</u>	<u>58.56\pm0.00</u>	58.56 \pm 0.00	58.56 \pm 0.00	<u>54.67\pm0.00</u>	<u>54.67\pm0.00</u>	<u>54.67\pm0.00</u>	<u>54.67\pm0.00</u>
DeepAFL ($T = 5$)	64.72 \pm 0.07	64.72 \pm 0.07	64.72 \pm 0.07	64.72 \pm 0.07	60.31 \pm 0.04	60.31 \pm 0.04	60.31 \pm 0.04	60.31 \pm 0.04
DeepAFL ($T = 10$)	65.96 \pm 0.05	65.96 \pm 0.05	65.96 \pm 0.05	65.96 \pm 0.05	61.37 \pm 0.07	61.37 \pm 0.07	61.37 \pm 0.07	61.37 \pm 0.07
DeepAFL ($T = 20$)	66.98\pm0.04	66.98\pm0.04	66.98\pm0.04	66.98\pm0.04	62.35\pm0.01	62.35\pm0.01	62.35\pm0.01	62.35\pm0.01
Improvement \uparrow	8.42 $p < 0.05$	8.42 $p < 0.05$	5.89 $p < 0.05$	7.64 $p < 0.05$	7.68 $p < 0.05$			

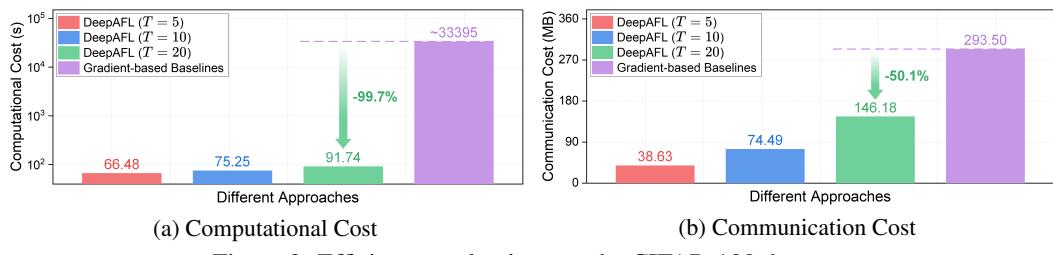


Figure 2: Efficiency evaluations on the CIFAR-100 dataset.

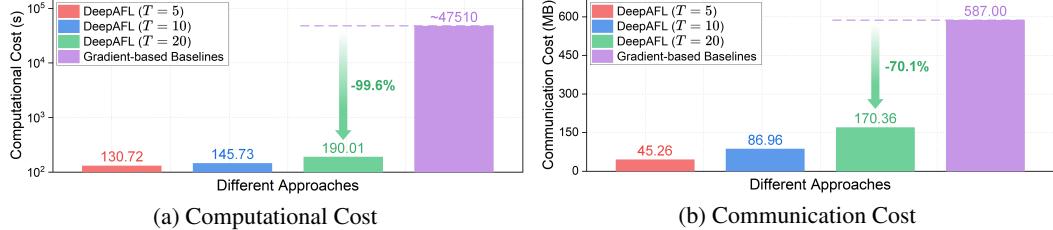


Figure 3: Efficiency evaluations on the Tiny-ImageNet dataset.

4.2 EXPERIMENTAL RESULTS

418 **Main Comparisons.** Here, we compare our DeepAFL extensively against baselines across various
419 datasets and settings, as shown in Tables 1 and 2. For a fair comparison, we included the benchmark
420 of the AFL (He et al., 2025b) and show our DeepAFL’s performance under $T \in \{5, 10, 20\}$. Notably,
421 even with a small value of $T = 5$, our DeepAFL already surpasses all SOTA baselines, thanks to its
422 dual advantages in heterogeneity invariance and representation learning. As T increases ($5 \rightarrow 20$),
423 the testing accuracy of our DeepAFL consistently enhances and eventually reaches 86.43%, 66.98%,
424 and 62.35% for the CIFAR-10, CIFAR-100, and Tiny-ImageNet, respectively, achieving impressive
425 improvements of up to 5.68%–8.42% over the SOTA baselines. See Appendix E.3.1 for more details.

426 **Invariance Analyses.** In Tables 1 and 2, our DeepAFL shows invariant performance across various
427 Non-IID settings on the same dataset, akin to the AFL. In contrast, other gradient-based baselines
428 commonly suffer performance degradation as the degree of heterogeneity increases. This invariance
429 property of our DeepAFL can be further extended to the number of clients. As shown in Figure 7, the
430 superiority of our DeepAFL over the gradient-based baseline progressively widens with an increased
431 number of clients ($100 \rightarrow 1000$) for the large-scale scenarios. These results empirically validate the
432 ideal invariance of our DeepAFL as shown in Theorem 1. See Appendix E.3.2 for more details.

432 **Representation Analyses.** Another key advantage of our DeepAFL is its capability of deep representation learning. As shown in Tables 9–11, our DeepAFL exhibits a stable improvement in both
 433 training and testing accuracy as the number of layers T increases. This consistent gain shows its
 434 effective ability to learn deep representations, handling the common issue of underfitting in traditional
 435 analytic learning. Moreover, while our main results only report our DeepAFL’s performance
 436 at $T \in \{5, 10, 20\}$, we observe its performance continues to improve as T further increases. Specif-
 437 ically, on the complex CIFAR-100 and Tiny-ImageNet, our DeepAFL’s testing accuracy at $T = 50$
 438 improves by over 1.5% compared to that at $T = 20$. See Appendix E.3.3 for more details.
 439

440 **Efficiency Evaluations.** We further give efficiency evaluations to show DeepAFL’s superior balance
 441 between accuracy and efficiency. As shown in Figures 2–3, we show that, in addition to its SOTA
 442 performance, our DeepAFL also achieves superior efficiency compared to existing gradient-based
 443 baselines. In Tables 9–11 and Figures 4–6, the total time cost of DeepAFL for 100 clients shows a
 444 minimal marginal increase with each additional layer, adding 1–2s per layer for the CIFAR-10 and
 445 CIFAR-100, and 3s per layer for the Tiny-ImageNet. Thus, when set to $T = 50$, our DeepAFL can
 446 achieve more than 9% performance improvements compared to AFL on the complex CIFAR-100 and
 447 Tiny-ImageNet, with a time cost that is less than twice that of AFL, as the primary time-consuming
 448 backbone’s forward pass only needs to be performed once. See Appendix E.3.4 for more details.
 449

450 **Parameter Analyses.** Then, we provide comprehensive parameter sensitivity analyses of our Deep-
 451 AFL, including the regularization parameters λ , γ , the activation function $\sigma(\cdot)$, and the projection
 452 dimensions d_Φ , d_F . Specifically, as shown in Tables 3–4, for the regularization parameters, our Deep-
 453 AFL is insensitive to λ but more sensitive to γ . For the simple CIFAR-10, we observe satisfactory
 454 performance with $\gamma \in [0.1, 0.5]$. For the more complex datasets, CIFAR-100 and Tiny-ImageNet,
 455 the optimal value for γ is a smaller 0.01. In Table 5, the activation function exhibits a notable im-
 456 pact on performance, with a difference of about 2%. Here, GELU is the optimal activation function,
 457 while Softshrink is the poorest yet still beats the model with no activation. As shown in Tables 6–8,
 458 larger dimensions for both d_Φ and d_F lead to better performance. Nonetheless, overly large dimen-
 459 sions can not only incur significant overhead due to their quadratic complexity, but also impact the
 460 numerical stability, causing crashing when $d_\Phi = d_F = 2^{13}$. See Appendix E.3.5 for more details.
 461

462 **Ablation Studies.** Finally, we conduct extensive ablation studies for the individual contributions of
 463 our DeepAFL’s key components. Specifically, we compared our DeepAFL and four distinct ablation
 464 models across different layers on various datasets, as shown in Tables 12–17 and Figures 8–10. First,
 465 we observe that the residual skip connection in our DeepAFL has a dual role: it not only accelerates
 466 performance improvement as the number of layers T increases, but also enhances the final converged
 467 performance. Second, ablating the random projection \mathbf{B}_t of each layer with an identity matrix shows
 468 that the **stochasticity** provided by \mathbf{B}_t is vital for DeepAFL to escape local optima and saddle points.
 469 Third, ablating the activation function $\sigma(\cdot)$ proves the necessity of its **nonlinearity** in our DeepAFL.
 470 Fourth, ablating the trainable transformation Ω_{t+1} renders the ablation model ineffective, indicating
 471 that its provided **learnability** is also indispensable. See Appendix E.3.6 for more details. Meanwhile,
 472 we also provide comparisons of our DeepAFL against other deepening strategies in Appendix E.3.7.
 473

5 CONCLUSION AND DISCUSSION

474 In this paper, we propose DeepAFL, as the first attempt in FL to achieve gradient-free representation
 475 learning while preserving strong invariance to data heterogeneity. By identifying and addressing a
 476 fundamental limitation in existing analytic-learning-based approaches, our DeepAFL exhibits dual
 477 advantages in both heterogeneity invariance and representation learning. By virtue of its gradient-
 478 free nature, our DeepAFL also achieves high efficiency by eliminating costly iterative computations
 479 and communications on gradients. Given its theoretical and empirical superiority, we believe our
 480 DeepAFL means a great advancement for the SOTA in the fields of both analytic learning and FL.
 481

482 For future work, a natural and promising direction is to extend our deep residual analytic models in
 483 our DeepAFL to other well-suited fields for analytic learning. Specifically, we will go on to aim at
 484 continual learning, where all data arrives online and historical data is not accessible. While analytic
 485 learning has been proven to be effective in solving the catastrophic forgetting problem of continual
 486 learning, applying our DeepAFL to this field remains non-trivial. This is because the features learned
 487 from a previous phase may become invalid in a subsequent one if updating the weights layer by layer.
 488 As a result, we would need a new approach to incorporate the phase-wise recursions as well. Some
 489 further discussion is provided in Appendix F, and our usage of LLMs is declared in Appendix G.
 490

486 ETHICS STATEMENT
487488 The authors have read and adhere to the ICLR Code of Ethics. We have carefully considered the
489 potential ethical implications of our work and believe that our proposed approach, DeepAFL, aligns
490 with the principles outlined in the provided ICLR Code of Ethics.491 Our research is situated within the FL paradigm, which is specifically designed to address privacy
492 and security concerns by enabling collaborative model training without requiring the sharing of raw,
493 sensitive user data. Our proposed DeepAFL operates exclusively in a decentralized setting, with
494 all local client data remaining on the respective devices. We provide detailed privacy analyses in
495 our paper and discuss how our DeepAFL can be enhanced with existing proven privacy-preserving
496 protocols. The empirical evaluations presented in this paper are based on well-established, publicly
497 available benchmark datasets (i.e., CIFAR-10, CIFAR-100, and Tiny-ImageNet). This work did not
498 involve any human subjects, the collection of new private data, or the use of sensitive information.499 Furthermore, a core contribution of our proposed DeepAFL is its superior performance in FL, with
500 its dual advantages in heterogeneity invariance and representation learning. To the best of our knowl-
501 edge, the process of our DeepAFL does not introduce additional fairness and bias issues. On the
502 contrary, its ability to achieve superior and invariant performance across various heterogeneous FL
503 environments can help foster more equitable and generalizable models in real-world applications.504 Lastly, in line with the principles of research reproducibility, we are committed to making our code
505 open-source upon the paper’s acceptance. We believe this will allow the broader research community
506 to verify our findings, build upon our work, and ensure full transparency of our methodology.508 REPRODUCIBILITY STATEMENT
509510 To ensure the full reproducibility of our research, we have made extensive efforts to provide all the
511 necessary details. We provide detailed procedures of our proposed approach in our paper, including
512 comprehensive descriptions of the algorithmic steps and design choices. For our theoretical claims,
513 the complete proofs and analyses are included in the Appendix. The experimental results, including
514 hyperparameters and data preprocessing steps for the benchmark datasets, are elaborated in the
515 Appendix. Upon acceptance, our codes will be made publicly available as open-source, allowing
516 the broader research community to verify our findings and build upon our work.518 REFERENCES
519520 Durmus Alp Emre Acar, Yue Zhao, Ramon Matas, Matthew Mattina, Paul Whatmough, and
521 Venkatesh Saligrama. Federated learning based on dynamic regularization. In International
522 Conference on Learning Representations, 2021. URL <https://openreview.net/forum?id=B7v4QMR6Z9w>.524 Guillaume Alain and Yoshua Bengio. Understanding intermediate layers using linear classifier
525 probes. In International Conference on Learning Representations, 2016.527 Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new
528 perspectives. IEEE Transactions on Pattern Analysis and Machine Intelligence, 35(8):1798–1828,
529 2013. doi: 10.1109/TPAMI.2013.50.530 Erik L Bolager, Iryna Burak, Chinmay Datar, Qing Sun, and Felix Dietrich. Sampling weights of
531 deep neural networks. Advances in Neural Information Processing Systems, 36:63075–63116,
532 2023.533 Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar
534 Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for federated
535 learning on user-held data. arXiv preprint arXiv:1611.04482, 2016.537 Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar
538 Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-
539 preserving machine learning. In proceedings of the 2017 ACM SIGSAC Conference on Computer
and Communications Security (CCS), pp. 1175–1191, 2017.

540 Di Chai, Leye Wang, Liu Yang, Junxue Zhang, Kai Chen, and Qiang Yang. A survey for fed-
 541 erated learning evaluations: Goals and measures. *IEEE Transactions on Knowledge and Data
 542 Engineering*, 36(10):5007–5024, 2024. doi: 10.1109/TKDE.2024.3382002.

543

544 Randall E Cline. Representations for the generalized inverse of a partitioned matrix. *Journal of the
 545 Society for Industrial and Applied Mathematics*, 12(3):588–600, 1964.

546 Liam Collins, Hamed Hassani, Aryan Mokhtari, and Sanjay Shakkottai. Exploiting shared repre-
 547 sentations for personalized federated learning. In *International Conference on Machine Learning*,
 548 pp. 2089–2099. PMLR, 2021.

549

550 Thomas M Cover. Geometrical and statistical properties of systems of linear inequalities with ap-
 551 plications in pattern recognition. *IEEE transactions on electronic computers*, (3):326–334, 2006.

552

553 Feng Ding, Peter X. Liu, and Jie Ding. Iterative solutions of the generalized sylvester ma-
 554 trix equations by using the hierarchical identification principle. *Applied Mathematics and
 555 Computation*, 197(1):41–50, 2008. ISSN 0096-3003. doi: <https://doi.org/10.1016/j.amc.2007.07.040>. URL <https://www.sciencedirect.com/science/article/pii/S0096300307007515>.

556

557

558 Guang-Ren Duan. *Generalized Sylvester equations: unified parametric solutions*. Crc Press, 2015.

559

560 Kejia Fan, Yajiang Huang, Jingyu He, Feijiang Han, Jianheng Tang, Huiping Zhuang, Anfeng
 561 Liu, Tian Wang, Mianxiong Dong, Houbing Herbert Song, and Yunhuai Liu. CALM: A ubiq-
 562 uitous crowdsourced analytic learning mechanism for continual service construction with data
 563 privacy preservation. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous
 564 Technologies*, 9(2), 2025a. doi: 10.1145/3729473.

565

566 Kejia Fan, Jianheng Tang, Zhirui Yang, Feijiang Han, Jiaxu Li, Run He, Yajiang Huang, Anfeng
 567 Liu, Houbing Herbert Song, Yunhuai Liu, and Huiping Zhuang. APFL: Analytic personalized
 568 federated learning via dual-stream least squares. *arXiv preprint*, 2025b.

569

570 Tao Fan, Hanlin Gu, Xuemei Cao, Chee Seng Chan, Qian Chen, Yiqiang Chen, Yihui Feng, Yang
 571 Gu, Jiaxiang Geng, Bing Luo, Shuoling Liu, Win Kent Ong, Chao Ren, Jiaqi Shao, Chuan Sun,
 572 Xiaoli Tang, Hong Xi Tae, Yongxin Tong, Shuyue Wei, Fan Wu, Wei Xi, Mingcong Xu, He Yang,
 573 Xin Yang, Jiangpeng Yan, Hao Yu, Han Yu, Teng Zhang, Yifei Zhang, Xiaojin Zhang, Zhenzhe
 574 Zheng, Lixin Fan, and Qiang Yang. Ten challenging problems in federated foundation models,
 575 2025c. URL <https://arxiv.org/abs/2502.12176>.

576

577 Eros Fanì, Raffaello Camoriano, Barbara Caputo, and Marco Ciccone. Accelerating heterogeneous
 578 federated learning with closed-form classifiers. In *International Conference on Machine Learning*,
 579 pp. 13029–13048. PMLR, 2024.

580

581 Ping Guo, Michael R Lyu, and NE Mastorakis. Pseudoinverse learning algorithm for feedforward
 582 neural networks. *Advances in Neural Networks and Applications*, 1(321-326), 2001.

583

584 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-
 585 nition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition
 586 (CVPR)*, June 2016.

587

588 Run He, Di Fang, Yicheng Xu, Yawen Cui, Ming Li, Cen Chen, Ziqian Zeng, and Huiping Zhuang.
 589 Semantic shift estimation via dual-projection and classifier reconstruction for exemplar-free class-
 590 incremental learning. In *Forty-second International Conference on Machine Learning (ICML)*,
 591 2025a.

592

593 Run He, Kai Tong, Di Fang, Han Sun, Ziqian Zeng, Haoran Li, Tianyi Chen, and Huiping
 594 Zhuang. AFL: A single-round analytic approach for federated learning with pre-trained models. In
 595 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025b.

596

597 Chenghao Hu and Baochun Li. Maskcrypt: Federated learning with selective homomorphic en-
 598 cryption. *IEEE Transactions on Dependable and Secure Computing*, 22(1):221–233, 2025. doi:
 599 10.1109/TDSC.2024.3392424.

594 Rui Hu, Yuanxiong Guo, and Yanmin Gong. Federated learning with sparsified model perturba-
 595 tion: Improving accuracy under client-level differential privacy. *IEEE Transactions on Mobile*
 596 *Computing*, 23(8):8242–8255, 2024. doi: 10.1109/TMC.2023.3343288.

597 Like Hui and Mikhail Belkin. Evaluation of neural architectures trained with square loss vs cross-
 598 entropy in classification tasks. In *International Conference on Learning Representations*, 2021.
 599 URL <https://openreview.net/forum?id=hsFN92eQEl>.

600 601 Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images.
 602 *Technical Report*, 2009.

603 604 Guannan Lai, Yujie Li, Xiangkun Wang, Tianrui Li Junbo Zhang, and Xin Yang. Order-robust
 605 class incremental learning: Graph-driven dynamic similarity grouping. In *Proceedings of the*
 606 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.

607 608 Samuel Lanthaler and Nicholas H Nelsen. Error bounds for learning with vector-valued random
 609 features. *Advances in Neural Information Processing Systems*, 36:71834–71861, 2023.

610 611 Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.

612 613 Gihun Lee, Minchan Jeong, Yongjin Shin, Sangmin Bae, and Se-Young Yun. Preservation of the
 614 global knowledge by not-true distillation in federated learning. In Alice H. Oh, Alekh Agar-
 615 wal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing*
 616 *Systems*, 2022. URL <https://openreview.net/forum?id=qw3Mzb1Juo>.

617 618 Jiaxu Li, Rui Li, Jianyu Qi, Songning Lai, Linpu Lv, Kejia Fan, Jianheng Tang, Yutao Yue,
 619 Dongzhan Zhou, Yuanhuai Liu, and Huiping Zhuang. CFSSeg: Closed-form solution for class-
 620 incremental semantic segmentation of 2d images and 3d point clouds. *arXiv preprint*, 2025.

621 622 Qinbin Li, Bingsheng He, and Dawn Song. Model-contrastive federated learning. In *Proceedings*
 623 *of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10713–10722,
 624 2021a.

625 626 Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith.
 627 Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and*
 628 *Systems*, 2:429–450, 2020.

629 630 Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. Ditto: Fair and robust federated
 631 learning through personalization. In *International Conference on Machine Learning*, pp. 6357–
 632 6368. PMLR, 2021b.

633 634 Xinting Liao, Weiming Liu, Chaochao Chen, Pengyang Zhou, Fengyuan Yu, Huabin Zhu, Binhui
 635 Yao, Tao Wang, Xiaolin Zheng, and Yanchao Tan. Rethinking the representation in federated un-
 636 supervised learning with non-iid data. In *Proceedings of the IEEE/CVF Conference on Computer*
 637 *Vision and Pattern Recognition*, pp. 22841–22850, 2024.

638 639 Tao Lin, Lingjing Kong, Sebastian U Stich, and Martin Jaggi. Ensemble distillation for robust model
 640 fusion in federated learning. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin
 641 (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 2351–2363. Cur-
 642 ran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/18df51b97cccd68128e994804f3ecc87-Paper.pdf.

643 644 Xiang Liu, Liangxi Liu, Feiyang Ye, Yunheng Shen, Xia Li, Linshan Jiang, and Jialin Li.
 645 FedLPA: One-shot federated learning with layer-wise posterior aggregation. *Advances in Neural*
 646 *Information Processing Systems*, 37:81510–81548, 2024a.

647 648 Yiqi Liu, Shan Chang, Ye Liu, Bo Li, and Cong Wang. FairFed: Improving fairness and efficiency of
 649 contribution evaluation in federated learning via cooperative shapley value. In *IEEE INFOCOM*
 650 *2024 - IEEE Conference on Computer Communications*, pp. 621–630, 2024b. doi: 10.1109/
 651 INFOCOM52122.2024.10621438.

652 653 Zichen Liu, Chao Du, Wee Sun Lee, and Min Lin. Locality sensitive sparse encoding for learning
 654 world models online. In *The Twelfth International Conference on Learning Representations*,
 655 2024c. URL <https://openreview.net/forum?id=i8PjQT3Uig>.

648 Cheng-Yaw Low, Jaewoo Park, and Andrew Beng-Jin Teoh. Stacking-based deep neural network:
 649 deep analytic network for pattern classification. *IEEE Transactions on Cybernetics*, 50(12):5021–
 650 5034, 2019.

651 Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera Y Ar-
 652 cas. Communication-efficient learning of deep networks from decentralized data. In *Artificial*
 653 *Intelligence and Statistics*, pp. 1273–1282. PMLR, 2017.

654 John Nguyen, Jianyu Wang, Kshitiz Malik, Maziar Sanjabi, and Michael Rabbat. Where to begin?
 655 on the impact of pre-training and initialization in federated learning. In *The Eleventh International*
 656 *Conference on Learning Representations*, 2023. URL [https://openreview.net/forum?](https://openreview.net/forum?id=Mpa3tRJFBb)
 657 [id=Mpa3tRJFBb](https://openreview.net/forum?id=Mpa3tRJFBb).

658 Liangzu Peng, Juan Elenter, Joshua Agterberg, Alejandro Ribeiro, and Rene Vidal. TSVD: Bridging
 659 theory and practice in continual learning with pre-trained models. In *The Thirteenth International*
 660 *Conference on Learning Representations*, 2025. URL [https://openreview.net/forum?](https://openreview.net/forum?id=bqv7M0wc4x)
 661 [id=bqv7M0wc4x](https://openreview.net/forum?id=bqv7M0wc4x).

662 Hongming Piao, Yichen Wu, Dapeng Wu, and Ying Wei. Federated continual learning via prompt-
 663 based dual knowledge transfer. In *Forty-first International Conference on Machine Learning*,
 664 2024. URL <https://openreview.net/forum?id=Kqa5JakTjB>.

665 Ameya Prabhu, Shiven Sinha, Ponnurangam Kumaraguru, Philip Torr, Ozan Sener, and Puneet
 666 Dokania. RanDumb: random representations outperform online continually learned represen-
 667 tations. *Advances in Neural Information Processing Systems*, 37:37988–38006, 2024.

668 Chao Ren, Han Yu, Hongyi Peng, Xiaoli Tang, Bo Zhao, Liping Yi, Alysa Ziying Tan, Yulan Gao,
 669 Anran Li, Xiaoxiao Li, Zengxiang Li, and Qiang Yang. Advances and open challenges in fed-
 670 erated foundation models. *IEEE Communications Surveys & Tutorials*, pp. 1–41, 2025. doi:
 671 10.1109/COMST.2025.3552524.

672 Xuankun Rong, Jianshu Zhang, Kun He, and Mang Ye. CAN: Leveraging clients as navigators for
 673 generative replay in federated continual learning. In *Forty-second International Conference on*
 674 *Machine Learning*, 2025.

675 Changlong Shi, He Zhao, Bingjie Zhang, Mingyuan Zhou, Dandan Guo, and Yi Chang. FedAWA:
 676 Adaptive optimization of aggregation weights in federated learning using client vectors. In
 677 *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 30651–30660,
 2025.

678 Jinhyun So, Ramy E Ali, Başak Güler, Jiantao Jiao, and A Salman Avestimehr. Securing secure
 679 aggregation: Mitigating multi-round privacy leakage in federated learning. In *Proceedings of the*
 680 *AAAI Conference on Artificial Intelligence*, volume 37, pp. 9864–9873, 2023.

681 Alysa Ziying Tan, Han Yu, Lizhen Cui, and Qiang Yang. Towards personalized federated learning.
 682 *IEEE transactions on Neural Networks and Learning Systems*, 34(12):9587–9603, 2022a.

683 Yue Tan, Guodong Long, Lu Liu, Tianyi Zhou, Qinghua Lu, Jing Jiang, and Chengqi Zhang. Fed-
 684 proto: Federated prototype learning across heterogeneous clients. *Proceedings of the AAAI*
 685 *Conference on Artificial Intelligence*, 36(8):8432–8440, Jun. 2022b. doi: 10.1609/aaai.v36i8.
 686 20819. URL <https://ojs.aaai.org/index.php/AAAI/article/view/20819>.

687 Jianheng Tang, Huiping Zhuang, Di Fang, Jiaxu Li, Feijiang Han, Yajiang Huang, Kejia Fan, Leye
 688 Wang, Zhanxing Zhu, Shanghang Zhang, Houbing Herbert Song, and Yunhuai Liu. ACU: An-
 689 alytic continual unlearning for efficient and exact forgetting with privacy preservation. *arXiv*
 690 *preprint*, 2025a.

691 Jianheng Tang, Huiping Zhuang, Jingyu He, Run He, Jingchao Wang, Kejia Fan, Anfeng Liu, Tian
 692 Wang, Leye Wang, Zhanxing Zhu, Shanghang Zhang, Houbing Herbert Song, and Yunhuai Liu.
 693 AFCL: Analytic federated continual learning for spatio-temporal invariance of Non-IID data.
 694 *arXiv preprint*, 2025b.

695 Kar-Ann Toh. Learning from the kernel and the range space. In *2018 IEEE/ACIS 17th International*
 696 *Conference on Computer and Information Science (ICIS)*, pp. 1–6. IEEE, 2018.

702 Quyen Tran, Tung Lam Tran, Khanh Doan, Toan Tran, Dinh Phung, Khoat Than, and Trung Le.
 703 Boosting multiple views for pretrained-based continual learning. In *The Thirteenth International*
 704 *Conference on Learning Representations*, 2025. URL [https://openreview.net/forum?](https://openreview.net/forum?id=AZR4R31w7y)
 705 [id=AZR4R31w7y](https://openreview.net/forum?id=AZR4R31w7y).

706 V.N. Vapnik. An overview of statistical learning theory. *IEEE Transactions on Neural Networks*, 10
 707 (5):988–999, 1999. doi: 10.1109/72.788640.

708 Ziyao Wang, Jianyu Wang, and Ang Li. FedHyper: A universal and robust learning rate
 709 scheduler for federated learning with hypergradient descent. In *The Twelfth International*
 710 *Conference on Learning Representations*, 2024. URL [https://openreview.net/forum?](https://openreview.net/forum?id=K19CqKf7h6)
 711 [id=K19CqKf7h6](https://openreview.net/forum?id=K19CqKf7h6).

712 Ai-Guo Wu, Guang-Ren Duan, and Bin Zhou. Solution to generalized sylvester matrix equations.
 713 *IEEE Transactions on Automatic Control*, 53(3):811–815, 2008. doi: 10.1109/TAC.2008.919562.

714 Panlong Wu, Kangshuo Li, Ting Wang, Yanjie Dong, Victor C. M. Leung, and Fangxin Wang.
 715 FedFMSL: Federated learning of foundation models with sparsely activated lora. *IEEE*
 716 *Transactions on Mobile Computing*, 23(12):15167–15181, 2024. doi: 10.1109/TMC.2024.
 717 [3454634](https://doi.org/10.1109/TMC.2024.3454634).

718 Xiyuan Yang, Wenke Huang, and Mang Ye. FedAS: Bridging inconsistency in personalized fed-
 719 erated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
 720 *Recognition*, pp. 11986–11995, 2024.

721 Zhiqin Yang, Yonggang Zhang, Yu Zheng, Xinmei Tian, Hao Peng, Tongliang Liu, and Bo Han.
 722 FedFed: Feature distillation against data heterogeneity in federated learning. In *Advances in*
 723 *Neural Information Processing Systems*, volume 36, pp. 60397–60428. Curran Associates, Inc.,
 724 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/bdcdf38389d7fcefc73c4c3720217155-Paper-Conference.pdf.

725 Mang Ye, Xiuwen Fang, Bo Du, Pong C. Yuen, and Dacheng Tao. Heterogeneous federated learning:
 726 State-of-the-art and research challenges. *ACM Computing Survey*, 56(3), October 2023a. ISSN
 727 0360-0300. doi: 10.1145/3625558. URL <https://doi.org/10.1145/3625558>.

728 Rui Ye, Mingkai Xu, Jianyu Wang, Chenxin Xu, Siheng Chen, and Yanfeng Wang. FedDisco:
 729 Federated learning with discrepancy-aware collaboration. In Andreas Krause, Emma Brunskill,
 730 Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of*
 731 *the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine*
 732 *Learning Research*, pp. 39879–39902. PMLR, 23–29 Jul 2023b.

733 Hao Yu, Xin Yang, Xin Gao, Yan Kang, Hao Wang, Junbo Zhang, and Tianrui Li. Personalized
 734 federated continual learning via multi-granularity prompt. In *Proceedings of the 30th ACM*
 735 *SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD ’24, pp. 4023–4034,
 736 New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704901. doi:
 737 10.1145/3637528.3671948. URL <https://doi.org/10.1145/3637528.3671948>.

738 Hao Yu, Xin Yang, Le Zhang, Hanlin Gu, Tianrui Li, Lixin Fan, and Qiang Yang. Handling spatial-
 739 temporal data heterogeneity for federated continual learning via tail anchor. In *Proceedings of the*
 740 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.

741 Xiang Zhang, Run He, Chen Jiao, Di Fang, Ming Li, Ziqian Zeng, Cen Chen, and Huiping
 742 Zhuang. L3A: Label-augmented analytic adaptation for multi-label class incremental learn-
 743 ing. In *Forty-second International Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=bBPnGYbypF>.

744 Xiaokang Zhou, Qiuyue Yang, Xuzhe Zheng, Wei Liang, Kevin I-Kai Wang, Jianhua Ma, Yi Pan,
 745 and Qun Jin. Personalized federated learning with model-contrastive learning for multi-modal
 746 user modeling in human-centric metaverse. *IEEE Journal on Selected Areas in Communications*,
 747 42(4):817–831, 2024.

748 Zhuangdi Zhu, Junyuan Hong, and Jiayu Zhou. Data-free knowledge distillation for heterogeneous
 749 federated learning. In *International Conference on Machine Learning*, pp. 12878–12889. PMLR,
 750 2021.

756 Huiping Zhuang, Zhiping Lin, and Kar-Ann Toh. Correlation projection for analytic learning of a
757 classification network. *Neural Processing Letters*, 53(6):3893–3914, 2021.

758

759 Huiping Zhuang, Zhenyu Weng, Hongxin Wei, RENCHUNZI Xie, Kar-Ann Toh, and Zhiping Lin.
760 ACIL: Analytic class-incremental learning with absolute memorization and privacy protection.
761 *Advances in Neural Information Processing Systems*, 35:11602–11614, 2022.

762

763 Huiping Zhuang, Zhenyu Weng, Run He, Zhiping Lin, and Ziqian Zeng. GKEAL: Gaussian kernel
764 embedded analytic learning for few-shot class incremental task. In *Proceedings of the IEEE/CVF*
765 conference on computer vision and pattern recognition, pp. 7746–7755, 2023.

766

767 Huiping Zhuang, Yizhu Chen, Di Fang, Run He, Kai Tong, Hongxin Wei, Ziqian Zeng, and
768 Cen Chen. GACL: Exemplar-free generalized analytic continual learning. *Advances in Neural*
769 *Information Processing Systems*, 37:83024–83047, 2024a.

770

771 Huiping Zhuang, Yuchen Liu, Run He, Kai Tong, Ziqian Zeng, Cen Chen, Yi Wang, and Lap-Pui
772 Chau. F-OAL: Forward-only online analytic learning with fast training and low memory footprint
773 in class incremental learning. In *The Thirty-eighth Annual Conference on Neural Information*
774 *Processing Systems*, 2024b.

775

776 Huiping Zhuang, Zhiping Lin, Yimin Yang, and Kar-Ann Toh. An analytic formulation of convolutional
777 neural network learning for pattern recognition. *Information Sciences*, 686:121317, 2025.
778 ISSN 0020-0255.

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810 A APPENDIX FOR PROCEDURES OF OUR DEEPAFL
811812 **Algorithm 1** The Training Procedure of our Proposed DeepAFL

813 **Input:** The clients' local datasets $\{\mathcal{D}^1, \mathcal{D}^2, \dots, \mathcal{D}^K\}$.
 814 **Output:** The transformation matrices $\{\Omega_t\}_{t=1}^T$ and the analytic classifiers $\{\mathbf{W}_t\}_{t=0}^T$.
 815 1: **for** each layer $t \in \{0, 1, 2, \dots, T\}$ **do**
 816 2: // (a) Local Computation for Analytic Classifier Construction (Client side)
 817 3: **for** each client $k \in \{1, 2, \dots, K\}$ **do**
 818 4: **if** constructing the first layer, i.e., $t = 0$ **then**
 819 5: Extract and construct its local zero-layer feature matrix Φ_0^k via (10)
 820 6: Compute its local *Feature Auto-Correlation Matrix* \mathbf{G}_t^k using Φ_t^k via (11);
 821 7: Compute its local *Label Cross-Correlation Matrix* \mathbf{H}_t^k using Φ_t^k and \mathbf{Y}_t^k via (11);
 822 8: Transmit $\{\mathbf{G}_t^k, \mathbf{H}_t^k\}$ to the server;
 823 9: // (b) Information Aggregation for Analytic Classifier Construction (Server side)
 824 10: Aggregate $\{\mathbf{G}_t^k\}_{k=1}^K$ and $\{\mathbf{H}_t^k\}_{k=1}^K$ to obtain $\mathbf{G}_t^{1:K}$ and $\mathbf{H}_t^{1:K}$ via (12);
 825 11: Derive the global classifier \mathbf{W}_t using the obtained $\mathbf{G}_t^{1:K}$ and $\mathbf{H}_t^{1:K}$ via (13);
 826 12: Transmit the global classifier \mathbf{W}_t to all clients;
 827 13: // (c) Local Computation for Residual Block Construction (Client side)
 828 14: **for** each client $k \in \{1, 2, \dots, K\}$ **do**
 829 15: **if** layer requirement not satisfied, i.e., $t < T$ **then**
 830 16: Compute its local hidden random feature \mathbf{F}_t^k using Φ_t^k via (14);
 831 17: Compute its local residual matrix \mathbf{R}_t^k using Φ_t^k and \mathbf{W}_t via (14);
 832 18: Compute its local *Hidden Auto-Correlation Matrix* Π_t^k using \mathbf{F}_t^k via (15);
 833 19: Compute its local *Residual Cross-Correlation Matrix* Υ_t^k using \mathbf{F}_t^k and \mathbf{R}_t^k via (15);
 834 20: Transmit $\{\Pi_t^k, \Upsilon_t^k\}$ to the server;
 835 21: // (d) Information Aggregation for Residual Block Construction (Server side)
 836 22: **if** layer requirement not satisfied, i.e., $t < T$ **then**
 837 23: Aggregate $\{\Pi_t^k\}_{k=1}^K$ and $\{\Upsilon_t^k\}_{k=1}^K$ to obtain $\Pi_t^{1:K}$ and $\Upsilon_t^{1:K}$ via (16);
 838 24: Derive the transformation matrix Ω_{t+1} using $\Pi_t^{1:K}$ and $\Upsilon_t^{1:K}$ via (17) and (18);
 839 25: Transmit the transformation matrix Ω_{t+1} to all clients;
 840 26: // (e) Feature Updating for Next Layer Construction (Client side)
 841 27: **if** layer requirement not satisfied, i.e., $t < T$ **then**
 842 28: **for** each client $k \in \{1, 2, \dots, K\}$ **do**
 843 29: Update its local feature matrix Φ_t^k to obtain Φ_{t+1}^k via (19);
 844 30: **Return:** The transformation matrices $\{\Omega_t\}_{t=1}^T$ and the analytic classifiers $\{\mathbf{W}_t\}_{t=0}^T$.

850 **Algorithm 2** The Inference Procedure of our Proposed DeepAFL

851 **Input:** Local sample \mathbf{x}^i , transformation matrices $\{\Omega_t\}_{t=1}^T$, and analytic classifiers $\{\mathbf{W}_t\}_{t=0}^T$.
 852 **Output:** Predicted label \hat{y}^i .
 853 1: Extract and construct its local zero-layer feature Φ_0^i using \mathbf{x}^i via (10)
 854 2: **for** each layer $t \in \{1, 2, \dots, T\}$ **do**
 855 3: Compute its local hidden random feature \mathbf{F}_t^i via (14);
 856 4: Update its local feature Φ_{t-1}^i to obtain Φ_t^i using Ω_t and \mathbf{F}_t^i via (19);
 857 5: Calculate the predicted score vector $\hat{\mathbf{y}}^i = \Phi_T^i \mathbf{W}_T$;
 858 6: Calculate the predicted label $\hat{y}^i = \arg \max(\hat{\mathbf{y}}^i)$;
 859 7: **Return:** Predicted label \hat{y}^i .

864 **B APPENDIX FOR VALIDITY ANALYSES**
 865

866 In this section, we thoroughly analyze the validity of our DeepAFL, theoretically establishing several key propositions. First of all, in Appendix B.1, we derived the analytical solutions for the least
 867 squares problems addressed in this paper, which also constitute the centralized analytical solution.
 868 Next, in Appendix B.2, we demonstrate that the deep residual analytic networks derived from our
 869 DeepAFL is equivalent to the centralized closed-form solutions and remain invariant to data hetero-
 870 geneity. Finally, in Appendix B.3, we validate DeepAFL’s representation learning capacity, proving
 871 that the empirical risk of our DeepAFL is monotonically non-increasing with increasing layer depth.
 872

873 **B.1 DERIVATION OF ANALYTICAL LEAST SQUARES SOLUTIONS**
 874

875 Here, we theoretically derive the analytical solutions to the least squares problems (4) and (7) ad-
 876 dressed in this paper. In fact, these solutions constitute the centralized analytical solutions for both
 877 the global classifier and the transformation matrix. The detailed proofs are as follows.
 878

879 **Lemma 1:** For any least squares problem with the following form:

880
$$\mathbf{W}^* = \arg \min_{\mathbf{W}} \|\mathbf{Y} - \Phi \mathbf{W}\|_F^2 + \lambda \|\mathbf{W}\|_F^2, \quad (26)$$

 881

882 it yields a distinct analytical (i.e., closed-form) solution, which can be formulated as:

883
$$\mathbf{W}^* = (\Phi^\top \Phi + \lambda \mathbf{I})^{-1} \Phi^\top \mathbf{Y}. \quad (27)$$

 884

885 *Proof.* Our proof strategy involves computing the derivative of the loss function in (26) with respect
 886 to \mathbf{W} and subsequently setting it to zero to derive the optimal \mathbf{W}^* . Specifically, we first denote the
 887 loss function of the aforementioned least squares problem (26) as follows:
 888

889
$$\mathcal{J}(\mathbf{W}) = \|\mathbf{Y} - \Phi \mathbf{W}\|_F^2 + \lambda \|\mathbf{W}\|_F^2. \quad (28)$$

 890

891 Subsequently, to facilitate subsequent differentiation, we expand the loss function (28) as follows:

892
$$\begin{aligned} \mathcal{J}(\mathbf{W}) &= \|\mathbf{Y} - \Phi \mathbf{W}\|_F^2 + \lambda \|\mathbf{W}\|_F^2 \\ &= \text{Tr}((\mathbf{Y} - \Phi \mathbf{W})^\top (\mathbf{Y} - \Phi \mathbf{W})) + \lambda \text{Tr}(\mathbf{W}^\top \mathbf{W}) \\ &= \text{Tr}(\mathbf{Y}^\top \mathbf{Y}) + 2\text{Tr}(\mathbf{Y}^\top \Phi \mathbf{W}) + \text{Tr}(\mathbf{W}^\top \Phi^\top \Phi \mathbf{W}) + \lambda \text{Tr}(\mathbf{W}^\top \mathbf{W}), \end{aligned} \quad (29)$$

 893

894 where $\text{Tr}(\cdot)$ represents the trace of a matrix. To minimize the objective function, we further compute
 895 the derivative of $\mathcal{J}(\mathbf{W})$ with respect to \mathbf{W} as follows:
 896

897
$$\frac{\partial \mathcal{J}(\mathbf{W})}{\partial \mathbf{W}} = -2\Phi^\top \mathbf{Y} + 2\Phi^\top \Phi \mathbf{W} + 2\lambda \mathbf{W}. \quad (30)$$

 898

899 By setting the derivative to zero, we can obtain:
 900

901
$$(\Phi^\top \Phi + \lambda \mathbf{I}) \mathbf{W}^* = \Phi^\top \mathbf{Y}. \quad (31)$$

902 Since $(\Phi^\top \Phi + \lambda \mathbf{I})$ is full rank and hence invertible, the closed-form solution for the optimal \mathbf{W}^*
 903 can be obtained as follows:
 904

905
$$\mathbf{W}^* = (\Phi^\top \Phi + \lambda \mathbf{I})^{-1} \Phi^\top \mathbf{Y}. \quad (32)$$

 906

907 \blacksquare

908 **Lemma 2:** For any least squares problem with the following form:
 909

910
$$\Omega^* = \arg \min_{\Omega} \|\mathbf{R} - \mathbf{F} \Omega \mathbf{W}\|_F^2 + \gamma \|\Omega\|_F^2, \quad (33)$$

 911

912 it yields a distinct analytical (i.e., closed-form) solution, which can be formulated as:
 913

914
$$\Omega^* = \mathbf{V} [(\mathbf{V}^\top \mathbf{F}^\top \mathbf{R} \mathbf{W}^\top \mathbf{U}) \oslash (\gamma \mathbf{1} + \text{diag}(\Lambda^F) \otimes \text{diag}(\Lambda^W))] \mathbf{U}^\top, \quad (34)$$

 915

916 where \oslash and \otimes represent element-wise division and outer product, respectively, while $\mathbf{F}^\top \mathbf{F} =$
 917 $\mathbf{V} \Lambda^F \mathbf{V}^\top$ and $\mathbf{W} \mathbf{W}^\top = \mathbf{U} \Lambda^W \mathbf{U}^\top$ are spectral decompositions.
 918

918 *Proof.* Similar to the proof of **Lemma 1**, we begin by computing the derivative of the loss function
 919 in (33) with respect to Ω and then set this derivative to zero to derive the optimal Ω^* . Specifically,
 920 we denote the loss function of the aforementioned least squares problem (33) as follows:
 921

$$922 \quad \mathcal{L}(\Omega) = \|\mathbf{R} - \mathbf{F}\Omega\mathbf{W}\|_{\text{F}}^2 + \gamma\|\Omega\|_{\text{F}}^2, \quad (35)$$

923 which can be further expanded into the following form:
 924

$$925 \quad \begin{aligned} \mathcal{L}(\Omega) &= \|\mathbf{R} - \mathbf{F}\Omega\mathbf{W}\|_{\text{F}}^2 + \gamma\|\Omega\|_{\text{F}}^2 \\ 926 &= \text{Tr}((\mathbf{R} - \mathbf{F}\Omega\mathbf{W})^{\top}(\mathbf{R} - \mathbf{F}\Omega\mathbf{W})) + \gamma\text{Tr}(\Omega^{\top}\Omega) \\ 927 &= \text{Tr}(\mathbf{R}^{\top}\mathbf{R}) - 2\text{Tr}(\mathbf{R}^{\top}\mathbf{F}\Omega\mathbf{W}) + \text{Tr}(\mathbf{W}^{\top}\Omega^{\top}\mathbf{F}^{\top}\mathbf{F}\Omega\mathbf{W}) + \gamma\text{Tr}(\Omega^{\top}\Omega). \end{aligned} \quad (36)$$

928 Subsequently, by further differentiating (36), we can obtain:
 929

$$930 \quad \frac{\partial \mathcal{L}(\Omega)}{\partial \Omega} = -2\mathbf{F}^{\top}\mathbf{R}\mathbf{W}^{\top} + 2\mathbf{F}^{\top}\mathbf{F}\Omega\mathbf{W}\mathbf{W}^{\top} + 2\gamma\Omega. \quad (37)$$

934 By setting the derivative to zero, we can obtain:
 935

$$936 \quad \mathbf{F}^{\top}\mathbf{F}\Omega^*\mathbf{W}\mathbf{W}^{\top} + \gamma\Omega^* = \mathbf{F}^{\top}\mathbf{R}\mathbf{W}^{\top}. \quad (38)$$

938 Since $\mathbf{F}^{\top}\mathbf{F}$ and $\mathbf{W}\mathbf{W}^{\top}$ are real symmetric matrices, we can spectrally decompose them to obtain:
 939

$$940 \quad \mathbf{F}^{\top}\mathbf{F} = \mathbf{V}\Lambda^{\text{F}}\mathbf{V}^{\top}, \quad \mathbf{W}\mathbf{W}^{\top} = \mathbf{U}\Lambda^{\text{W}}\mathbf{U}^{\top}, \quad (39)$$

942 where $\mathbf{V}^{\top}\mathbf{V} = \mathbf{I}$ and $\mathbf{U}^{\top}\mathbf{U} = \mathbf{I}$, and \mathbf{I} denotes the identity matrix. By denoting $\mathbf{S} = \mathbf{V}^{\top}\Omega^*\mathbf{U}$ and
 943 substituting (39) into (38), we obtain:

$$944 \quad \begin{aligned} \mathbf{F}^{\top}\mathbf{F}\Omega^*\mathbf{W}\mathbf{W}^{\top} + \gamma\Omega^* &= \mathbf{F}^{\top}\mathbf{R}\mathbf{W}^{\top} \\ 945 &\Rightarrow \mathbf{V}\Lambda^{\text{F}}\mathbf{V}^{\top}\Omega^*\mathbf{U}\Lambda^{\text{W}}\mathbf{U}^{\top} + \gamma\Omega^* = \mathbf{F}^{\top}\mathbf{R}\mathbf{W}^{\top} \\ 946 &\Rightarrow \mathbf{V}^{\top}\mathbf{V}\Lambda^{\text{F}}\mathbf{S}\Lambda^{\text{W}}\mathbf{U}^{\top}\mathbf{U} + \gamma\mathbf{V}^{\top}\Omega^*\mathbf{U} = \mathbf{V}^{\top}\mathbf{F}^{\top}\mathbf{R}\mathbf{W}^{\top}\mathbf{U} \\ 947 &\Rightarrow \Lambda^{\text{F}}\mathbf{S}\Lambda^{\text{W}} + \gamma\mathbf{S} = \mathbf{V}^{\top}\mathbf{F}^{\top}\mathbf{R}\mathbf{W}^{\top}\mathbf{U}. \end{aligned} \quad (40)$$

950 Since Λ^{F} and Λ^{W} are both diagonal matrices, we proceed to expand them element-wise. Specifically,
 951 let λ_i^{F} and λ_j^{W} denote the i -th and j -th diagonal elements of Λ^{F} and Λ^{W} , respectively. Con-
 952 currently, we denote the (i, j) -th entry of \mathbf{S} and $\mathbf{V}^{\top}\mathbf{F}^{\top}\mathbf{R}\mathbf{W}^{\top}\mathbf{U}$ by $S_{i,j}$ and $(\mathbf{V}^{\top}\mathbf{F}^{\top}\mathbf{R}\mathbf{W}^{\top}\mathbf{U})_{i,j}$,
 953 respectively. According to (40), we obtain:
 954

$$955 \quad \lambda_i^{\text{F}}S_{i,j}\lambda_j^{\text{W}} + \gamma S_{i,j} = (\mathbf{V}^{\top}\mathbf{F}^{\top}\mathbf{R}\mathbf{W}^{\top}\mathbf{U})_{i,j}. \quad (41)$$

957 Meanwhile, we can further obtain:
 958

$$959 \quad S_{i,j} = \frac{(\mathbf{V}^{\top}\mathbf{F}^{\top}\mathbf{R}\mathbf{W}^{\top}\mathbf{U})_{i,j}}{\lambda_i^{\text{F}}\lambda_j^{\text{W}} + \gamma}. \quad (42)$$

962 Based on (42), we can extend $S_{i,j}$ to the entire matrix \mathbf{S} , yielding:
 963

$$964 \quad \mathbf{S} = (\mathbf{V}^{\top}\mathbf{F}^{\top}\mathbf{R}\mathbf{W}^{\top}\mathbf{U}) \oslash (\gamma\mathbf{1} + \text{diag}(\Lambda^{\text{F}}) \otimes \text{diag}(\Lambda^{\text{W}})), \quad (43)$$

966 where $\mathbf{1}$ is the all-ones matrix, while \oslash and \otimes represent element-wise division and outer product.
 967 Finally, leveraging the orthogonality of \mathbf{V} and \mathbf{U} , we can obtain $\Omega^* = \mathbf{V}\mathbf{S}\mathbf{U}^{\top}$. Substituting (43)
 968 into this expression, we can obtain:
 969

$$970 \quad \Omega^* = \mathbf{V}(\mathbf{V}^{\top}\mathbf{F}^{\top}\mathbf{R}\mathbf{W}^{\top}\mathbf{U}) \oslash (\gamma\mathbf{1} + \text{diag}(\Lambda^{\text{F}}) \otimes \text{diag}(\Lambda^{\text{W}}))\mathbf{U}^{\top}. \quad (44)$$

971

■

972 B.2 OUR DEEP AFL'S INVARIANCE TO DATA HETEROGENEITY
973

974 Here, we demonstrate that the deep residual analytic network derived from our DeepAFL exhibits
975 ideal invariance to data heterogeneity, remaining invariant to any data partition across clients with
976 different heterogeneity. Specifically, as established in Lemma 1 and Lemma 2, we have derived the
977 centralized analytical solutions for both the global classifier and the transformation matrix. Here,
978 we additionally prove that the solution obtained from our DeepAFL is identical to these centralized
979 analytical solutions. The detailed proofs are as follows.

980 **Theorem 1 (Invariance to Data Heterogeneity):** In the FL scenario, let \mathcal{D} be the full dataset, and
981 $\mathcal{P} = \{\mathcal{D}^1, \mathcal{D}^2, \dots, \mathcal{D}^K\}$ be any heterogeneous partition of \mathcal{D} among K clients, where $\mathcal{D}_i \cap \mathcal{D}_j = \emptyset$
982 for $i \neq j$ and $\bigcup_{i=1}^N \mathcal{D}_i = \mathcal{D}$. Given fixed seeds for the random projection matrices \mathbf{A} and $\{\mathbf{B}_t\}_{t=0}^{T-1}$,
983 the global model weights $\{\mathbf{W}_t\}_{t=0}^T$ and $\{\Omega_t\}_{t=1}^T$ derived from DeepAFL are invariant to any data
984 partition \mathcal{P} with different heterogeneity, being identical to the centralized analytical solutions on \mathcal{D} .
985

986 *Proof.* As established in Lemma 1 and Lemma 2, for the complete dataset $\mathcal{D} = \{\mathbf{X}^{1:K}, \mathbf{Y}^{1:K}\}$, the
987 centralized analytical solutions for $\{\mathbf{W}_t\}_{t=0}^T$ and $\{\Omega_t\}_{t=1}^T$ can be expressed as:

$$988 \mathbf{W}_t = [(\Phi_t^{1:K})^\top \Phi_t^{1:K} + \lambda \mathbf{I}]^{-1} (\Phi_t^{1:K})^\top \mathbf{Y}, \quad (45)$$

$$989 \Omega_{t+1} = \mathbf{V}_t [(\mathbf{V}_t^\top (\mathbf{F}_t^{1:K})^\top \mathbf{R}_t^{1:K} \mathbf{W}_t^\top \mathbf{U}_t) \oslash (\gamma \mathbf{1} + \text{diag}(\Lambda_t^F) \otimes \text{diag}(\Lambda_t^W))] \mathbf{U}_t^\top, \quad (46)$$

990 where $(\mathbf{F}_t^{1:K})^\top \mathbf{F}_t^{1:K} = \mathbf{V}_t \Lambda_t^F \mathbf{V}_t^\top$ and $\mathbf{W}_t \mathbf{W}_t^\top = \mathbf{U}_t \Lambda_t^W \mathbf{U}_t^\top$ are spectral decompositions, while \oslash
991 and \otimes represent element-wise division and outer product, respectively.
992

993 In fact, any heterogeneous partition \mathcal{P} can be viewed as a permutation of the sample ordering within
994 the complete dataset \mathcal{D} . Accordingly, for any heterogeneous partition $\mathcal{P} = \{\mathcal{D}^1, \mathcal{D}^2, \dots, \mathcal{D}^K\}$, we
995 define the corresponding permuted complete dataset as $\tilde{\mathcal{D}} = \{\tilde{\mathbf{X}}^{1:K}, \tilde{\mathbf{Y}}^{1:K}\}$, which satisfies:

$$996 \tilde{\mathbf{X}}^{1:K} = \pi \mathbf{X}^{1:K} \quad \tilde{\mathbf{Y}}^{1:K} = \pi \mathbf{Y}^{1:K}, \quad (47)$$

997 where π is the corresponding permutation matrix and satisfies $\pi^\top \pi = \mathbf{I}$. Next, we employ mathematical
998 induction to rigorously demonstrate that the global model weights $\{\tilde{\mathbf{W}}_t\}_{t=0}^T$ and $\{\tilde{\Omega}_t\}_{t=1}^T$
999 obtained by our DeepAFL under any heterogeneous partition \mathcal{P} are identical to their centralized
1000 analytical solutions, as presented in (45) and (46).
1001

1002 For the base case, we demonstrate that the initial $\tilde{\mathbf{W}}_0$ and $\tilde{\Omega}_1$ obtained via our DeepAFL precisely
1003 satisfy (45) and (46). It can be readily deduced that the complete zero-layer feature matrix similarly
1004 satisfies $\tilde{\Phi}_0^{1:K} = \pi \Phi_0^{1:K}$. According to (13), the expression for $\tilde{\mathbf{W}}_0$ can be given by:
1005

$$1006 \tilde{\mathbf{W}}_0 = (\tilde{\mathbf{G}}_0^{1:K} + \lambda \mathbf{I})^{-1} \tilde{\mathbf{H}}_0^{1:K} \\ 1007 = [(\tilde{\Phi}_0^{1:K})^\top \tilde{\Phi}_0^{1:K} + \lambda \mathbf{I}]^{-1} (\tilde{\Phi}_0^{1:K})^\top \tilde{\mathbf{Y}}^{1:K} \\ 1008 = [(\Phi_0^{1:K})^\top \pi^\top \pi \Phi_0^{1:K} + \lambda \mathbf{I}]^{-1} (\tilde{\Phi}_0^{1:K})^\top \tilde{\mathbf{Y}}^{1:K} \\ 1009 = [(\Phi_0^{1:K})^\top \Phi_0^{1:K} + \lambda \mathbf{I}]^{-1} (\tilde{\Phi}_0^{1:K})^\top \tilde{\mathbf{Y}}^{1:K}. \quad (48)$$

1010 It is evident that $\tilde{\mathbf{W}}_0$ coincides precisely with the corresponding centralized analytical solution (45).
1011 Consequently, the residual corresponding to each sample matches that obtained through centralized
1012 training, differing only in sample order. Thus, the complete residual matrix satisfies $\tilde{\mathbf{R}}_0^{1:K} = \pi \mathbf{R}_0^{1:K}$.
1013 Similarly, the complete hidden random features also satisfy $\tilde{\mathbf{F}}_0^{1:K} = \pi \mathbf{F}_0^{1:K}$. Drawing from (17) and
1014 (18), the $\tilde{\Omega}_1$ is given by:
1015

$$1017 \tilde{\Omega}_1 = \mathbf{V}_0 [(\mathbf{V}_0^\top \tilde{\mathbf{Y}}_0^{1:K} \mathbf{W}_0^\top \mathbf{U}_0) \oslash (\gamma \mathbf{1} + \text{diag}(\Lambda_0^F) \otimes \text{diag}(\Lambda_0^W))] \mathbf{U}_0^\top \\ 1018 = \mathbf{V}_0 [(\mathbf{V}_0^\top (\tilde{\mathbf{F}}_0^{1:K})^\top \tilde{\mathbf{R}}_0^{1:K} \mathbf{W}_0^\top \mathbf{U}_0) \oslash (\gamma \mathbf{1} + \text{diag}(\Lambda_0^F) \otimes \text{diag}(\Lambda_0^W))] \mathbf{U}_0^\top \\ 1019 = \mathbf{V}_0 [(\mathbf{V}_0^\top (\mathbf{F}_0^{1:K})^\top \pi^\top \pi \mathbf{R}_0^{1:K} \mathbf{W}_0^\top \mathbf{U}_0) \oslash (\gamma \mathbf{1} + \text{diag}(\Lambda_0^F) \otimes \text{diag}(\Lambda_0^W))] \mathbf{U}_0^\top \\ 1020 = \mathbf{V}_0 [(\mathbf{V}_0^\top (\mathbf{F}_0^{1:K})^\top \mathbf{R}_0^{1:K} \mathbf{W}_0^\top \mathbf{U}_0) \oslash (\gamma \mathbf{1} + \text{diag}(\Lambda_0^F) \otimes \text{diag}(\Lambda_0^W))] \mathbf{U}_0^\top, \quad (49)$$

1021 where the spectral decomposition results are also identical to those in the preceding formula (46):
1022 $(\tilde{\mathbf{F}}_0^{1:K})^\top \tilde{\mathbf{F}}_0^{1:K} = (\mathbf{F}_0^{1:K})^\top \pi^\top \pi \mathbf{F}_0^{1:K} = (\mathbf{F}_0^{1:K})^\top \mathbf{F}_0^{1:K} = \mathbf{V}_0 \Lambda_0^F \mathbf{V}_0^\top$ and $\tilde{\mathbf{W}}_0 \tilde{\mathbf{W}}_0^\top = \mathbf{W}_0 \mathbf{W}_0^\top =$
1023 $\mathbf{U}_0 \Lambda_0^W \mathbf{U}_0^\top$. It can be observed that the initial $\tilde{\Omega}_1$ derived from our DeepAFL also precisely matches
1024 the centralized analytical solution. Thus, the base case holds for the theorem.
1025

For the inductive step, assume the preceding $(t-1)$ layers $\{\tilde{\mathbf{W}}_i\}_{i=0}^{t-1}$ and $\{\tilde{\Omega}_i\}_{i=1}^t$ satisfy (45) and (46). We further demonstrate that, based on these assumptions, the subsequent layer $\tilde{\mathbf{W}}_t$ and $\tilde{\Omega}_{t+1}$ obtained via our DeepAFL also satisfy (45) and (46). Since the previous $(t-1)$ layers' global model weights exactly match the centralized analytical solution, the complete feature matrix $\tilde{\Phi}_t^{1:K}$ and the complete hidden random features $\tilde{\mathbf{F}}_t^{1:K}$ satisfy $\tilde{\Phi}_t^{1:K} = \pi\Phi_t^{1:K}$ and $\tilde{\mathbf{F}}_t^{1:K} = \pi\mathbf{F}_t^{1:K}$. According to (13), the expression for $\tilde{\mathbf{W}}_t$ is thus obtained as follows:

$$\begin{aligned} \tilde{\mathbf{W}}_t &= (\tilde{\mathbf{G}}_t^{1:K} + \lambda\mathbf{I})^{-1}\tilde{\mathbf{H}}_t^{1:K} \\ &= [(\tilde{\Phi}_t^{1:K})^\top\tilde{\Phi}_t^{1:K} + \lambda\mathbf{I}]^{-1}(\tilde{\Phi}_t^{1:K})^\top\tilde{\mathbf{Y}}^{1:K} \\ &= [(\Phi_t^{1:K})^\top\pi^\top\pi\Phi_t^{1:K} + \lambda\mathbf{I}]^{-1}(\tilde{\Phi}_t^{1:K})^\top\tilde{\mathbf{Y}}^{1:K} \\ &= [(\Phi_t^{1:K})^\top\Phi_t^{1:K} + \lambda\mathbf{I}]^{-1}(\tilde{\Phi}_t^{1:K})^\top\tilde{\mathbf{Y}}^{1:K}. \end{aligned} \quad (50)$$

It is evident that the global classifier $\tilde{\mathbf{W}}_t$ satisfy (45), and thus the complete residual matrix satisfies $\tilde{\mathbf{R}}_t^{1:K} = \pi\mathbf{R}_t^{1:K}$. Subsequently, analogously to (49), we derive the analytical expression for $\tilde{\Omega}_{t+1}$ yielded by our DeepAFL as follows:

$$\begin{aligned} \tilde{\Omega}_{t+1} &= \mathbf{V}_t[(\mathbf{V}_t^\top\tilde{\mathbf{Y}}_t^{1:K}\mathbf{W}_t^\top\mathbf{U}_t) \odot (\gamma\mathbf{1} + \text{diag}(\Lambda_t^F) \otimes \text{diag}(\Lambda_t^W))] \mathbf{U}_t^\top \\ &= \mathbf{V}_t[(\mathbf{V}_t^\top(\tilde{\mathbf{F}}_t^{1:K})^\top\tilde{\mathbf{R}}_t^{1:K}\mathbf{W}_t^\top\mathbf{U}_t) \odot (\gamma\mathbf{1} + \text{diag}(\Lambda_t^F) \otimes \text{diag}(\Lambda_t^W))] \mathbf{U}_t^\top \\ &= \mathbf{V}_t[(\mathbf{V}_t^\top(\mathbf{F}_t^{1:K})^\top\pi^\top\pi\mathbf{R}_t^{1:K}\mathbf{W}_t^\top\mathbf{U}_t) \odot (\gamma\mathbf{1} + \text{diag}(\Lambda_t^F) \otimes \text{diag}(\Lambda_t^W))] \mathbf{U}_t^\top \\ &= \mathbf{V}_t[(\mathbf{V}_t^\top(\mathbf{F}_t^{1:K})^\top\mathbf{R}_t^{1:K}\mathbf{W}_t^\top\mathbf{U}_t) \odot (\gamma\mathbf{1} + \text{diag}(\Lambda_t^F) \otimes \text{diag}(\Lambda_t^W))] \mathbf{U}_t^\top, \end{aligned} \quad (51)$$

where the spectral decomposition results are also identical to those in the preceding formula (46): $(\tilde{\mathbf{F}}_t^{1:K})^\top\tilde{\mathbf{F}}_t^{1:K} = (\mathbf{F}_t^{1:K})^\top\pi^\top\pi\mathbf{F}_t^{1:K} = (\mathbf{F}_t^{1:K})^\top\mathbf{F}_t^{1:K} = \mathbf{V}_t\Lambda_t^F\mathbf{V}_t^\top$ and $\tilde{\mathbf{W}}_t\tilde{\mathbf{W}}_t^\top = \mathbf{W}_t\mathbf{W}_t^\top = \mathbf{U}_t\Lambda_t^W\mathbf{U}_t^\top$. As a result, the transformation matrix $\tilde{\Omega}_{t+1}$ derived from our DeepAFL also precisely matches the centralized analytical solution. Thus, the inductive step also holds for the theorem.

In summary, based on the aforementioned base case and inductive step, we establish via mathematical induction that the global model weights $\{\mathbf{W}_t\}_{t=0}^T$ and $\{\Omega_t\}_{t=1}^T$ derived from our DeepAFL are invariant to any data partition \mathcal{P} with different heterogeneity. In fact, the results obtained from our DeepAFL are identical to the centralized analytical solutions (45) and (46). ■

B.3 OUR DEEPAFL'S CAPABILITY OF REPRESENTATION LEARNING

Here, we further demonstrate that the empirical risk monotonically non-increases with increasing layer depth and is guaranteed to converge to a limit, thus validating DeepAFL's representation learning capability. To begin, we analyze the general scenario without considering the regularization term in Theorem 2, proving that under this scenario, the empirical risk derived from our DeepAFL monotonically non-increases with increasing layer depth and converges to a limit.

Theorem 2 (Capability of Representation Learning): Let the general empirical risk $\mathcal{H}(\Phi, \mathbf{W})$ (i.e., the loss function), for a given feature representation Φ and a classifier \mathbf{W} be defined as:

$$\mathcal{H}(\Phi, \mathbf{W}) = \|\mathbf{Y} - \Phi\mathbf{W}\|_F^2, \quad (52)$$

where \mathbf{Y} denotes the ground truth labels. Our DeepAFL yields a sequence of feature-classifier pairs (Φ_t, \mathbf{W}_t) at each layer $t \in [0, T]$. When setting regularization parameters γ and λ to 0, the sequence of the empirical risks $\{\mathcal{H}(\Phi_t, \mathbf{W}_t)\}_{t=0}^T$ in our DeepAFL keeps monotonically non-increasing, i.e.,

$$\mathcal{H}(\Phi_t, \mathbf{W}_t) \geq \mathcal{H}(\Phi_{t+1}, \mathbf{W}_{t+1}), \forall t \in [0, T]. \quad (53)$$

Furthermore, as the number of layers T within our DeepAFL increases (i.e., $T \rightarrow \infty$), the sequence of empirical risks is guaranteed to converge to a limit $\mathcal{H}^* \leq \mathcal{H}(\Phi_0, \mathbf{W}_0)$.

Proof. According to Algorithm 1, the construction of each feature-classifier pair $(\Phi_{t+1}, \mathbf{W}_{t+1})$ proceeds in two alternating steps: (1) fixing the classifier \mathbf{W}_t to construct the feature representation Φ_{t+1} ; (2) fixing the updated feature representation Φ_{t+1} to optimize the classifier \mathbf{W}_{t+1} . We will prove that the empirical risks in each of these substeps are non-increasing, thereby demonstrating that the sequence of empirical risks $\{\mathcal{H}(\Phi_t, \mathbf{W}_t)\}_{t=0}^T$ is monotonically non-increasing.

1080 First, let's focus on the substep of constructing the feature representation Φ_{t+1} while keeping the
 1081 global classifier \mathbf{W}_t fixed. At this stage, the intermediate empirical risk for the feature-classifier
 1082 pair $(\Phi_{t+1}, \mathbf{W}_t)$ can be expressed as:
 1083

$$\mathcal{H}(\Phi_{t+1}, \mathbf{W}_t) = \|\mathbf{Y} - \Phi_{t+1}\mathbf{W}_t\|_F^2, \quad (54)$$

1085 where Φ_{t+1} obtained through our DeepAFL satisfies $\Phi_{t+1} = \Phi_t + \mathbf{F}_t\Omega_{t+1}$. Therefore, the afore-
 1086 mentioned empirical risk in our DeepAFL can be further expressed as:
 1087

$$\mathcal{H}(\Phi_{t+1}, \mathbf{W}_t) = \|\mathbf{Y} - (\Phi_t + \mathbf{F}_t\Omega_{t+1})\mathbf{W}_t\|_F^2. \quad (55)$$

1089 According to Lemma 2, the Ω_{t+1} derived through our DeepAFL represents the optimal solution that
 1090 minimizes the aforementioned empirical risk, thereby satisfying:
 1091

$$\mathcal{H}(\Phi_{t+1}, \mathbf{W}_t) = \|\mathbf{Y} - (\Phi_t + \mathbf{F}_t\Omega_{t+1})\mathbf{W}_t\|_F^2 \leq \|\mathbf{Y} - (\Phi_t + \mathbf{F}_t\Omega)\mathbf{W}_t\|_F^2, \forall \Omega, \quad (56)$$

1093 where, as a special case when $\Omega = \mathbf{0}$, it holds that:
 1094

$$\|\mathbf{Y} - (\Phi_t + \mathbf{F}_t\mathbf{0})\mathbf{W}_t\|_F^2 = \|\mathbf{Y} - \Phi_t\mathbf{W}_t\|_F^2 = \mathcal{H}(\Phi_t, \mathbf{W}_t). \quad (57)$$

1096 Thus, it follows that:
 1097

$$\mathcal{H}(\Phi_{t+1}, \mathbf{W}_t) \leq \mathcal{H}(\Phi_t, \mathbf{W}_t). \quad (58)$$

1099 Second, we further analyze the substep of constructing the global classifier \mathbf{W}_{t+1} while keeping the
 1100 feature representation Φ_{t+1} fixed. When the regularization parameter λ is set to zero, the \mathbf{W}_{t+1}
 1101 derived through our DeepAFL is obtained by minimizing the following empirical risk:
 1102

$$\mathbf{W}_{t+1} = \arg \min_{\mathbf{W}} \|\mathbf{Y} - \Phi_{t+1}\mathbf{W}\|_F^2 = \arg \min_{\mathbf{W}} \mathcal{H}(\Phi_{t+1}, \mathbf{W}). \quad (59)$$

1104 According to Lemma 1, the \mathbf{W}_{t+1} derived through our DeepAFL represents the optimal solution
 1105 that minimizes the aforementioned empirical risk. We can obtain:
 1106

$$\mathcal{H}(\Phi_{t+1}, \mathbf{W}_{t+1}) = \|\mathbf{Y} - \Phi_{t+1}\mathbf{W}_{t+1}\|_F^2 \leq \|\mathbf{Y} - \Phi_{t+1}\mathbf{W}\|_F^2, \forall \mathbf{W}, \quad (60)$$

1108 where, as a special case when $\mathbf{W} = \mathbf{W}_t$, it holds that:
 1109

$$\|\mathbf{Y} - \Phi_{t+1}\mathbf{W}_t\|_F^2 = \mathcal{H}(\Phi_{t+1}, \mathbf{W}_t). \quad (61)$$

1111 Thus, it follows that:
 1112

$$\mathcal{H}(\Phi_{t+1}, \mathbf{W}_{t+1}) \leq \mathcal{H}(\Phi_{t+1}, \mathbf{W}_t). \quad (62)$$

1113 In summary, based on (58) and (62), it holds for each feature-classifier pair $(\Phi_{t+1}, \mathbf{W}_{t+1})$ that:
 1114

$$\mathcal{H}(\Phi_t, \mathbf{W}_t) \geq \mathcal{H}(\Phi_{t+1}, \mathbf{W}_t) \geq \mathcal{H}(\Phi_{t+1}, \mathbf{W}_{t+1}), \forall t \in [0, T]. \quad (63)$$

1116 Therefore, the sequence of empirical risks $\{\mathcal{H}(\Phi_t, \mathbf{W}_t)\}_{t=0}^T$ is monotonically non-increasing.
 1117 Given that $\mathcal{H}(\Phi_t, \mathbf{W}_t) \geq 0$ and the sequence is monotonically non-increasing, by the Monotone
 1118 Convergence Theorem, it necessarily converges to a limit $\mathcal{H}^* \geq 0$. Moreover, since the sequence
 1119 begins at $\mathcal{H}(\Phi_0, \mathbf{W}_0)$ and decreases at each layer, it holds that $\mathcal{H}^* \leq \mathcal{H}(\Phi_0, \mathbf{W}_0)$.
 1120 ■

1121 Considering our employment of regularization to enhance generalization performance and numerical
 1122 stability, we further extend Theorem 2 to incorporate the regularized setting. To this end, we begin
 1123 by analyzing the regularization terms in the regularized empirical risk, thereby defining it for the
 1124 subsequent theorem, as detailed below:
 1125

1126 (1) From the perspective of empirical risk, the empirical risk is a function of the global classifier \mathbf{W}_t
 1127 and the feature representation Φ_t . Accordingly, both should be subject to regularization. The feature
 1128 representation Φ_t is obtained through layer-by-layer updates and can essentially be expressed as:
 1129

$$\Phi_t = \Phi_{t-1} + (\mathbf{F}_{t-1}\Omega_t) = \Phi_{t-2} + (\mathbf{F}_{t-2}\Omega_t + \mathbf{F}_{t-1}\Omega_{t-1}) = \dots = \Phi_0 + \sum_{i=1}^t \mathbf{F}_{i-1}\Omega_i, \quad (64)$$

1131 where \mathbf{F}_{i-1} is not trainable, providing stochasticity and nonlinearity to the features via fixed acti-
 1132 vation functions and random projections, whereas Ω_i constitutes the trainable component for rep-
 1133 resentation learning. Thus, regularizing all $\{\Omega_i\}_{i=1}^t$ can be seen as a form of regularizing Φ_t . In
 addition, the global classifier \mathbf{W}_t is itself trainable, and should be regularized directly.

(2) From the perspective of the model architecture, the global model derived from our DeepAFL after constructing the first t layers encompasses the final global classifier \mathbf{W}_t and all transformation matrices $\{\Omega_i\}_{i=1}^t$. Therefore, regularization terms need to be introduced for both \mathbf{W}_t and $\{\Omega_i\}_{i=1}^t$.

Building upon the foregoing analyses of the regularized empirical risk, we further provide its explicit definition in (65). Moreover, in Theorem 3, we theoretically prove that the empirical risk of the resulting model is monotonically non-increasing as the layer depth increases.

Theorem 3 (Regularized Capability of Representation Learning): In our DeepAFL, we define the regularized empirical risk $\mathcal{G}(t)$ (i.e., the regularized loss function) for each layer $t \in [0, T]$ as:

$$\mathcal{G}(t) = \underbrace{\|\mathbf{Y} - \Phi_t \mathbf{W}_t\|_F^2}_{(1)} + \lambda \underbrace{\|\mathbf{W}_t\|_F^2}_{(2)} + \sum_{i=1}^t \gamma \underbrace{\|\Omega_i\|_F^2}_{(3)}, \quad (65)$$

where $\gamma > 0$ and $\lambda > 0$ are non-negative regularization parameters. Moreover, Φ_t , \mathbf{W}_t , and Ω_t are the feature representation, the analytic classifier, and the learned transformation in our DeepAFL at the layer t . Here, the term (1) is the original empirical risk $\mathcal{H}(\Phi_t, \mathbf{W}_t)$, while the terms (2) and (3) are the regularization in our DeepAFL. The sequence of the regularized empirical risks $\{\mathcal{G}(t)\}_{t=0}^T$ in our DeepAFL remains monotonically non-increasing, i.e.,

$$\mathcal{G}(t) \geq \mathcal{G}(t+1), \forall t \in [0, T]. \quad (66)$$

Furthermore, as the number of layers T within our DeepAFL increases (i.e., $T \rightarrow \infty$), the sequence of the regularized empirical risks is guaranteed to converge to a limit $\mathcal{G}^* \leq \mathcal{G}(0)$.

Proof. Analogous to Theorem 2, we sequentially analyze the construction of the feature representation Φ_{t+1} (i.e., deriving the transformation matrix Ω_{t+1}) and the global classifier \mathbf{W}_{t+1} at each layer, thereby demonstrating that the regularized empirical risks $\{\mathcal{G}(t)\}_{t=0}^T$ in our DeepAFL remain monotonically non-increasing.

First, let's focus on the substep of constructing the feature representation Φ_{t+1} and the learned transformation Ω_{t+1} while fixing the global classifier \mathbf{W}_t and other transformations $\{\Omega_i\}_{i=1}^t$. For notational convenience, we denote the intermediate regularized empirical risk between layers t and $t+1$ (i.e., the risk associated with Φ_{t+1} , \mathbf{W}_t , and $\{\Omega_i\}_{i=1}^{t+1}$) as $\hat{\mathcal{G}}(t)$, which can be expressed as:

$$\hat{\mathcal{G}}(t) = \|\mathbf{Y} - \Phi_{t+1} \mathbf{W}_t\|_F^2 + \lambda \|\mathbf{W}_t\|_F^2 + \sum_{i=1}^{t+1} \gamma \|\Omega_i\|_F^2, \quad (67)$$

where Φ_{t+1} obtained through our DeepAFL satisfies $\Phi_{t+1} = \Phi_t + \mathbf{F}_t \Omega_{t+1}$. Therefore, the aforementioned regularized empirical risk can be further expressed as:

$$\hat{\mathcal{G}}(t) = \|\mathbf{Y} - (\Phi_t + \mathbf{F}_t \Omega_{t+1}) \mathbf{W}_t\|_F^2 + \lambda \|\mathbf{W}_t\|_F^2 + \sum_{i=1}^{t+1} \gamma \|\Omega_i\|_F^2. \quad (68)$$

According to Lemma 2, the Ω_{t+1} derived through our DeepAFL represents the optimal solution that minimizes the regularized empirical risk:

$$\Omega_{t+1} = \arg \min_{\Omega} \|\mathbf{Y} - (\Phi_t + \mathbf{F}_t \Omega) \mathbf{W}_t\|_F^2 + \gamma \|\Omega\|_F^2. \quad (69)$$

Therefore, we can obtain:

$$\|\mathbf{Y} - (\Phi_t + \mathbf{F}_t \Omega_{t+1}) \mathbf{W}_t\|_F^2 + \gamma \|\Omega_{t+1}\|_F^2 \leq \|\mathbf{Y} - (\Phi_t + \mathbf{F}_t \Omega) \mathbf{W}_t\|_F^2 + \gamma \|\Omega\|_F^2, \forall \Omega, \quad (70)$$

where $\Omega = \mathbf{0}$ serves as a special case, from which we further derive:

$$\|\mathbf{Y} - (\Phi_t + \mathbf{F}_t \Omega_{t+1}) \mathbf{W}_t\|_F^2 + \gamma \|\Omega_{t+1}\|_F^2 \leq \|\mathbf{Y} - \Phi_t \mathbf{W}_t\|_F^2. \quad (71)$$

By adding $\lambda \|\mathbf{W}_t\|_F^2$ and $\sum_{i=1}^t \gamma \|\Omega_i\|_F^2$ to both sides of inequality (71), we obtain:

$$\begin{aligned} \hat{\mathcal{G}}(t) &= \|\mathbf{Y} - (\Phi_t + \mathbf{F}_t \Omega_{t+1}) \mathbf{W}_t\|_F^2 + \lambda \|\mathbf{W}_t\|_F^2 + \sum_{i=1}^{t+1} \gamma \|\Omega_i\|_F^2 \\ &\leq \|\mathbf{Y} - \Phi_t \mathbf{W}_t\|_F^2 + \lambda \|\mathbf{W}_t\|_F^2 + \sum_{i=1}^t \gamma \|\Omega_i\|_F^2 = \mathcal{G}(t). \end{aligned} \quad (72)$$

Second, we further analyze the substep of constructing the global classifier \mathbf{W}_{t+1} while keeping the feature representation Φ_{t+1} and transformation matrices $\{\Omega_i\}_{i=1}^{t+1}$ fixed. According to Lemma 1, the \mathbf{W}_{t+1} represents the optimal solution that minimizes the regularized empirical risk:

$$\mathbf{W}_{t+1} = \arg \min_{\mathbf{W}} \|\mathbf{Y} - \Phi_{t+1} \mathbf{W}\|_F^2 + \lambda \|\mathbf{W}\|_F^2. \quad (73)$$

1188 Therefore, we can obtain:
 1189

$$\| \mathbf{Y} - \Phi_{t+1} \mathbf{W}_{t+1} \|_{\text{F}}^2 + \lambda \| \mathbf{W}_{t+1} \|_{\text{F}}^2 \leq \| \mathbf{Y} - \Phi_{t+1} \mathbf{W} \|_{\text{F}}^2 + \lambda \| \mathbf{W} \|_{\text{F}}^2, \forall \mathbf{W}, \quad (74)$$

1190 where $\mathbf{W} = \mathbf{W}_t$ serves as a special case, from which we further derive:
 1191

$$\| \mathbf{Y} - \Phi_{t+1} \mathbf{W}_{t+1} \|_{\text{F}}^2 + \lambda \| \mathbf{W}_{t+1} \|_{\text{F}}^2 \leq \| \mathbf{Y} - \Phi_{t+1} \mathbf{W}_t \|_{\text{F}}^2 + \lambda \| \mathbf{W}_t \|_{\text{F}}^2. \quad (75)$$

1192 By adding $\sum_{i=1}^{t+1} \gamma \| \Omega_i \|_{\text{F}}^2$ to both sides of inequality (75), we obtain:
 1193

$$\begin{aligned} \mathcal{G}(t+1) &= \| \mathbf{Y} - \Phi_{t+1} \mathbf{W}_{t+1} \|_{\text{F}}^2 + \lambda \| \mathbf{W}_{t+1} \|_{\text{F}}^2 + \sum_{i=1}^{t+1} \gamma \| \Omega_i \|_{\text{F}}^2 \\ &\leq \| \mathbf{Y} - \Phi_{t+1} \mathbf{W}_t \|_{\text{F}}^2 + \lambda \| \mathbf{W}_t \|_{\text{F}}^2 + \sum_{i=1}^{t+1} \gamma \| \Omega_i \|_{\text{F}}^2 = \hat{\mathcal{G}}(t). \end{aligned} \quad (76)$$

1194 In summary, based on (72) and (76), it holds for each layer that:
 1195

$$\mathcal{G}(t) \geq \hat{\mathcal{G}}(t) \geq \mathcal{G}(t+1), \forall t \in [0, T]. \quad (77)$$

1196 Therefore, the sequence of the regularized empirical risks $\{\mathcal{G}(t)\}_{t=0}^T$ in our DeepAFL remains
 1197 monotonically non-increasing. Given that $\mathcal{G}(t) \geq 0$, by the well-known Monotone Convergence
 1198 Theorem, it necessarily converges to a limit $\mathcal{G}^* \geq 0$. Moreover, since the sequence begins at $\mathcal{G}(0)$
 1199 and decreases at each layer, it holds that $\mathcal{G}^* \leq \mathcal{G}(0)$.
 1200

■

1210 C APPENDIX FOR PRIVACY ANALYSES

1211 In this section, we provide detailed privacy analyses of our DeepAFL, providing multi-layered analyses to substantiate our DeepAFL's robust privacy preservation, as follows:
 1212

1213 First, in our DeepAFL, all clients need only submit their locally computed *Auto-Correlation* and
 1214 *Cross-Correlation Matrices*, and numerous studies (Tan et al., 2022b; He et al., 2025b; Fan et al.,
 1215 2025b) have shown the privacy advantages of uploading such information. Specifically, according
 1216 to (15), the uploaded *Cross-Correlation Matrices* Π_t^k and Υ_t^k are essentially prototype matrices,
 1217 where each column corresponds to the unaveraged class and residual prototypes for the correspond-
 1218 ing class. Consequently, the information uploaded within DeepAFL is akin to that of prototype-
 1219 based FL methods (Tan et al., 2022b), which have provided a detailed introduction to the inher-
 1220 ent privacy benefits of transmitting prototypes. Additionally, existing analytical learning-based FL
 1221 methods (He et al., 2025b; Fan et al., 2025b) also require the uploading of *Auto-Correlation* and
 1222 *Cross-Correlation Matrices* similarly. Thus, our DeepAFL shares the same privacy advantages as
 1223 these existing methods.
 1224

1225 Second, since clients are not required to upload their local dataset sizes N_k , it is pretty hard to
 1226 infer the clients' private data (i.e., \mathbf{X}_t^k , \mathbf{Y}_t^k , Φ_t^k , \mathbf{F}_t^k , and \mathbf{R}_t^k) from the information they upload.
 1227 Specifically, in our DeepAFL, each client uploads only its local matrices $\mathbf{G}_t^k \in \mathbb{R}^{d_{\Phi} \times d_{\Phi}}$, $\mathbf{H}_t^k \in$
 1228 $\mathbb{R}^{d_{\Phi} \times C}$, $\Pi_t^k \in \mathbb{R}^{d_{\Phi} \times d_{\Phi}}$, and $\Upsilon_t^k \in \mathbb{R}^{d_{\Phi} \times C}$. Evidently, clients need not upload their local dataset sizes
 1229 N_k , which also cannot be deduced from the dimensions of the aforementioned matrices. Moreover,
 1230 the dimensions of all private data are directly tied to N_k , i.e., $\mathbf{X}_t^k \in \mathbb{R}^{N_k \times d_{\text{H}} \times d_{\text{W}} \times d_{\text{C}}}$, $\mathbf{Y}_t^k \in \mathbb{R}^{N_k \times C}$,
 1231 $\Phi_t^k \in \mathbb{R}^{N_k \times d_{\Phi}}$, $\mathbf{F}_t^k \in \mathbb{R}^{N_k \times d_{\Phi}}$, and $\mathbf{R}_t^k \in \mathbb{R}^{N_k \times C}$. Therefore, without knowledge of N_k , the
 1232 dimensions of the aforementioned private data cannot be determined. This results in infinitely many
 1233 possible instantiations, rendering it fundamentally impossible to infer the clients' private data.
 1234

1235 Third, many existing privacy-preserving techniques can be directly integrated into our DeepAFL, en-
 1236 abling the server to operate without requiring access to each client's specific matrices. Specifically,
 1237 in our DeepAFL, the server only needs to utilize the aggregated values of the clients' uploaded ma-
 1238 trices for constructing the global classifier and residual block, without needing the individual client
 1239 upload results. Therefore, this can be regarded as a form of *Secure Multi-Party Computation*. Existing
 1240 *Secure Aggregation Protocols* (Bonawitz et al., 2016; 2017; So et al., 2023) can be employed to
 1241 aggregate clients' uploaded matrices without accessing their specific matrices. Additionally, tech-
 1242 niques such as *Homomorphic Encryption* (Hu & Li, 2025) and *Differential Privacy* (Hu et al., 2024)
 1243 can also be integrated into our DeepAFL to further preserve client privacy.
 1244

1242 **D APPENDIX FOR EFFICIENCY ANALYSES**
 1243

1244 In this section, we provide detailed efficiency analyses of our DeepAFL, focusing on the com-
 1245 putation and communication complexities on both the client and server sides. Overall, for con-
 1246 structing all T layers, the total computation and communication complexities of each client are
 1247 $\mathcal{O}(TN_k(d_\Phi^2 + d_\Phi d_F + d_F^2))$ and $\mathcal{O}(T(d_\Phi^2 + d_F^2))$, while those for the server are
 1248 $\mathcal{O}(T(d_\Phi^3 + d_F^3 + N_k d_F^2))$ and $\mathcal{O}(TKC(d_\Phi + d_F))$ respectively. Notably, although the complexities of our DeepAFL are ap-
 1249 proximately T times greater than those of AFL to construct a deep network for superior performance,
 1250 they remain far lower than those of gradient-based methods. The detailed analyses are as follows:

1251 **Analysis of Computation Complexity:** Here, we thoroughly analyze the computation complex-
 1252 ity within our DeepAFL. Notably, each layer’s construction within our DeepAFL merely requires
 1253 single-round lightweight analytic computation by the clients and server, obviating the need for multi-
 1254 round gradient-based iterative updates.

1255 • **The Client Side:** First, for the construction of each layer t , each client k computes its local
 1256 *Feature Auto-Correlation Matrix* \mathbf{G}_t^k and *Label Cross-Correlation Matrix* \mathbf{H}_t^k via (11). Given
 1257 that $\Phi_t^k \in \mathbb{R}^{N_k \times d_\Phi}$ and $\mathbf{Y}_t^k \in \mathbb{R}^{N_k \times C}$, the computation complexity of the above calculations is
 1258 $\mathcal{O}(N_k d_\Phi^2 + CN_k d_\Phi)$. Second, each client further constructs its local hidden random features \mathbf{F}_t^k and
 1259 local residual matrix \mathbf{R}_t^k via (14), with a complexity of $\mathcal{O}(N_k d_\Phi d_F + CN_k d_\Phi)$. Third, each client
 1260 computes its local *Hidden Auto-Correlation Matrix* Π_t^k and *Residual Cross-Correlation Matrix* Υ_t^k
 1261 via (15). Given that $\mathbf{F}_t^k \in \mathbb{R}^{N_k \times d_F}$ and $\mathbf{R}_t^k \in \mathbb{R}^{d_F \times C}$, the computation complexity of executing (15)
 1262 can be calculated $\mathcal{O}(N_k d_F^2 + CN_k d_F)$. Fourth, the client update the next-layer feature Φ_{t+1}^k with a
 1263 complexity of $\mathcal{O}(N_k d_\Phi d_F)$. In summary, the complexity for each client to construct each layer is
 1264 $\mathcal{O}(CN_k(d_F + d_\Phi) + N_k(d_\Phi^2 + d_\Phi d_F + d_F^2))$, and the overall complexity for constructing all T layers
 1265 is $\mathcal{O}(TCN_k(d_F + d_\Phi) + TN_k(d_\Phi^2 + d_\Phi d_F + d_F^2))$.

1266 • **The Server Side:** First, for the construction of each layer t , the server aggregates all received
 1267 $\{\mathbf{G}_t^k\}_{k=1}^K$ and $\{\mathbf{H}_t^k\}_{k=1}^K$ and subsequently derive the global classifier \mathbf{W}_t via (12) and (13). Given
 1268 that $\mathbf{G}_t^k \in \mathbb{R}^{d_\Phi \times d_\Phi}$ and $\mathbf{H}_t^k \in \mathbb{R}^{d_\Phi \times C}$, the computation complexity of these calculations is $\mathcal{O}(d_\Phi^3 +$
 1269 $Kd_\Phi^2 + Cd_\Phi^2 + KCd_\Phi)$. Second, the server aggregates all received $\{\Pi_t^k\}_{k=1}^K$ and $\{\Upsilon_t^k\}_{k=1}^K$ via
 1270 (16) with a complexity of $\mathcal{O}(Kd_F^2 + KCd_F)$. Third, the server performs spectral decompositions
 1271 and subsequently derives the transformation matrix Ω_{t+1} via (17) and (18) with a computation
 1272 complexity of $\mathcal{O}(d_F^3 + Cd_F^2 + N_k d_F^2 + CN_k d_F)$. In summary, the complexity for the server to
 1273 construct each layer is $\mathcal{O}(d_\Phi^3 + d_F^3 + N_k d_F^2 + (K+C)(d_F^2 + d_\Phi^2) + C(Kd_\Phi + kd_F + N_k d_F))$, and the
 1274 overall complexity for constructing all T layers is $\mathcal{O}(T(d_\Phi^3 + d_F^3 + N_k d_F^2) + T(K+C)(d_F^2 + d_\Phi^2) +$
 1275 $TC(Kd_\Phi + kd_F + N_k d_F))$.

1276 **Analysis of Communication Complexity:** Here, we further thoroughly analyze the communica-
 1277 tion complexity within our DeepAFL. Notably, DeepAFL requires only a single communication
 1278 exchange between the client and server for constructing each layer, in contrast to gradient-based
 1279 methods that depend on multiple rounds of communication.

1280 • **The Client Side:** According to Algorithm 1, during the construction of each layer t , each client k is
 1281 only required to upload its local *Feature Auto-Correlation Matrix* \mathbf{G}_t^k , *Label Cross-Correlation Ma-
 1282 trix* \mathbf{H}_t^k , *Hidden Auto-Correlation Matrix* Π_t^k , and *Residual Cross-Correlation Matrix* Υ_t^k . Given
 1283 that $\mathbf{G}_t^k \in \mathbb{R}^{d_\Phi \times d_\Phi}$, $\mathbf{H}_t^k \in \mathbb{R}^{d_\Phi \times C}$, $\Pi_t^k \in \mathbb{R}^{d_F \times d_F}$, and $\Upsilon_t^k \in \mathbb{R}^{d_F \times C}$, the client’s total communica-
 1284 tion complexity for constructing all T layers can be derived as $\mathcal{O}(T(d_\Phi^2 + d_F^2) + TC(d_\Phi + d_F))$.

1285 • **The Server Side:** For each layer t , the server similarly needs only to transmit once to all clients the
 1286 global classifier $\mathbf{W}_t \in \mathbb{R}^{d_\Phi \times C}$ and the transformation matrix $\Omega_t \in \mathbb{R}^{d_F \times C}$, with a communication
 1287 complexity of $\mathcal{O}(KC(d_\Phi + d_F))$. Consequently, the server’s total communication complexity for
 1288 constructing all T layers can be expressed as $\mathcal{O}(TKC(d_\Phi + d_F))$.

1289 **Simplification of Complexity Results:** Considering the complexity of the complexity results de-
 1290 rived earlier, we further simplify them based on the relative magnitudes of their constituent terms.
 1291 Specifically, in practical scenarios, the total number of layers T , the number of clients K , and the
 1292 number of classes C are substantially smaller than the sample size N_k in each client’s local dataset,
 1293 as well as the feature dimensions d_Φ and d_F . Consequently, the computation complexity for each
 1294 client and the server can be simplified to $\mathcal{O}(TN_k(d_\Phi^2 + d_\Phi d_F + d_F^2))$ and $\mathcal{O}(T(d_\Phi^3 + d_F^3 + N_k d_F^2))$.
 1295 Similarly, the communication complexity for each client can also be simplified to $\mathcal{O}(T(d_\Phi^2 + d_F^2))$.

1296 E APPENDIX FOR EXPERIMENTAL EVALUATIONS

1298 E.1 ADDITIONAL RESULTS FOR EXPERIMENTAL EVALUATIONS

1300 Table 2: Performance comparisons of the top-1 accuracy (%) among our DeepAFL and the baselines,
 1301 on the CIFAR-10. The best result is highlighted in **bold**, and the second-best result is underlined. All
 1302 the experiments were conducted three times, and the results are shown as Mean \pm Standard Error. All
 1303 the improvements of our DeepAFL were validated by Chi-squared tests at the level of $p < 0.05$. The
 1304 results of all baselines are directly obtained from the given benchmark in AFL (He et al., 2025b).

Baseline	CIFAR-10			
	Non-IID-1		Non-IID-2	
	$\alpha = 0.1$	$\alpha = 0.05$	$s = 4$	$s = 2$
FedAvg (2017)	64.02 \pm 0.18	60.52 \pm 0.39	68.47 \pm 0.13	57.81 \pm 0.03
FedProx (2020)	64.07 \pm 0.08	60.39 \pm 0.09	68.46 \pm 0.08	57.61 \pm 0.12
MOON (2021a)	63.84 \pm 0.03	60.28 \pm 0.17	68.47 \pm 0.15	57.72 \pm 0.15
FedGen (2021)	64.14 \pm 0.24	60.65 \pm 0.19	68.24 \pm 0.28	57.02 \pm 0.18
FedDyn (2021)	64.77 \pm 0.11	60.35 \pm 0.54	73.50 \pm 0.11	64.07 \pm 0.09
FedNTD (2022)	64.64 \pm 0.02	61.16 \pm 0.33	70.24 \pm 0.11	58.77 \pm 0.18
FedDisco (2023b)	63.83 \pm 0.08	59.90 \pm 0.05	65.04 \pm 0.11	58.78 \pm 0.02
AFL (2025b)	<u>80.75</u> \pm 0.00	<u>80.75</u> \pm 0.00	<u>80.75</u> \pm 0.00	<u>80.75</u> \pm 0.00
DeepAFL ($T = 5$)	85.20 \pm 0.05	85.20 \pm 0.05	85.20 \pm 0.05	85.20 \pm 0.05
DeepAFL ($T = 10$)	85.93 \pm 0.09	85.93 \pm 0.09	85.93 \pm 0.09	85.93 \pm 0.09
DeepAFL ($T = 20$)	86.43 \pm 0.07	86.43 \pm 0.07	86.43 \pm 0.07	86.43 \pm 0.07
Improvement \uparrow	5.68 $p < 0.05$			

1322 Table 3: Accuracy of our DeepAFL under varying regularization parameters λ .

Datasets	Layers	$\lambda = 0.01$	$\lambda = 0.05$	$\lambda = 0.1$	$\lambda = 0.5$	$\lambda = 1$	$\lambda = 5$	$\lambda = 10$
CIFAR-10	$T = 5$	85.14 \pm 0.05	85.15 \pm 0.04	85.14 \pm 0.03	85.11 \pm 0.05	85.13 \pm 0.04	85.17 \pm 0.05	85.20 \pm 0.05
	$T = 10$	85.87 \pm 0.08	85.90 \pm 0.09	85.86 \pm 0.08	85.92 \pm 0.08	85.91 \pm 0.08	85.92 \pm 0.10	85.93 \pm 0.09
	$T = 20$	86.34 \pm 0.07	86.31 \pm 0.08	86.37 \pm 0.07	86.27 \pm 0.06	86.33 \pm 0.08	86.39 \pm 0.08	86.43 \pm 0.07
CIFAR-100	$T = 5$	64.61 \pm 0.01	64.62 \pm 0.02	64.59 \pm 0.02	64.61 \pm 0.02	64.62 \pm 0.02	64.64 \pm 0.01	64.66 \pm 0.02
	$T = 10$	65.95 \pm 0.07	65.96 \pm 0.06	65.95 \pm 0.08	65.94 \pm 0.06	65.94 \pm 0.06	65.91 \pm 0.04	65.91 \pm 0.05
	$T = 20$	66.91 \pm 0.06	66.91 \pm 0.06	66.92 \pm 0.06	66.96 \pm 0.05	66.94 \pm 0.03	66.87 \pm 0.03	66.95 \pm 0.03
Tiny-ImageNet	$T = 5$	60.23 \pm 0.02	60.24 \pm 0.02	60.25 \pm 0.01	60.24 \pm 0.02	60.22 \pm 0.02	60.26 \pm 0.02	60.23 \pm 0.02
	$T = 10$	61.34 \pm 0.08	61.33 \pm 0.08	61.35 \pm 0.07	61.36 \pm 0.07	61.34 \pm 0.08	61.32 \pm 0.08	61.32 \pm 0.07
	$T = 20$	62.34 \pm 0.02	62.33 \pm 0.04	62.34 \pm 0.02	62.32 \pm 0.02	62.35 \pm 0.03	62.33 \pm 0.02	62.34 \pm 0.01

1336 Table 4: Accuracy of our DeepAFL under varying regularization parameters γ .

Datasets	Layers	$\gamma = 0.01$	$\gamma = 0.05$	$\gamma = 0.1$	$\gamma = 0.5$	$\gamma = 1$	$\gamma = 5$	$\gamma = 10$
CIFAR-10	$T = 5$	85.06 \pm 0.05	85.14 \pm 0.06	85.15 \pm 0.04	85.10 \pm 0.02	85.10 \pm 0.04	85.04 \pm 0.04	84.99 \pm 0.04
	$T = 10$	85.87 \pm 0.06	85.90 \pm 0.07	85.91 \pm 0.26	85.86 \pm 0.10	85.84 \pm 0.08	85.78 \pm 0.08	85.65 \pm 0.07
	$T = 20$	86.21 \pm 0.09	86.21 \pm 0.08	86.32 \pm 0.08	86.37 \pm 0.06	86.33 \pm 0.08	86.26 \pm 0.07	86.23 \pm 0.07
CIFAR-100	$T = 5$	64.72 \pm 0.07	64.66 \pm 0.02	64.63 \pm 0.01	64.58 \pm 0.04	64.49 \pm 0.00	63.77 \pm 0.07	63.21 \pm 0.10
	$T = 10$	65.96 \pm 0.05	65.94 \pm 0.06	65.95 \pm 0.08	65.81 \pm 0.08	65.69 \pm 0.06	64.82 \pm 0.15	64.26 \pm 0.15
	$T = 20$	66.98 \pm 0.04	66.94 \pm 0.04	66.94 \pm 0.03	66.82 \pm 0.07	66.74 \pm 0.07	66.20 \pm 0.09	65.61 \pm 0.16
Tiny-ImageNet	$T = 5$	60.31 \pm 0.04	60.28 \pm 0.02	60.24 \pm 0.02	60.13 \pm 0.04	59.89 \pm 0.01	59.04 \pm 0.04	58.41 \pm 0.05
	$T = 10$	61.37 \pm 0.07	61.36 \pm 0.07	61.35 \pm 0.08	61.28 \pm 0.12	61.19 \pm 0.11	60.44 \pm 0.04	59.65 \pm 0.07
	$T = 20$	62.35 \pm 0.01	62.33 \pm 0.09	62.26 \pm 0.02	62.24 \pm 0.05	62.14 \pm 0.02	61.53 \pm 0.03	60.98 \pm 0.02

1350
1351Table 5: Accuracy of our DeepAFL under varying activation functions $\sigma(\cdot)$.

Datasets	Layers	None	GELU	ReLU	LeakyReLU	Tanh	Hardwish	Softshrink
CIFAR-10	$T = 5$	83.23 ± 0.10	85.13 ± 0.04	84.37 ± 0.09	84.49 ± 0.12	84.42 ± 0.07	84.78 ± 0.03	83.77 ± 0.12
	$T = 10$	83.08 ± 0.17	85.89 ± 0.10	85.04 ± 0.12	84.54 ± 0.49	84.92 ± 0.12	85.41 ± 0.04	84.13 ± 0.05
	$T = 20$	82.99 ± 0.06	86.34 ± 0.05	85.47 ± 0.17	83.85 ± 1.48	85.62 ± 0.05	86.05 ± 0.06	84.25 ± 0.07
CIFAR-100	$T = 5$	60.82 ± 0.09	64.62 ± 0.01	64.26 ± 0.02	64.29 ± 0.01	63.28 ± 0.12	64.02 ± 0.02	62.89 ± 0.13
	$T = 10$	60.82 ± 0.09	65.98 ± 0.08	65.39 ± 0.09	65.48 ± 0.10	64.27 ± 0.17	65.07 ± 0.20	63.91 ± 0.14
	$T = 20$	60.83 ± 0.09	66.95 ± 0.02	66.68 ± 0.08	66.72 ± 0.10	64.02 ± 1.02	66.04 ± 0.06	64.78 ± 0.08
Tiny-ImageNet	$T = 5$	56.71 ± 0.07	60.23 ± 0.01	60.17 ± 0.04	60.16 ± 0.02	58.77 ± 0.04	59.40 ± 0.04	58.80 ± 0.07
	$T = 10$	56.70 ± 0.08	61.33 ± 0.07	61.35 ± 0.05	61.36 ± 0.08	59.68 ± 0.07	60.48 ± 0.09	59.83 ± 0.04
	$T = 20$	56.70 ± 0.08	62.33 ± 0.02	62.31 ± 0.06	62.31 ± 0.07	60.74 ± 0.04	61.47 ± 0.01	60.70 ± 0.09

1362

Table 6: Accuracy of our DeepAFL under varying dimensions d_Φ . For a comprehensive and consistent analysis, all the experimental results were obtained by fixing the other dimension d_F at 1024.

Datasets	Layers	$d_\Phi = 2^7$	$d_\Phi = 2^8$	$d_\Phi = 2^9$	$d_\Phi = 2^{10}$	$d_\Phi = 2^{11}$	$d_\Phi = 2^{12}$	$d_\Phi = 2^{13}$
CIFAR-10	$T = 5$	81.48 ± 0.30	83.31 ± 0.12	84.58 ± 0.08	85.07 ± 0.02	85.48 ± 0.03	85.76 ± 0.7	85.57 ± 0.33
	$T = 10$	82.28 ± 0.30	84.41 ± 0.01	85.55 ± 0.11	85.96 ± 0.10	85.94 ± 0.01	85.94 ± 0.12	85.77 ± 0.28
	$T = 20$	82.68 ± 0.27	84.95 ± 0.13	86.01 ± 0.01	86.28 ± 0.15	86.41 ± 0.01	86.48 ± 0.11	83.12 ± 0.39
CIFAR-100	$T = 5$	57.61 ± 0.23	62.21 ± 0.14	64.06 ± 0.07	64.63 ± 0.02	65.39 ± 0.05	66.46 ± 0.12	67.46 ± 0.01
	$T = 10$	58.20 ± 0.19	63.45 ± 0.19	65.29 ± 0.07	65.93 ± 0.10	66.36 ± 0.04	67.20 ± 0.06	67.19 ± 0.44
	$T = 20$	58.51 ± 0.25	64.29 ± 0.36	66.64 ± 0.10	66.93 ± 0.04	67.39 ± 0.06	67.85 ± 0.05	67.94 ± 0.02
Tiny-ImageNet	$T = 5$	44.94 ± 0.24	56.91 ± 0.14	59.29 ± 0.26	60.21 ± 0.04	60.77 ± 0.23	61.84 ± 0.19	62.75 ± 0.22
	$T = 10$	44.11 ± 0.11	57.51 ± 0.18	60.68 ± 0.24	61.43 ± 0.04	61.90 ± 0.15	62.42 ± 0.24	63.13 ± 0.41
	$T = 20$	43.10 ± 0.19	58.09 ± 0.15	61.70 ± 0.19	62.33 ± 0.03	62.88 ± 0.18	63.11 ± 0.29	63.70 ± 0.19

1376

Table 7: Accuracy of our DeepAFL under varying dimensions d_F . For a comprehensive and consistent analysis, all the experimental results were obtained by fixing the other dimension d_Φ at 1024.

Datasets	Layers	$d_F = 2^7$	$d_F = 2^8$	$d_F = 2^9$	$d_F = 2^{10}$	$d_F = 2^{11}$	$d_F = 2^{12}$	$d_F = 2^{13}$
CIFAR-10	$T = 5$	83.48 ± 0.10	83.76 ± 0.09	84.24 ± 0.14	85.14 ± 0.04	86.16 ± 0.06	86.51 ± 0.04	85.55 ± 0.11
	$T = 10$	83.58 ± 0.11	84.03 ± 0.10	84.84 ± 0.08	85.89 ± 0.09	86.50 ± 0.13	86.42 ± 0.04	84.90 ± 0.09
	$T = 20$	83.83 ± 0.11	84.49 ± 0.09	85.34 ± 0.04	86.31 ± 0.08	86.53 ± 0.05	85.59 ± 0.10	83.40 ± 0.10
CIFAR-100	$T = 5$	61.22 ± 0.06	61.82 ± 0.11	63.11 ± 0.11	64.63 ± 0.01	66.49 ± 0.12	67.81 ± 0.01	67.38 ± 0.09
	$T = 10$	61.54 ± 0.10	62.57 ± 0.10	64.08 ± 0.04	65.96 ± 0.08	67.59 ± 0.06	67.92 ± 0.08	66.12 ± 0.04
	$T = 20$	61.92 ± 0.14	63.41 ± 0.12	65.23 ± 0.18	66.94 ± 0.03	68.20 ± 0.03	67.14 ± 0.07	64.10 ± 0.10
Tiny-ImageNet	$T = 5$	56.92 ± 0.02	57.55 ± 0.11	58.69 ± 0.05	60.24 ± 0.02	61.82 ± 0.07	63.36 ± 0.08	63.82 ± 0.04
	$T = 10$	57.23 ± 0.07	58.30 ± 0.09	59.58 ± 0.08	61.34 ± 0.07	62.89 ± 0.06	63.84 ± 0.10	62.89 ± 0.04
	$T = 20$	57.79 ± 0.05	59.18 ± 0.08	60.86 ± 0.03	62.34 ± 0.03	63.85 ± 0.11	63.48 ± 0.03	60.81 ± 0.08

1390

Table 8: Accuracy of our DeepAFL under varying dimensions d_Φ and d_F . For a comprehensive and consistent analysis, all the experimental results were obtained by ensuring $d_\Phi = d_F = d$.

Datasets	Layers	$d = 2^7$	$d = 2^8$	$d = 2^9$	$d = 2^{10}$	$d = 2^{11}$	$d = 2^{12}$	$d = 2^{13}$
CIFAR-10	$T = 5$	77.01 ± 0.31	80.62 ± 0.17	83.54 ± 0.09	85.04 ± 0.12	86.19 ± 0.03	86.54 ± 0.04	/
	$T = 10$	77.96 ± 0.21	81.62 ± 0.19	84.24 ± 0.08	85.77 ± 0.04	86.58 ± 0.04	86.47 ± 0.12	/
	$T = 20$	78.69 ± 0.11	82.38 ± 0.08	85.07 ± 0.07	86.22 ± 0.02	86.71 ± 0.13	84.76 ± 0.34	/
CIFAR-100	$T = 5$	49.88 ± 0.27	57.15 ± 0.10	61.74 ± 0.04	64.63 ± 0.00	66.86 ± 0.13	68.59 ± 0.08	/
	$T = 10$	51.36 ± 0.13	58.76 ± 0.12	63.22 ± 0.02	65.96 ± 0.06	68.20 ± 0.06	69.10 ± 0.03	/
	$T = 20$	53.21 ± 0.10	60.26 ± 0.07	64.72 ± 0.11	66.93 ± 0.03	68.75 ± 0.06	69.12 ± 0.08	/
Tiny-ImageNet	$T = 5$	44.94 ± 0.23	53.11 ± 0.13	57.44 ± 0.05	60.22 ± 0.03	62.29 ± 0.05	63.95 ± 0.07	/
	$T = 10$	44.20 ± 0.09	54.47 ± 0.21	58.72 ± 0.07	61.17 ± 0.01	63.13 ± 0.06	64.63 ± 0.10	/
	$T = 20$	43.83 ± 0.12	56.18 ± 0.11	59.92 ± 0.02	62.34 ± 0.02	64.08 ± 0.06	64.68 ± 0.04	/

1403

Table 9: Analyses of our DeepAFL’s representation learning capability on the CIFAR-10 with increasing network depth. “Training Acc” and “Testing Acc” show the accuracy (%) on the training and test sets. “Time Cost” denotes the required training time (s). The symbol Δ represents the difference between two consecutive cases, which reflects the marginal effect of deepening the network.

Layers	Training Acc	$\Delta_{\text{Training Acc}}$	Testing Acc	$\Delta_{\text{Testing Acc}}$	Time Cost	$\Delta_{\text{Time Cost}}$
AFL	83.81 \pm 0.00	/	80.75 \pm 0.00	/	52.36 \pm 0.24	/
$T = 0$	85.35 \pm 0.02	1.54 \pm 0.02	83.29 \pm 0.20	2.54 \pm 0.20	58.02 \pm 0.17	5.66 \pm 0.13
$T = 5$	89.03 \pm 0.03	3.68 \pm 0.02	85.13 \pm 0.08	1.84 \pm 0.08	65.08 \pm 0.39	7.06 \pm 0.22
$T = 10$	90.61 \pm 0.05	1.58 \pm 0.03	85.89 \pm 0.09	0.75 \pm 0.11	71.91 \pm 0.28	6.83 \pm 0.29
$T = 15$	91.80 \pm 0.04	1.19 \pm 0.04	86.18 \pm 0.08	0.29 \pm 0.04	78.86 \pm 0.24	6.95 \pm 0.15
$T = 20$	92.74 \pm 0.03	0.94 \pm 0.01	86.34 \pm 0.06	0.16 \pm 0.06	85.71 \pm 0.22	6.85 \pm 0.17
$T = 25$	93.49 \pm 0.02	0.75 \pm 0.01	86.40 \pm 0.06	0.06 \pm 0.06	92.57 \pm 0.19	6.86 \pm 0.11
$T = 30$	94.10 \pm 0.03	0.61 \pm 0.01	86.46 \pm 0.06	0.06 \pm 0.05	99.29 \pm 0.33	6.72 \pm 0.19
$T = 35$	94.63 \pm 0.03	0.53 \pm 0.03	86.72 \pm 0.06	0.26 \pm 0.01	105.79 \pm 0.31	6.50 \pm 0.21
$T = 40$	95.08 \pm 0.03	0.44 \pm 0.02	86.76 \pm 0.06	0.04 \pm 0.02	112.40 \pm 0.21	6.61 \pm 0.06
$T = 45$	95.43 \pm 0.03	0.35 \pm 0.01	86.77 \pm 0.06	0.01 \pm 0.00	119.00 \pm 0.17	6.60 \pm 0.13
$T = 50$	95.79 \pm 0.03	0.36 \pm 0.00	86.72 \pm 0.06	-0.05 \pm 0.11	125.72 \pm 0.31	6.72 \pm 0.08

Table 10: Analyses of our DeepAFL’s representation learning capability on the CIFAR-100 with increasing network depth. “Training Acc” and “Testing Acc” show the accuracy (%) on the training and test sets. “Time Cost” denotes the required training time (s). The symbol Δ represents the difference between two consecutive cases, which reflects the marginal effect of deepening the network.

Layers	Training Acc	$\Delta_{\text{Training Acc}}$	Testing Acc	$\Delta_{\text{Testing Acc}}$	Time Cost	$\Delta_{\text{Time Cost}}$
AFL	61.55 \pm 0.00	/	58.56 \pm 0.00	/	50.05 \pm 0.29	/
$T = 0$	65.66 \pm 0.04	4.11 \pm 0.04	60.81 \pm 0.15	2.25 \pm 0.15	56.90 \pm 0.17	6.85 \pm 0.19
$T = 5$	74.93 \pm 0.04	9.27 \pm 0.04	64.62 \pm 0.02	3.81 \pm 0.08	66.48 \pm 0.21	9.58 \pm 0.23
$T = 10$	79.41 \pm 0.08	4.48 \pm 0.04	65.98 \pm 0.13	1.36 \pm 0.09	75.25 \pm 0.33	8.77 \pm 0.03
$T = 15$	82.69 \pm 0.09	3.28 \pm 0.04	66.59 \pm 0.03	0.61 \pm 0.09	83.50 \pm 0.12	8.25 \pm 0.09
$T = 20$	85.15 \pm 0.03	2.45 \pm 0.06	66.95 \pm 0.04	0.36 \pm 0.02	91.74 \pm 0.29	8.24 \pm 0.22
$T = 25$	87.21 \pm 0.02	2.06 \pm 0.03	67.40 \pm 0.10	0.45 \pm 0.05	100.36 \pm 0.37	8.62 \pm 0.18
$T = 30$	88.84 \pm 0.06	1.63 \pm 0.03	67.71 \pm 0.06	0.31 \pm 0.03	108.37 \pm 0.09	8.01 \pm 0.03
$T = 35$	90.14 \pm 0.06	1.30 \pm 0.00	67.97 \pm 0.12	0.26 \pm 0.07	116.50 \pm 0.19	8.13 \pm 0.11
$T = 40$	91.29 \pm 0.09	1.15 \pm 0.02	68.19 \pm 0.07	0.22 \pm 0.03	124.74 \pm 0.17	8.24 \pm 0.03
$T = 45$	92.15 \pm 0.11	0.86 \pm 0.01	68.32 \pm 0.11	0.13 \pm 0.05	133.06 \pm 0.09	8.32 \pm 0.02
$T = 50$	92.93 \pm 0.06	0.78 \pm 0.03	68.51 \pm 0.09	0.19 \pm 0.02	141.38 \pm 0.08	8.32 \pm 0.03

Table 11: Analyses of our DeepAFL’s representation learning capability on the Tiny-ImageNet with increasing network depth. “Training Acc” and “Testing Acc” show the accuracy (%) on the training and test sets. “Time Cost” denotes the required training time (s). The symbol Δ represents the difference between two consecutive cases, which reflects the marginal effect of deepening the network.

Layers	Training Acc	$\Delta_{\text{Training Acc}}$	Testing Acc	$\Delta_{\text{Testing Acc}}$	Time Cost	$\Delta_{\text{Time Cost}}$
AFL	57.39 \pm 0.00	/	54.67 \pm 0.00	/	125.31 \pm 0.38	/
$T = 0$	60.30 \pm 0.02	2.91 \pm 0.02	56.72 \pm 0.14	2.05 \pm 0.14	130.72 \pm 0.22	5.41 \pm 0.08
$T = 5$	66.89 \pm 0.06	6.59 \pm 0.02	60.24 \pm 0.02	3.52 \pm 0.07	145.73 \pm 0.13	15.01 \pm 0.11
$T = 10$	70.14 \pm 0.03	3.24 \pm 0.02	61.34 \pm 0.13	1.10 \pm 0.06	160.75 \pm 0.23	15.02 \pm 0.03
$T = 15$	72.64 \pm 0.05	2.50 \pm 0.02	61.94 \pm 0.08	0.60 \pm 0.11	175.76 \pm 0.34	15.01 \pm 0.02
$T = 20$	74.75 \pm 0.02	2.11 \pm 0.03	62.34 \pm 0.02	0.40 \pm 0.03	190.01 \pm 0.09	14.25 \pm 0.11
$T = 25$	76.58 \pm 0.03	1.83 \pm 0.01	62.66 \pm 0.05	0.32 \pm 0.02	204.64 \pm 0.12	14.63 \pm 0.13
$T = 30$	78.31 \pm 0.02	1.73 \pm 0.02	62.94 \pm 0.03	0.29 \pm 0.02	219.56 \pm 0.11	14.92 \pm 0.03
$T = 35$	79.87 \pm 0.03	1.56 \pm 0.01	63.24 \pm 0.16	0.30 \pm 0.09	234.99 \pm 0.23	15.53 \pm 0.07
$T = 40$	81.32 \pm 0.03	1.45 \pm 0.02	63.48 \pm 0.18	0.24 \pm 0.05	249.95 \pm 0.13	14.94 \pm 0.04
$T = 45$	82.62 \pm 0.04	1.30 \pm 0.03	63.65 \pm 0.18	0.16 \pm 0.01	264.31 \pm 0.05	14.36 \pm 0.11
$T = 50$	83.82 \pm 0.09	1.20 \pm 0.03	63.74 \pm 0.04	0.10 \pm 0.08	279.68 \pm 0.14	15.37 \pm 0.13

1458
 1459 Table 12: Ablation study of our DeepAFL’s training accuracy on the CIFAR-10. We evaluate four
 1460 distinct ablation models to explore the contributions of our key components in the representations:
 1461 (1) Ablating the residual skip connection, (2) Ablating the random projection \mathbf{B}_t , (3) Ablating the
 1462 activation function $\sigma(\cdot)$, and (4) Ablating the trainable transformation Ω_{t+1} . The value in parentheses
 1463 (\downarrow) indicates the performance drop compared to our full DeepAFL under identical conditions.

Layers	DeepAFL	Ablation (1)	Ablation (2)	Ablation (3)	Ablation (4)
$T = 0$	85.35 ± 0.02	85.35 ± 0.02 ($\downarrow 0.00$)			
$T = 5$	89.03 ± 0.03	85.23 ± 0.05 ($\downarrow 3.80$)	86.86 ± 0.03 ($\downarrow 2.17$)	85.35 ± 0.02 ($\downarrow 3.68$)	86.13 ± 0.03 ($\downarrow 2.90$)
$T = 10$	90.61 ± 0.05	85.28 ± 0.05 ($\downarrow 5.33$)	86.97 ± 0.03 ($\downarrow 3.64$)	85.35 ± 0.02 ($\downarrow 5.26$)	86.30 ± 0.04 ($\downarrow 4.31$)
$T = 15$	91.80 ± 0.04	85.32 ± 0.07 ($\downarrow 6.48$)	87.03 ± 0.04 ($\downarrow 4.77$)	85.35 ± 0.02 ($\downarrow 6.45$)	86.38 ± 0.04 ($\downarrow 5.42$)
$T = 20$	92.74 ± 0.03	85.33 ± 0.07 ($\downarrow 7.41$)	87.06 ± 0.03 ($\downarrow 5.68$)	85.35 ± 0.02 ($\downarrow 7.39$)	86.43 ± 0.04 ($\downarrow 6.31$)
$T = 25$	93.49 ± 0.02	85.34 ± 0.06 ($\downarrow 8.15$)	87.07 ± 0.03 ($\downarrow 6.42$)	85.35 ± 0.02 ($\downarrow 8.14$)	86.45 ± 0.03 ($\downarrow 7.04$)
$T = 30$	94.10 ± 0.03	85.37 ± 0.06 ($\downarrow 8.73$)	87.09 ± 0.03 ($\downarrow 7.01$)	85.35 ± 0.02 ($\downarrow 8.75$)	86.48 ± 0.02 ($\downarrow 7.62$)
$T = 35$	94.63 ± 0.03	85.39 ± 0.06 ($\downarrow 9.24$)	87.10 ± 0.03 ($\downarrow 7.53$)	85.36 ± 0.02 ($\downarrow 9.27$)	86.45 ± 0.04 ($\downarrow 8.18$)
$T = 40$	95.08 ± 0.03	85.41 ± 0.05 ($\downarrow 9.67$)	87.10 ± 0.03 ($\downarrow 7.98$)	85.35 ± 0.02 ($\downarrow 9.73$)	86.47 ± 0.01 ($\downarrow 8.61$)
$T = 45$	95.43 ± 0.03	85.40 ± 0.05 ($\downarrow 10.0$)	87.11 ± 0.03 ($\downarrow 8.32$)	85.35 ± 0.02 ($\downarrow 10.1$)	86.45 ± 0.02 ($\downarrow 8.98$)
$T = 50$	95.79 ± 0.03	85.42 ± 0.04 ($\downarrow 10.4$)	87.11 ± 0.02 ($\downarrow 8.68$)	85.35 ± 0.02 ($\downarrow 10.4$)	86.47 ± 0.04 ($\downarrow 9.32$)

1476
 1477 Table 13: Ablation study of our DeepAFL’s testing accuracy on the CIFAR-10. We evaluate four
 1478 distinct ablation models to explore the contributions of our key components in the representations:
 1479 (1) Ablating the residual skip connection, (2) Ablating the random projection \mathbf{B}_t , (3) Ablating the
 1480 activation function $\sigma(\cdot)$, and (4) Ablating the trainable transformation Ω_{t+1} . The value in parentheses
 1481 (\downarrow) indicates the performance drop compared to our full DeepAFL under identical conditions.

Layers	DeepAFL	Ablation (1)	Ablation (2)	Ablation (3)	Ablation (4)
$T = 0$	83.29 ± 0.20	83.29 ± 0.20 ($\downarrow 0.00$)			
$T = 5$	85.13 ± 0.08	83.03 ± 0.08 ($\downarrow 2.10$)	83.92 ± 0.21 ($\downarrow 1.21$)	83.33 ± 0.23 ($\downarrow 1.80$)	83.87 ± 0.16 ($\downarrow 1.26$)
$T = 10$	85.89 ± 0.09	83.06 ± 0.07 ($\downarrow 2.83$)	83.99 ± 0.21 ($\downarrow 1.90$)	83.34 ± 0.23 ($\downarrow 2.55$)	83.95 ± 0.14 ($\downarrow 1.94$)
$T = 15$	86.18 ± 0.08	83.05 ± 0.08 ($\downarrow 3.13$)	84.00 ± 0.24 ($\downarrow 2.18$)	83.33 ± 0.23 ($\downarrow 2.85$)	84.08 ± 0.28 ($\downarrow 2.10$)
$T = 20$	86.34 ± 0.06	83.03 ± 0.08 ($\downarrow 3.31$)	84.01 ± 0.25 ($\downarrow 2.33$)	83.32 ± 0.23 ($\downarrow 3.02$)	84.02 ± 0.22 ($\downarrow 2.32$)
$T = 25$	86.40 ± 0.06	82.99 ± 0.11 ($\downarrow 3.41$)	84.00 ± 0.29 ($\downarrow 2.40$)	83.33 ± 0.23 ($\downarrow 3.07$)	84.08 ± 0.27 ($\downarrow 2.32$)
$T = 30$	86.46 ± 0.06	83.02 ± 0.11 ($\downarrow 3.44$)	84.01 ± 0.27 ($\downarrow 2.45$)	83.33 ± 0.23 ($\downarrow 3.13$)	84.06 ± 0.34 ($\downarrow 2.40$)
$T = 35$	86.72 ± 0.06	82.94 ± 0.12 ($\downarrow 3.78$)	83.99 ± 0.27 ($\downarrow 2.73$)	83.33 ± 0.24 ($\downarrow 3.39$)	84.03 ± 0.33 ($\downarrow 2.69$)
$T = 40$	86.76 ± 0.06	82.93 ± 0.16 ($\downarrow 3.83$)	84.02 ± 0.25 ($\downarrow 2.74$)	83.33 ± 0.24 ($\downarrow 3.43$)	84.00 ± 0.28 ($\downarrow 2.76$)
$T = 45$	86.77 ± 0.06	82.96 ± 0.15 ($\downarrow 3.81$)	84.01 ± 0.28 ($\downarrow 2.76$)	83.34 ± 0.23 ($\downarrow 3.43$)	83.91 ± 0.31 ($\downarrow 2.86$)
$T = 50$	86.72 ± 0.06	82.97 ± 0.16 ($\downarrow 3.75$)	84.02 ± 0.28 ($\downarrow 2.70$)	83.34 ± 0.24 ($\downarrow 3.38$)	83.97 ± 0.25 ($\downarrow 2.75$)

1494
 1495 Table 14: Ablation study of our DeepAFL’s training accuracy on the CIFAR-100. We evaluate four
 1496 distinct ablation models to explore the contributions of our key components in the representations:
 1497 (1) Ablating the residual skip connection, (2) Ablating the random projection \mathbf{B}_t , (3) Ablating the
 1498 activation function $\sigma(\cdot)$, and (4) Ablating the trainable transformation Ω_{t+1} . The value in parentheses
 1499 (\downarrow) indicates the performance drop compared to our full DeepAFL under identical conditions.

Layers	DeepAFL	Ablation (1)	Ablation (2)	Ablation (3)	Ablation (4)
$T = 0$	65.66 ± 0.04	65.66 ± 0.04 ($\downarrow 0.00$)			
$T = 5$	74.93 ± 0.04	65.81 ± 0.04 ($\downarrow 9.12$)	69.80 ± 0.11 ($\downarrow 5.13$)	65.66 ± 0.04 ($\downarrow 9.27$)	65.67 ± 0.07 ($\downarrow 9.26$)
$T = 10$	79.41 ± 0.08	66.32 ± 0.03 ($\downarrow 13.1$)	70.16 ± 0.10 ($\downarrow 9.25$)	65.66 ± 0.04 ($\downarrow 13.8$)	65.70 ± 0.06 ($\downarrow 13.7$)
$T = 15$	82.69 ± 0.09	66.60 ± 0.11 ($\downarrow 16.1$)	70.38 ± 0.09 ($\downarrow 12.3$)	65.66 ± 0.03 ($\downarrow 17.0$)	65.72 ± 0.08 ($\downarrow 17.0$)
$T = 20$	85.15 ± 0.03	66.86 ± 0.08 ($\downarrow 18.3$)	70.46 ± 0.08 ($\downarrow 14.7$)	65.66 ± 0.03 ($\downarrow 19.5$)	65.72 ± 0.09 ($\downarrow 19.4$)
$T = 25$	87.21 ± 0.02	67.10 ± 0.12 ($\downarrow 20.1$)	70.57 ± 0.08 ($\downarrow 16.6$)	65.67 ± 0.04 ($\downarrow 21.5$)	65.74 ± 0.10 ($\downarrow 21.5$)
$T = 30$	88.84 ± 0.06	67.25 ± 0.09 ($\downarrow 21.6$)	70.63 ± 0.10 ($\downarrow 18.2$)	65.67 ± 0.03 ($\downarrow 23.2$)	65.77 ± 0.10 ($\downarrow 23.1$)
$T = 35$	90.14 ± 0.06	67.40 ± 0.11 ($\downarrow 22.7$)	70.65 ± 0.10 ($\downarrow 19.5$)	65.67 ± 0.04 ($\downarrow 24.5$)	65.76 ± 0.08 ($\downarrow 24.4$)
$T = 40$	91.29 ± 0.09	67.56 ± 0.12 ($\downarrow 23.7$)	70.69 ± 0.09 ($\downarrow 20.6$)	65.68 ± 0.04 ($\downarrow 25.6$)	65.78 ± 0.09 ($\downarrow 25.5$)
$T = 45$	92.15 ± 0.11	67.70 ± 0.08 ($\downarrow 24.5$)	70.73 ± 0.10 ($\downarrow 21.4$)	65.68 ± 0.04 ($\downarrow 26.5$)	65.77 ± 0.08 ($\downarrow 26.4$)
$T = 50$	92.93 ± 0.06	67.80 ± 0.11 ($\downarrow 25.1$)	70.78 ± 0.09 ($\downarrow 22.2$)	65.67 ± 0.04 ($\downarrow 27.3$)	65.76 ± 0.06 ($\downarrow 27.2$)

Table 15: Ablation study of our DeepAFL’s testing accuracy on the CIFAR-100. We evaluate four distinct ablation models to explore the contributions of our key components in the representations: (1) Ablating the residual skip connection, (2) Ablating the random projection \mathbf{B}_t , (3) Ablating the activation function $\sigma(\cdot)$, and (4) Ablating the trainable transformation Ω_{t+1} . The value in parentheses (\downarrow) indicates the performance drop compared to our full DeepAFL under identical conditions.

Layers	DeepAFL	Ablation (1)	Ablation (2)	Ablation (3)	Ablation (4)
$T = 0$	60.81 ± 0.15	60.81 ± 0.15 ($\downarrow 0.00$)			
$T = 5$	64.62 ± 0.02	60.91 ± 0.15 ($\downarrow 3.71$)	62.43 ± 0.09 ($\downarrow 2.19$)	60.80 ± 0.16 ($\downarrow 3.82$)	60.80 ± 0.10 ($\downarrow 3.82$)
$T = 10$	65.98 ± 0.13	61.17 ± 0.10 ($\downarrow 4.81$)	62.55 ± 0.06 ($\downarrow 3.43$)	60.81 ± 0.16 ($\downarrow 5.17$)	60.82 ± 0.05 ($\downarrow 5.16$)
$T = 15$	66.59 ± 0.03	61.48 ± 0.06 ($\downarrow 5.11$)	62.55 ± 0.08 ($\downarrow 4.04$)	60.82 ± 0.16 ($\downarrow 5.77$)	60.85 ± 0.05 ($\downarrow 5.74$)
$T = 20$	66.95 ± 0.04	61.82 ± 0.10 ($\downarrow 5.13$)	62.55 ± 0.09 ($\downarrow 4.40$)	60.82 ± 0.16 ($\downarrow 6.13$)	60.84 ± 0.12 ($\downarrow 6.11$)
$T = 25$	67.40 ± 0.10	61.95 ± 0.11 ($\downarrow 5.45$)	62.61 ± 0.05 ($\downarrow 4.79$)	60.82 ± 0.16 ($\downarrow 6.58$)	60.83 ± 0.07 ($\downarrow 6.57$)
$T = 30$	67.71 ± 0.06	62.12 ± 0.12 ($\downarrow 5.59$)	62.60 ± 0.09 ($\downarrow 5.11$)	60.82 ± 0.16 ($\downarrow 6.89$)	60.77 ± 0.04 ($\downarrow 6.94$)
$T = 35$	67.97 ± 0.12	62.26 ± 0.13 ($\downarrow 5.71$)	62.63 ± 0.11 ($\downarrow 5.34$)	60.83 ± 0.16 ($\downarrow 7.14$)	60.81 ± 0.06 ($\downarrow 7.16$)
$T = 40$	68.19 ± 0.07	62.28 ± 0.06 ($\downarrow 5.91$)	62.64 ± 0.08 ($\downarrow 5.55$)	60.82 ± 0.15 ($\downarrow 7.37$)	60.88 ± 0.05 ($\downarrow 7.31$)
$T = 45$	68.32 ± 0.11	62.37 ± 0.12 ($\downarrow 5.95$)	62.63 ± 0.12 ($\downarrow 5.69$)	60.83 ± 0.16 ($\downarrow 7.49$)	60.89 ± 0.11 ($\downarrow 7.43$)
$T = 50$	68.51 ± 0.09	62.39 ± 0.13 ($\downarrow 6.12$)	62.64 ± 0.11 ($\downarrow 5.87$)	60.84 ± 0.16 ($\downarrow 7.67$)	60.89 ± 0.06 ($\downarrow 7.62$)

Table 16: Ablation study of our DeepAFL’s training accuracy on the Tiny-ImageNet. We evaluate four distinct ablation models to explore the contributions of our key components in the representations: (1) Ablating the residual skip connection, (2) Ablating the random projection \mathbf{B}_t , (3) Ablating the activation function $\sigma(\cdot)$, and (4) Ablating the trainable transformation Ω_{t+1} . The value in parentheses (\downarrow) indicates the performance drop compared to our full DeepAFL under identical conditions.

Layers	DeepAFL	Ablation (1)	Ablation (2)	Ablation (3)	Ablation (4)
$T = 0$	60.30 ± 0.02	60.30 ± 0.02 ($\downarrow 0.00$)			
$T = 5$	66.89 ± 0.06	60.84 ± 0.03 ($\downarrow 6.05$)	63.21 ± 0.03 ($\downarrow 3.68$)	60.30 ± 0.01 ($\downarrow 6.59$)	60.31 ± 0.02 ($\downarrow 6.58$)
$T = 10$	70.14 ± 0.03	61.87 ± 0.05 ($\downarrow 8.27$)	63.46 ± 0.03 ($\downarrow 6.68$)	60.30 ± 0.01 ($\downarrow 9.84$)	60.31 ± 0.01 ($\downarrow 9.83$)
$T = 15$	72.64 ± 0.05	62.53 ± 0.04 ($\downarrow 10.1$)	63.59 ± 0.03 ($\downarrow 9.05$)	60.31 ± 0.01 ($\downarrow 12.3$)	60.32 ± 0.02 ($\downarrow 12.3$)
$T = 20$	74.75 ± 0.02	62.97 ± 0.03 ($\downarrow 11.8$)	63.68 ± 0.03 ($\downarrow 11.1$)	60.31 ± 0.01 ($\downarrow 14.4$)	60.33 ± 0.02 ($\downarrow 14.4$)
$T = 25$	76.58 ± 0.03	63.30 ± 0.02 ($\downarrow 13.3$)	63.74 ± 0.03 ($\downarrow 12.8$)	60.31 ± 0.01 ($\downarrow 16.3$)	60.33 ± 0.03 ($\downarrow 16.3$)
$T = 30$	78.31 ± 0.02	63.58 ± 0.04 ($\downarrow 14.7$)	63.77 ± 0.03 ($\downarrow 14.5$)	60.31 ± 0.01 ($\downarrow 18.0$)	60.35 ± 0.02 ($\downarrow 18.0$)
$T = 35$	79.87 ± 0.03	63.76 ± 0.02 ($\downarrow 16.1$)	63.80 ± 0.03 ($\downarrow 16.1$)	60.31 ± 0.01 ($\downarrow 19.6$)	60.36 ± 0.02 ($\downarrow 19.5$)
$T = 40$	81.32 ± 0.11	63.85 ± 0.02 ($\downarrow 17.5$)	63.82 ± 0.02 ($\downarrow 17.5$)	60.31 ± 0.01 ($\downarrow 21.0$)	60.35 ± 0.02 ($\downarrow 21.0$)
$T = 45$	82.62 ± 0.04	63.98 ± 0.03 ($\downarrow 18.6$)	63.84 ± 0.02 ($\downarrow 18.8$)	60.31 ± 0.01 ($\downarrow 22.3$)	60.35 ± 0.02 ($\downarrow 22.3$)
$T = 50$	83.82 ± 0.09	64.07 ± 0.01 ($\downarrow 19.8$)	63.86 ± 0.01 ($\downarrow 20.0$)	60.31 ± 0.01 ($\downarrow 23.5$)	60.36 ± 0.01 ($\downarrow 23.5$)

Table 17: Ablation study of our DeepAFL’s testing accuracy on the Tiny-ImageNet. We evaluate four distinct ablation models to explore the contributions of our key components in the representations: (1) Ablating the residual skip connection, (2) Ablating the random projection \mathbf{B}_t , (3) Ablating the activation function $\sigma(\cdot)$, and (4) Ablating the trainable transformation Ω_{t+1} . The value in parentheses (\downarrow) indicates the performance drop compared to our full DeepAFL under identical conditions.

Layers	DeepAFL	Ablation (1)	Ablation (2)	Ablation (3)	Ablation (4)
$T = 0$	56.72 ± 0.14	56.72 ± 0.14 ($\downarrow 0.00$)			
$T = 5$	60.24 ± 0.02	57.39 ± 0.11 ($\downarrow 2.85$)	58.16 ± 0.24 ($\downarrow 2.08$)	56.69 ± 0.13 ($\downarrow 3.55$)	56.69 ± 0.11 ($\downarrow 3.55$)
$T = 10$	61.34 ± 0.13	58.42 ± 0.16 ($\downarrow 2.92$)	58.22 ± 0.17 ($\downarrow 3.12$)	56.70 ± 0.13 ($\downarrow 4.64$)	56.69 ± 0.14 ($\downarrow 4.65$)
$T = 15$	61.94 ± 0.08	59.01 ± 0.10 ($\downarrow 2.93$)	58.20 ± 0.14 ($\downarrow 3.74$)	56.70 ± 0.13 ($\downarrow 5.24$)	56.69 ± 0.12 ($\downarrow 5.25$)
$T = 20$	62.34 ± 0.02	59.21 ± 0.08 ($\downarrow 3.13$)	58.25 ± 0.19 ($\downarrow 4.09$)	56.70 ± 0.13 ($\downarrow 5.64$)	56.68 ± 0.12 ($\downarrow 5.66$)
$T = 25$	62.66 ± 0.05	59.37 ± 0.14 ($\downarrow 3.29$)	58.25 ± 0.16 ($\downarrow 4.41$)	56.70 ± 0.14 ($\downarrow 5.96$)	56.70 ± 0.10 ($\downarrow 5.96$)
$T = 30$	62.94 ± 0.03	59.49 ± 0.10 ($\downarrow 3.45$)	58.24 ± 0.17 ($\downarrow 4.70$)	56.70 ± 0.14 ($\downarrow 6.24$)	56.70 ± 0.12 ($\downarrow 6.24$)
$T = 35$	63.24 ± 0.16	59.54 ± 0.07 ($\downarrow 3.70$)	58.27 ± 0.18 ($\downarrow 4.97$)	56.70 ± 0.14 ($\downarrow 6.54$)	56.72 ± 0.12 ($\downarrow 6.52$)
$T = 40$	63.48 ± 0.18	59.66 ± 0.08 ($\downarrow 3.82$)	58.34 ± 0.16 ($\downarrow 5.14$)	56.70 ± 0.14 ($\downarrow 6.78$)	56.73 ± 0.12 ($\downarrow 6.75$)
$T = 45$	63.65 ± 0.18	59.72 ± 0.10 ($\downarrow 3.93$)	58.35 ± 0.18 ($\downarrow 5.30$)	56.70 ± 0.14 ($\downarrow 6.95$)	56.72 ± 0.10 ($\downarrow 6.93$)
$T = 50$	63.74 ± 0.04	59.81 ± 0.11 ($\downarrow 3.93$)	58.39 ± 0.15 ($\downarrow 5.35$)	56.70 ± 0.14 ($\downarrow 7.04$)	56.72 ± 0.08 ($\downarrow 7.02$)

1566 Table 18: Training accuracy comparison of our DeepAFL against the other four deepening strategies
 1567 for the analytic networks on the CIFAR-10, including (1) cascades w/ random feature, (2) cascades
 1568 w/ random feature & label encoding, (3) cascades w/ activated random feature, (4) cascades w/
 1569 activated random feature & label encoding. All the strategies share the same zero-layer features to
 1570 ensure fairness. The value in parentheses (↓) indicates the performance lags behind our DeepAFL.

Layers	DeepAFL	Strategy (1)	Strategy (2)	Strategy (3)	Strategy (4)
$T = 0$	85.35 ± 0.02	85.35 ± 0.02 (↓ 0.00)			
$T = 5$	89.03 ± 0.03	85.35 ± 0.02 (↓ 3.68)	85.36 ± 0.03 (↓ 3.67)	85.68 ± 0.01 (↓ 3.35)	85.68 ± 0.02 (↓ 3.35)
$T = 10$	90.61 ± 0.05	85.35 ± 0.03 (↓ 5.26)	85.36 ± 0.03 (↓ 5.25)	85.68 ± 0.01 (↓ 4.93)	85.67 ± 0.02 (↓ 4.94)
$T = 15$	91.80 ± 0.04	85.35 ± 0.03 (↓ 6.45)	85.35 ± 0.03 (↓ 6.45)	85.70 ± 0.00 (↓ 6.10)	85.68 ± 0.03 (↓ 6.12)
$T = 20$	92.74 ± 0.03	85.35 ± 0.03 (↓ 7.39)	85.35 ± 0.03 (↓ 7.39)	85.70 ± 0.01 (↓ 7.04)	85.67 ± 0.01 (↓ 7.07)
$T = 25$	93.49 ± 0.02	85.35 ± 0.03 (↓ 8.14)	85.35 ± 0.03 (↓ 8.14)	85.71 ± 0.01 (↓ 7.78)	85.68 ± 0.02 (↓ 7.81)
$T = 30$	94.10 ± 0.03	85.36 ± 0.03 (↓ 8.74)	85.35 ± 0.03 (↓ 8.75)	85.71 ± 0.01 (↓ 8.39)	85.67 ± 0.03 (↓ 8.43)
$T = 35$	94.63 ± 0.03	85.36 ± 0.03 (↓ 9.27)	85.35 ± 0.03 (↓ 9.28)	85.71 ± 0.01 (↓ 8.92)	85.68 ± 0.03 (↓ 8.95)
$T = 40$	95.08 ± 0.03	85.35 ± 0.02 (↓ 9.73)	85.35 ± 0.03 (↓ 9.73)	85.71 ± 0.01 (↓ 9.37)	85.68 ± 0.03 (↓ 9.40)
$T = 45$	95.43 ± 0.03	85.35 ± 0.03 (↓ 10.1)	85.35 ± 0.03 (↓ 10.1)	85.71 ± 0.00 (↓ 9.72)	85.69 ± 0.03 (↓ 9.74)
$T = 50$	95.79 ± 0.03	85.36 ± 0.03 (↓ 10.4)	85.35 ± 0.03 (↓ 10.4)	85.70 ± 0.01 (↓ 10.1)	85.68 ± 0.03 (↓ 10.1)

1584 Table 19: Testing accuracy comparison of our DeepAFL against the other four deepening strategies
 1585 for the analytic networks on the CIFAR-10, including (1) cascades w/ random feature, (2) cascades
 1586 w/ random feature & label encoding, (3) cascades w/ activated random feature, (4) cascades w/
 1587 activated random feature & label encoding. All the strategies share the same zero-layer features to
 1588 ensure fairness. The value in parentheses (↓) indicates the performance lags behind our DeepAFL.

Layers	DeepAFL	Strategy (1)	Strategy (2)	Strategy (3)	Strategy (4)
$T = 0$	83.29 ± 0.20	83.30 ± 0.20 (↓ 0.00)			
$T = 5$	85.13 ± 0.08	83.30 ± 0.22 (↓ 1.83)	83.31 ± 0.24 (↓ 1.82)	83.62 ± 0.20 (↓ 1.51)	83.61 ± 0.20 (↓ 1.52)
$T = 10$	85.89 ± 0.09	83.29 ± 0.23 (↓ 2.60)	83.31 ± 0.23 (↓ 2.58)	83.63 ± 0.23 (↓ 2.26)	83.60 ± 0.20 (↓ 2.29)
$T = 15$	86.18 ± 0.08	83.32 ± 0.23 (↓ 2.86)	83.31 ± 0.23 (↓ 2.87)	83.62 ± 0.20 (↓ 2.56)	83.60 ± 0.19 (↓ 2.58)
$T = 20$	86.34 ± 0.06	83.29 ± 0.24 (↓ 3.05)	83.31 ± 0.24 (↓ 3.03)	83.62 ± 0.22 (↓ 2.72)	83.62 ± 0.23 (↓ 2.72)
$T = 25$	86.40 ± 0.06	83.29 ± 0.24 (↓ 3.11)	83.31 ± 0.23 (↓ 3.09)	83.63 ± 0.22 (↓ 2.77)	83.63 ± 0.21 (↓ 2.77)
$T = 30$	86.46 ± 0.06	83.31 ± 0.24 (↓ 3.15)	83.29 ± 0.23 (↓ 3.17)	83.64 ± 0.22 (↓ 2.82)	83.60 ± 0.22 (↓ 2.86)
$T = 35$	86.72 ± 0.06	83.30 ± 0.23 (↓ 3.42)	83.31 ± 0.24 (↓ 3.41)	83.64 ± 0.22 (↓ 3.08)	83.59 ± 0.20 (↓ 3.13)
$T = 40$	86.76 ± 0.06	83.31 ± 0.25 (↓ 3.45)	83.32 ± 0.24 (↓ 3.44)	83.66 ± 0.20 (↓ 3.10)	83.62 ± 0.20 (↓ 3.14)
$T = 45$	86.77 ± 0.06	83.31 ± 0.24 (↓ 3.46)	83.30 ± 0.24 (↓ 3.47)	83.65 ± 0.23 (↓ 3.12)	83.60 ± 0.22 (↓ 3.17)
$T = 50$	86.72 ± 0.06	83.31 ± 0.23 (↓ 3.41)	83.30 ± 0.24 (↓ 3.42)	83.64 ± 0.22 (↓ 3.08)	83.60 ± 0.22 (↓ 3.12)

1602 Table 20: Training accuracy comparison of our DeepAFL against the other four deepening strategies
 1603 for the analytic networks on the CIFAR-100, including (1) cascades w/ random feature, (2) cascades
 1604 w/ random feature & label encoding, (3) cascades w/ activated random feature, (4) cascades w/
 1605 activated random feature & label encoding. All the strategies share the same zero-layer features to
 1606 ensure fairness. The value in parentheses (↓) indicates the performance lags behind our DeepAFL.

Layers	DeepAFL	Strategy (1)	Strategy (2)	Strategy (3)	Strategy (4)
$T = 0$	65.66 ± 0.04	65.66 ± 0.04 (↓ 0.00)			
$T = 5$	74.93 ± 0.04	65.64 ± 0.06 (↓ 9.29)	65.58 ± 0.07 (↓ 9.35)	66.18 ± 0.08 (↓ 8.75)	65.30 ± 0.05 (↓ 9.63)
$T = 10$	79.41 ± 0.08	65.64 ± 0.05 (↓ 13.8)	65.58 ± 0.06 (↓ 13.8)	66.67 ± 0.08 (↓ 12.7)	65.39 ± 0.06 (↓ 14.0)
$T = 15$	82.69 ± 0.09	65.64 ± 0.05 (↓ 17.0)	65.58 ± 0.06 (↓ 17.1)	67.11 ± 0.11 (↓ 15.6)	65.51 ± 0.04 (↓ 17.2)
$T = 20$	85.15 ± 0.03	65.64 ± 0.06 (↓ 19.5)	65.58 ± 0.07 (↓ 19.6)	67.52 ± 0.11 (↓ 17.6)	65.60 ± 0.07 (↓ 19.6)
$T = 25$	87.21 ± 0.02	65.64 ± 0.05 (↓ 21.6)	65.57 ± 0.07 (↓ 21.6)	67.86 ± 0.13 (↓ 19.4)	65.70 ± 0.12 (↓ 21.5)
$T = 30$	88.84 ± 0.06	65.64 ± 0.05 (↓ 23.2)	65.56 ± 0.07 (↓ 23.3)	68.14 ± 0.11 (↓ 20.7)	65.74 ± 0.12 (↓ 23.1)
$T = 35$	90.14 ± 0.06	65.64 ± 0.05 (↓ 24.5)	65.58 ± 0.06 (↓ 24.6)	68.40 ± 0.12 (↓ 21.7)	65.84 ± 0.09 (↓ 24.3)
$T = 40$	91.29 ± 0.09	65.64 ± 0.05 (↓ 25.6)	65.58 ± 0.05 (↓ 25.7)	68.64 ± 0.14 (↓ 22.6)	65.95 ± 0.15 (↓ 25.3)
$T = 45$	92.15 ± 0.11	65.64 ± 0.05 (↓ 26.5)	65.58 ± 0.07 (↓ 26.6)	68.82 ± 0.16 (↓ 23.3)	66.04 ± 0.17 (↓ 26.1)
$T = 50$	92.93 ± 0.06	65.64 ± 0.05 (↓ 27.3)	65.59 ± 0.07 (↓ 27.3)	68.99 ± 0.19 (↓ 23.9)	66.16 ± 0.13 (↓ 26.8)

1620
 1621 Table 21: Testing accuracy comparison of our DeepAFL against the other four deepening strategies
 1622 for the analytic networks on the CIFAR-100, including (1) cascades w/ random feature, (2) cascades
 1623 w/ random feature & label encoding, (3) cascades w/ activated random feature, (4) cascades w/
 1624 activated random feature & label encoding. All the strategies share the same zero-layer features to
 1625 ensure fairness. The value in parentheses (↓) indicates the performance lags behind our DeepAFL.

Layers	DeepAFL	Strategy (1)	Strategy (2)	Strategy (3)	Strategy (4)
$T = 0$	60.81 ± 0.15	60.81 ± 0.15 (↓ 0.00)			
$T = 5$	64.62 ± 0.02	60.80 ± 0.13 (↓ 3.82)	60.75 ± 0.14 (↓ 3.87)	61.23 ± 0.07 (↓ 3.39)	60.57 ± 0.04 (↓ 4.05)
$T = 10$	65.98 ± 0.13	60.78 ± 0.10 (↓ 5.20)	60.74 ± 0.11 (↓ 5.24)	61.67 ± 0.09 (↓ 4.31)	60.69 ± 0.05 (↓ 5.29)
$T = 15$	66.59 ± 0.03	60.78 ± 0.12 (↓ 5.81)	60.74 ± 0.11 (↓ 5.85)	61.98 ± 0.03 (↓ 4.61)	60.85 ± 0.02 (↓ 5.74)
$T = 20$	66.95 ± 0.04	60.79 ± 0.10 (↓ 6.16)	60.73 ± 0.11 (↓ 6.22)	62.21 ± 0.02 (↓ 4.74)	60.95 ± 0.07 (↓ 6.00)
$T = 25$	67.40 ± 0.10	60.78 ± 0.10 (↓ 6.62)	60.72 ± 0.11 (↓ 6.68)	62.31 ± 0.04 (↓ 5.09)	60.99 ± 0.06 (↓ 6.41)
$T = 30$	67.71 ± 0.06	60.78 ± 0.12 (↓ 6.93)	60.73 ± 0.11 (↓ 6.98)	62.53 ± 0.08 (↓ 5.18)	61.12 ± 0.08 (↓ 6.59)
$T = 35$	67.97 ± 0.12	60.78 ± 0.12 (↓ 7.19)	60.72 ± 0.10 (↓ 7.25)	62.73 ± 0.12 (↓ 5.24)	61.16 ± 0.06 (↓ 6.81)
$T = 40$	68.19 ± 0.07	60.78 ± 0.12 (↓ 7.41)	60.74 ± 0.10 (↓ 7.45)	62.82 ± 0.11 (↓ 5.37)	61.20 ± 0.01 (↓ 6.99)
$T = 45$	68.32 ± 0.11	60.78 ± 0.11 (↓ 7.54)	60.73 ± 0.11 (↓ 7.59)	62.96 ± 0.11 (↓ 5.36)	61.27 ± 0.01 (↓ 7.05)
$T = 50$	68.51 ± 0.09	60.77 ± 0.12 (↓ 7.74)	60.71 ± 0.10 (↓ 7.80)	63.05 ± 0.11 (↓ 5.46)	61.37 ± 0.06 (↓ 7.14)

1638
 1639 Table 22: Training accuracy comparison of our DeepAFL against the other four deepening strategies
 1640 for the analytic networks on the Tiny-ImageNet, including (1) cascades w/ random feature, (2)
 1641 cascades w/ random feature & label encoding, (3) cascades w/ activated random feature, (4) cascades
 1642 w/ activated random feature & label encoding. All the strategies share the same zero-layer features to
 1643 ensure fairness. The value in parentheses (↓) indicates the performance lags behind our DeepAFL.

Layers	DeepAFL	Strategy (1)	Strategy (2)	Strategy (3)	Strategy (4)
$T = 0$	60.30 ± 0.02	60.30 ± 0.02 (↓ 0.00)			
$T = 5$	66.89 ± 0.06	60.20 ± 0.02 (↓ 6.69)	59.84 ± 0.02 (↓ 7.05)	59.92 ± 0.03 (↓ 6.97)	56.14 ± 0.08 (↓ 10.8)
$T = 10$	70.14 ± 0.03	60.20 ± 0.01 (↓ 9.94)	59.83 ± 0.01 (↓ 10.3)	60.00 ± 0.02 (↓ 10.1)	54.78 ± 0.07 (↓ 15.4)
$T = 15$	72.64 ± 0.05	60.21 ± 0.01 (↓ 12.4)	59.82 ± 0.04 (↓ 12.8)	60.08 ± 0.03 (↓ 12.6)	54.40 ± 0.07 (↓ 18.2)
$T = 20$	74.75 ± 0.02	60.20 ± 0.01 (↓ 14.5)	59.81 ± 0.01 (↓ 14.9)	60.17 ± 0.04 (↓ 14.6)	54.23 ± 0.11 (↓ 20.5)
$T = 25$	76.58 ± 0.03	60.19 ± 0.01 (↓ 16.4)	59.80 ± 0.02 (↓ 16.8)	60.25 ± 0.04 (↓ 16.3)	54.12 ± 0.11 (↓ 22.5)
$T = 30$	78.31 ± 0.02	60.21 ± 0.01 (↓ 18.1)	59.81 ± 0.04 (↓ 18.5)	60.33 ± 0.04 (↓ 18.0)	54.03 ± 0.12 (↓ 24.3)
$T = 35$	79.87 ± 0.03	60.21 ± 0.01 (↓ 19.7)	59.81 ± 0.02 (↓ 20.1)	60.44 ± 0.07 (↓ 19.4)	54.06 ± 0.12 (↓ 25.8)
$T = 40$	81.32 ± 0.11	60.20 ± 0.01 (↓ 21.1)	59.80 ± 0.03 (↓ 21.5)	60.52 ± 0.05 (↓ 20.8)	53.99 ± 0.12 (↓ 27.3)
$T = 45$	82.62 ± 0.04	60.20 ± 0.01 (↓ 22.4)	59.80 ± 0.03 (↓ 22.8)	60.62 ± 0.06 (↓ 22.0)	54.07 ± 0.08 (↓ 28.6)
$T = 50$	83.82 ± 0.09	60.20 ± 0.02 (↓ 23.6)	59.78 ± 0.03 (↓ 24.0)	60.71 ± 0.05 (↓ 23.1)	54.05 ± 0.11 (↓ 29.8)

1656
 1657 Table 23: Testing accuracy comparison of our DeepAFL against the other four deepening strategies
 1658 for the analytic networks on the Tiny-ImageNet, including (1) cascades w/ random feature, (2)
 1659 cascades w/ random feature & label encoding, (3) cascades w/ activated random feature, (4) cascades
 1660 w/ activated random feature & label encoding. All the strategies share the same zero-layer features to
 1661 ensure fairness. The value in parentheses (↓) indicates the performance lags behind our DeepAFL.

Layers	DeepAFL	Strategy (1)	Strategy (2)	Strategy (3)	Strategy (4)
$T = 0$	56.72 ± 0.14	56.72 ± 0.14 (↓ 0.00)			
$T = 5$	60.24 ± 0.02	56.69 ± 0.12 (↓ 3.55)	56.32 ± 0.10 (↓ 3.92)	56.31 ± 0.10 (↓ 3.93)	53.26 ± 0.12 (↓ 6.98)
$T = 10$	61.34 ± 0.13	56.61 ± 0.13 (↓ 4.73)	56.26 ± 0.11 (↓ 5.08)	56.40 ± 0.06 (↓ 4.94)	52.01 ± 0.06 (↓ 9.33)
$T = 15$	61.94 ± 0.08	56.62 ± 0.14 (↓ 5.32)	56.25 ± 0.13 (↓ 5.69)	56.49 ± 0.10 (↓ 5.45)	51.74 ± 0.12 (↓ 10.2)
$T = 20$	62.34 ± 0.02	56.64 ± 0.14 (↓ 5.70)	56.23 ± 0.08 (↓ 6.11)	56.61 ± 0.10 (↓ 5.73)	51.63 ± 0.09 (↓ 10.7)
$T = 25$	62.66 ± 0.05	56.62 ± 0.14 (↓ 6.04)	56.27 ± 0.09 (↓ 6.39)	56.66 ± 0.12 (↓ 6.00)	51.66 ± 0.16 (↓ 11.0)
$T = 30$	62.94 ± 0.03	56.63 ± 0.09 (↓ 6.31)	56.23 ± 0.09 (↓ 6.71)	56.75 ± 0.21 (↓ 6.19)	51.47 ± 0.16 (↓ 11.5)
$T = 35$	63.24 ± 0.16	56.61 ± 0.14 (↓ 6.63)	56.22 ± 0.04 (↓ 7.02)	56.79 ± 0.15 (↓ 6.45)	51.51 ± 0.21 (↓ 11.7)
$T = 40$	63.48 ± 0.18	56.62 ± 0.12 (↓ 6.86)	56.26 ± 0.02 (↓ 7.22)	56.83 ± 0.24 (↓ 6.65)	51.47 ± 0.15 (↓ 12.0)
$T = 45$	63.65 ± 0.18	56.61 ± 0.13 (↓ 7.04)	56.22 ± 0.09 (↓ 7.43)	56.99 ± 0.20 (↓ 6.66)	51.42 ± 0.09 (↓ 12.2)
$T = 50$	63.74 ± 0.04	56.61 ± 0.14 (↓ 7.13)	56.23 ± 0.13 (↓ 7.51)	57.11 ± 0.18 (↓ 6.63)	51.47 ± 0.12 (↓ 12.3)

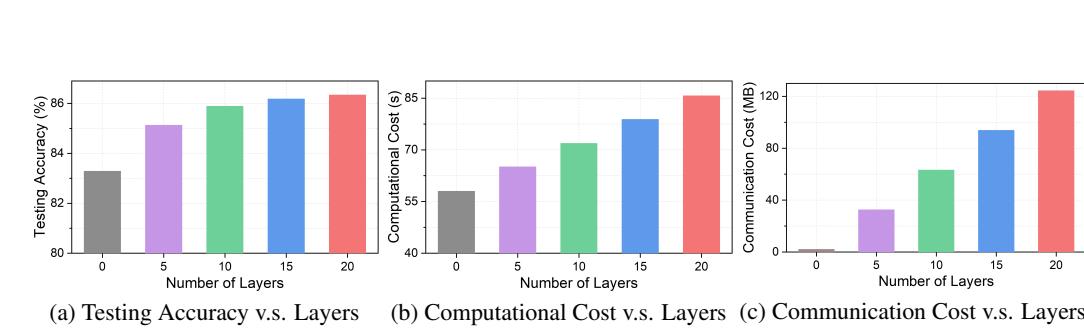


Figure 4: Accuracy-Efficiency balance of our DeepAFL on the CIFAR-10.

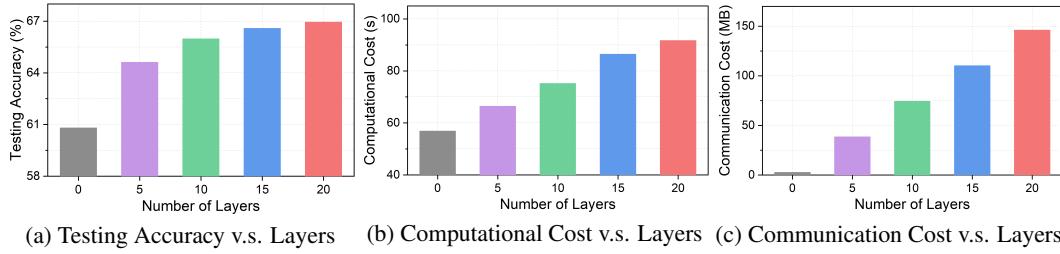


Figure 5: Accuracy-Efficiency balance of our DeepAFL on the CIFAR-100.

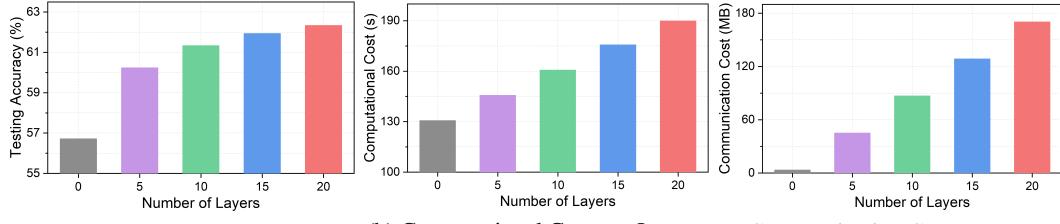


Figure 6: Accuracy-Efficiency balance of our DeepAFL on the Tiny-ImageNet.

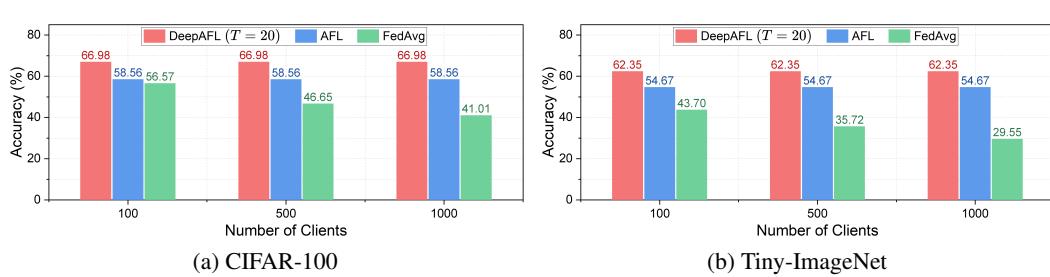


Figure 7: Scalability analyses of our DeepAFL against baselines with different numbers of clients.

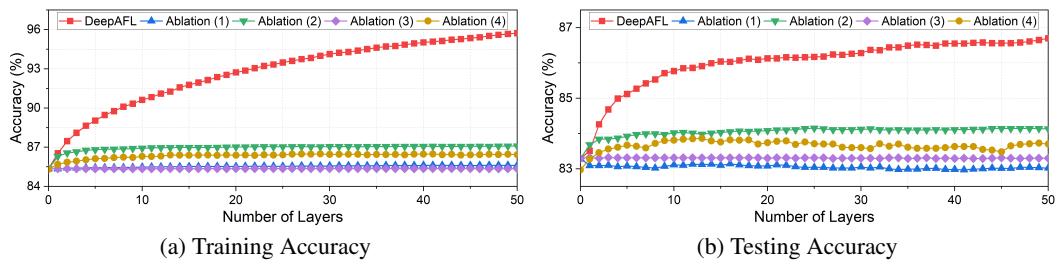


Figure 8: Ablation study of our DeepAFL on the CIFAR-10.

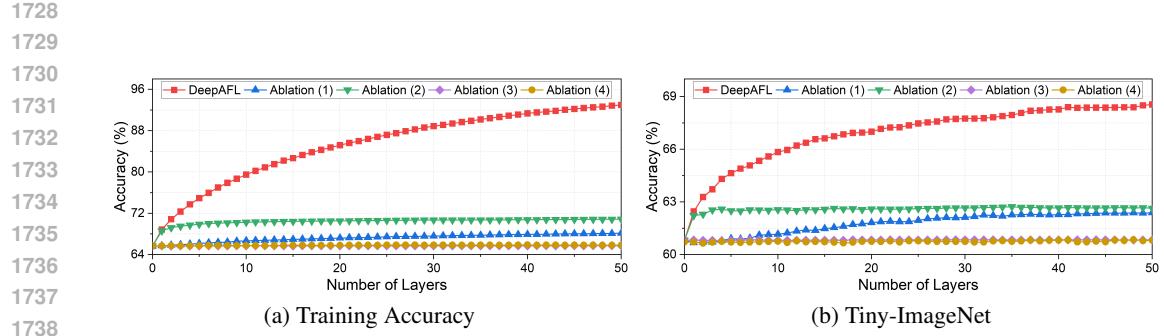


Figure 9: Ablation study of our DeepAFL on the CIFAR-100.

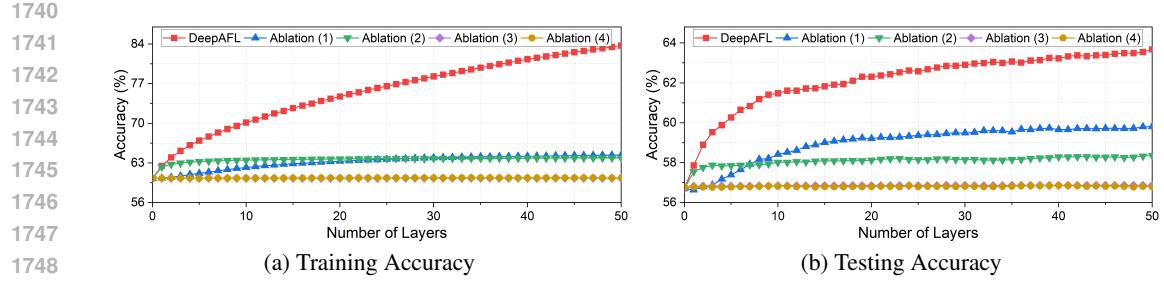


Figure 10: Ablation study of our DeepAFL on the Tiny-ImageNet.

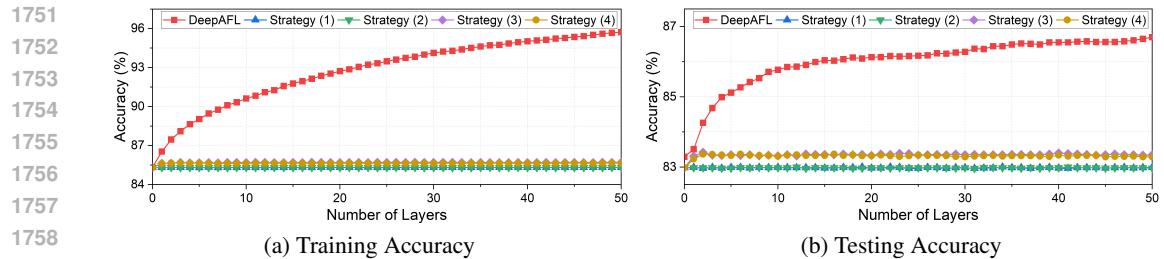


Figure 11: Comparison of our DeepAFL with other deepening strategies on the CIFAR-10.

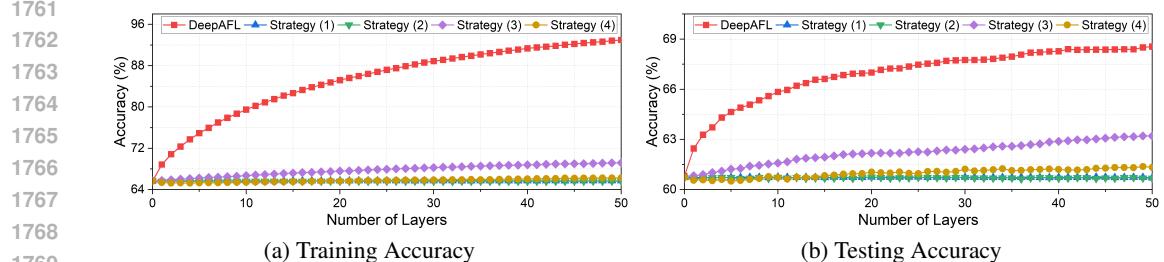


Figure 12: Comparison of our DeepAFL with other deepening strategies on the CIFAR-100.

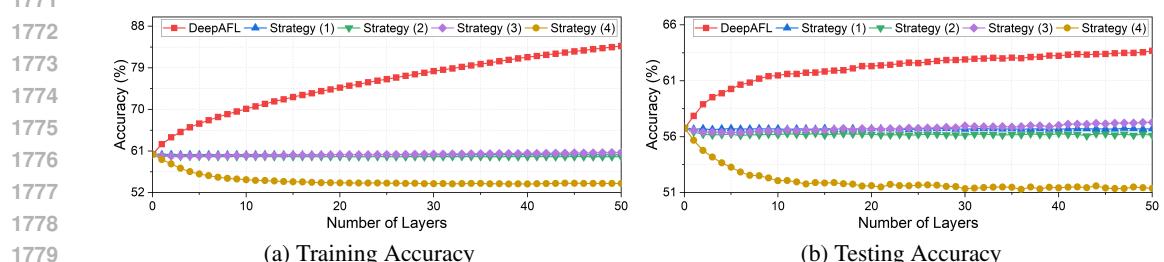


Figure 13: Comparison of our DeepAFL with other deepening strategies on the Tiny-ImageNet.

1782 E.2 IMPLEMENTATION DETAILS OF THE EXPERIMENTS
1783

1784 Here, we detail the implementation details of our experiments. Specifically, the number of clients is
1785 set to 100 for all methods. To simulate diverse data heterogeneity scenarios, two common non-IID
1786 partitioning settings are employed: LDA (Lin et al., 2020) (as Non-IID-1) and Sharding (Lin et al.,
1787 2020) (as Non-IID-2). Under the Non-IID-1 setting, the dataset is allocated to all clients via the
1788 Dirichlet distribution, with the parameter α modulating the heterogeneity level. Under the Non-IID-
1789 2 setting, the dataset is sorted by label and divided into equal-sized shards for distribution among
1790 clients, where the number of shards per client s controls the heterogeneity level. Smaller values of
1791 α and s both indicate more heterogeneous data distributions. We set $\alpha \in \{0.1, 0.05\}$ and $s \in \{2, 4\}$
1792 for CIFAR-10, while setting $\alpha \in \{0.1, 0.01\}$ and $s \in \{5, 10\}$ for CIFAR-100 and Tiny-ImageNet.
1793

1794 To ensure a fair comparison, in our main experiments, the results of all baselines are directly taken
1795 from the provided data in AFL (He et al., 2025b). Furthermore, for all three datasets (i.e., CIFAR-10,
1796 CIFAR-100, and Tiny-ImageNet), we adopted exactly the same pre-processing procedure as AFL,
1797 resizing all input images to 224×224 . Regarding the specific parameters of our DeepAFL, we
1798 adopt the following *default settings*. First, we employ GELU as the activation function and set the
1799 projection dimensions to $d_\Phi = d_F = 1024$ across all three datasets. Additionally, the regularization
1800 parameter λ is set to 10 for CIFAR-10, and 1 for both CIFAR-100 and Tiny-ImageNet. Then, for
1801 the regularization parameter γ , we set it to 0.1 for CIFAR-10, and 0.01 for both CIFAR-100 and
1802 Tiny-ImageNet. To ensure experimental transparency, we provide the detailed parameter settings
1803 for each experiment in the extended analyses of our evaluation in Section E.3.
1804

1805 As for the metrics employed in our experiments, we adopt the top-1 accuracy on the testing sets as
1806 the primary metric for evaluating the performance of all approaches. In addition, since the top-1
1807 accuracy on the training sets also serves as a strong indicator of the representation learning capabil-
1808 ity, we extensively employ this metric to analyze how the representations of different models evolve
1809 with increasing network depth. Moreover, to assess the efficiency of all approaches, we measure the
1810 required training time (s) and the volume of transmitted data (MB) as indicators of computational
1811 and communication cost, respectively. To analyze the marginal effect of increasing network depth,
1812 we use the symbol Δ to denote the difference between two consecutive cases of the the metric.
1813

1814 All the experiments in our paper are executed three times per setting, and we report the mean and
1815 standard error of the experimental results. Furthermore, all the experimental evaluations in this
1816 paper are conducted using PyTorch on NVIDIA RTX 4090 GPUs. Notably, for transparency, the
1817 related codes will be made publicly accessible as open-sourced upon the acceptance of this paper,
1818 allowing the broader research community to verify our findings and build upon our work.
1819

1820 E.3 DETAILED ANALYSES ON EXPERIMENTAL EVALUATIONS
18211822 E.3.1 MAIN COMPARISONS
1823

1824 In our main results reported in Tables 1 and 2, we employ GELU as the activation function and set the
1825 projection dimensions to $d_\Phi = d_F = 1024$ across all three datasets. The regularization parameter
1826 λ is fixed at 10 for CIFAR-10 and at 1 for both CIFAR-100 and Tiny-ImageNet. Similarly, the
1827 regularization parameter γ is set to 0.1 for CIFAR-10 and 0.01 for CIFAR-100 and Tiny-ImageNet.
1828 To ensure a fair comparison, the results of all baselines in Tables 1 and 2 are directly taken from
1829 the benchmark data provided in AFL (He et al., 2025b). It is worth noting that, since our DeepAFL
1830 involves random projections, its performance may vary under different random seeds. To mitigate
1831 the influence of randomness, all main experiments in this paper were repeated three times, and the
1832 results are reported as Mean \pm Standard Error. The improvements of our DeepAFL were validated
1833 by Chi-squared tests, all of which were found to be statistically significant at the $p = 0.05$ level.
1834

1835 From the results, we observe that the performance of DeepAFL is consistent across different levels
1836 of heterogeneity, which empirically supports the heterogeneity invariance in Theorem 1. Moreover,
1837 as the network depth T increases, DeepAFL consistently achieves significant performance gains,
1838 empirically supporting its representation learning capability in Theorems 2–3. It is also worth noting
1839 that although we report the performance of DeepAFL under $T \in \{5, 10, 20\}$ in the main results,
1840 its performance continues to improve as T increases. This is further confirmed by our extended
1841 experiments with $T \in [0, 50]$. Specifically, on the more complex CIFAR-100 and Tiny-ImageNet
1842 datasets, the testing accuracy of DeepAFL at $T = 50$ surpasses that at $T = 20$ by more than 1.5%.

1836 E.3.2 INVARIANCE ANALYSES
1837

1838 Here, we conduct detailed analyses to demonstrate DeepAFL’s ideal property of invariance to data
1839 heterogeneity. As shown in Tables 1 and 2, the performance of all gradient-based methods deterio-
1840 rates markedly as data heterogeneity increases (i.e., as the partitioning parameters α or s decrease).
1841 In contrast, due to its ideal property of invariance to data heterogeneity, our DeepAFL maintains
1842 stable performance. This property of DeepAFL can be further extended to invariance with respect
1843 to the number of clients. As shown in Figure 7, the performance of DeepAFL remains entirely
1844 consistent across all scenarios, while its advantage over gradient-based methods (e.g., FedAvg) be-
1845 comes increasingly pronounced as the number of clients K grows. Quantitatively, at $K = 100$,
1846 DeepAFL achieves performance gains of 10.41% and 18.65% over FedAvg on the CIFAR-100 and
1847 Tiny-ImageNet, respectively, which expand to 25.97% and 32.8% at $K = 1000$. Notably, the results
1848 of AFL and FedAvg in Figure 7 are also taken from the given data in AFL (He et al., 2025b).
1849

1850 E.3.3 REPRESENTATION ANALYSES
1851

1852 Here, we provide a comprehensive analysis of DeepAFL’s capability for deep representation learn-
1853 ing. To this end, we examine how its training accuracy and testing accuracy evolve as the number of
1854 layers T increases, as shown in Tables 9–11. As observed, on the simple CIFAR-10 dataset, Deep-
1855 AFL’s performance improves steadily, reaching optimality at $T = 45$. On the complex CIFAR-100
1856 and Tiny-ImageNet datasets, it scales without overfitting up to $T = 50$, yielding over 1.5% improve-
1857 ment relative to $T = 20$. This indicates that for datasets with higher complexity, deeper models can
1858 be constructed to enhance representation capacity and achieve superior performance. Moreover, a
1859 particularly interesting and evident observation is that AFL exhibits consistently low training accu-
1860 racy, which can likely be attributed to the underfitting limitations of its simple single-layer linear
1861 model. In contrast, DeepAFL is able to significantly improve training accuracy by deepening the
1862 network, showing that its representation learning capability can effectively overcome the underfitting
1863 issues in traditional analytic learning. In fact, our DeepAFL only requires minimal computational
1864 overhead when increasing depth, with more detailed analyses presented in Appendix E.3.4.
1865

1866 E.3.4 EFFICIENCY EVALUATIONS
1867

1868 Here, we present the comprehensive efficiency evaluations of our DeepAFL. To highlight its superior
1869 balance between accuracy and efficiency, Figures 2–3 compare DeepAFL with baselines in terms of
1870 computational and communication cost. Specifically, to ensure a fair comparison, we selected the
1871 baseline with the lowest computational and communication cost (i.e., FedAvg) as the representative
1872 benchmark to highlight the superiority of our DeepAFL. Notably, the computational costs of all
1873 baselines are directly adopted from the existing results reported in AFL (He et al., 2025b). More-
1874 over, Tables 9–11 and Figures 4–6 provide further details on how our DeepAFL achieves a balance
1875 between accuracy and efficiency through the flexible adjustment of T .
1876

1877 Specifically, as shown in Figures 2–3, compared with gradient-based baselines, our DeepAFL (with
1878 $T = 20$) achieves at least a 99.7% reduction in computational cost and a 50.2% reduction in commu-
1879 nication cost on the CIFAR-100. On the Tiny-ImageNet, the advantages of DeepAFL correspond to
1880 a 99.6% reduction in computational cost and a 70.1% reduction in communication cost. By flexibly
1881 adjusting the number of layers T , for example, setting $T = 5$, users can obtain greater efficiency ad-
1882 vantages. In summary, although the cost of our DeepAFL is inevitably higher than that of AFL due to
1883 the additional deep layers required for enhanced representations, it remains substantially lower than
1884 that of gradient-based baselines. This advantage stems from our DeepAFL’s gradient-free manner,
1885 which avoids the iterative costly computation and communication associated with gradients.
1886

1887 Subsequently, we also report the detailed accuracy-efficiency balance of our DeepAFL with varying
1888 numbers of layers, as shown in Tables 9–11 and Figures 4–6. Encouragingly, the construction of
1889 each additional layer in DeepAFL requires no more than 3s across all three datasets, providing an
1890 intuitive demonstration of its efficiency. Compared with AFL, even when building a 50-layer deep
1891 network, DeepAFL incurs less than a twofold increase in training time while delivering performance
1892 improvements of at most up to 31.38% on the training set and 9.95% on the testing set. It is worth
1893 noting that, there is considerable potential for our DeepAFL to further reduce communication cost
1894 in the future through compression techniques such as matrix factorization, as its primary communi-
1895 cation information is the *Auto-Correlation* and *Cross-Correlation* Matrices.
1896

1890
1891

E.3.5 PARAMETER ANALYSES

1892
1893
1894
1895

Here, we provide comprehensive analyses for the parameter sensitivity of our DeepAFL, encompassing the regularization parameters λ , γ , the activation function $\sigma(\cdot)$, and the projection dimensions d_Φ , d_F . Specifically, we keep all parameters' default settings, varying only the parameter under investigation to assess its sensitivity. The specific analyses for each parameter are detailed below:

1896
1897
1898
1899
1900
1901
1902

First, we explore the regularization parameters λ and γ , which regularize the global classifier and the transformation matrix, respectively. As illustrated in Tables 3–4, DeepAFL is insensitive to λ but sensitive to γ . This evidence may stem from the fact that our deep residual network employs only the final global classifier during inference, yet relies on all transformation matrices. Specifically, on the simple CIFAR-10 dataset, the best performance is achieved when $\lambda = 10$ and $\gamma \in [0.1, 0.5]$. For the complex CIFAR-100 and Tiny-ImageNet datasets, the optimal values of λ are found to be highly dispersed, while the optimal γ consistently occurs at the smaller value of $\gamma = 0.01$.

1903
1904
1905
1906
1907
1908
1909

Second, we further analyze the impact of different activation functions on DeepAFL's performance. The detailed experimental results are shown in Table 5. The selection of the activation function markedly influences DeepAFL's performance. More specifically, a performance variance of about 2% can be observed among different activation functions, with GELU emerging as the optimal choice and Softshrink as the least effective. Moreover, compared to omitting the activation function entirely, employing GELU yields up to 5% gains, underscoring its critical role. More comprehensive analysis of this significance is available in the corresponding ablation study in Appendix E.3.6.

1910
1911
1912
1913
1914
1915
1916
1917

Third, we study the effects of varying projection dimensions d_Φ and d_F on our DeepAFL in Tables 6–8. As these projection dimensions increase, DeepAFL's performance rises initially before declining. This behavior can be attributed to the fact that, while larger dimensions can enhance the model's expressive power by capturing more information, they also concurrently increase the propensity for overfitting and compromise numerical stability. Notably, the model crashes when $d_\Phi = d_F = 2^{13}$. Furthermore, as detailed in Appendix D, the complexity of DeepAFL scales at least quadratically with these dimensions, and excessively large values incur substantial overhead, cautioning against the indiscriminate pursuit of larger dimensions. That's why we select $d_\Phi = d_F = 2^{10}$ as default.

1918
1919

E.3.6 ABLATION STUDIES

1920
1921
1922
1923
1924
1925
1926

Here, we present the detailed analyses of our ablation studies to dissect the individual contributions of residual skip connections, random projections, activation functions, and trainable transformations in our DeepAFL. Specifically, we construct four ablation models, with each omitting one of these key components, and compare them against our full DeepAFL across varying layer depths on diverse datasets, as shown in Tables 12–17 and Figures 8–10. Moreover, the ablation studies are conducted under identical conditions for fair comparisons, with all parameters aligned with those used in the main experiments. The detailed analyses are provided below:

1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943

First, we focus on analyzing the contribution of the residual skip connections in our DeepAFL, with the corresponding ablation model denoted as Ablation (1). For the simple CIFAR-10, ablating the residual connections prevents performance improvements as the network depth increases, as shown in Tables 12–13 and Figure 8, indicating that this dataset is highly sensitive to such skip connections. On the CIFAR-100 and Tiny-ImageNet datasets, performance continues to improve even without skip connections, exhibiting an initial rise followed by convergence, as shown in Tables 14–17 and Figures 9–10. Meanwhile, the ablation model attains respectable performance on Tiny-ImageNet, yet exhibits markedly limited gains on CIFAR-100. These disparities across datasets may stem from their differing complexities, with CIFAR-10, CIFAR-100, and Tiny-ImageNet containing 10, 100, and 200 classes, respectively. Notably, compared to DeepAFL, the ablation model yields substantially inferior performance across all cases. This degradation stems from the fact that, without skip connections, the model encounters an information bottleneck akin to that in gradient-based deep learning. Crucially, the importance of residual skip connections is theoretically grounded, as their ablation would invalidate Theorems 2 and 3. Specifically, the residual skip connections are crucial because they enable a special case: if the transformation matrix is a zero matrix, no update is applied to the representation features. This functionality guarantees that the representation features learned at each layer will be at least as good as those from the preceding layer. Thus, without the residual skip connections, each layer essentially constructs new representation features from the hidden random features, thereby forfeiting these valuable theoretical guarantees.

1944 Second, we further analyze the contributions of the random projections, which give the *stochasticity*
 1945 to our DeepAFL. Ablating the random projections $\{\mathbf{B}_i\}_{i=1}^T$ can be achieved by simply setting them
 1946 to the identity matrices, and we denote this ablated model as Ablation (2). While this ablated model
 1947 still yields performance improvements across all datasets, its performance growth rate is markedly
 1948 slower compared to DeepAFL, and it reaches convergence earlier. Specifically, at $T = 10$, the abla-
 1949 tion model reaches convergence across all three datasets, exhibiting performance gaps of 3.64% to
 1950 9.25% on the training set and 1.90% to 3.43% on the testing set compared to our DeepAFL. Sub-
 1951 sequently, DeepAFL’s performance continues to improve with increasing layers, further widening the
 1952 performance gap. By $T = 50$, this gap further expands to 8.68% to 22.2% on the training set and
 1953 2.70% to 5.87% on the testing set. In Figures 8–10, we can observe that without random projections
 1954 to introduce the essential *stochasticity*, the model tends to converge to local optima or saddle points,
 1955 even though the projection itself does not alter the feature dimensionality, as $d_\Phi = d_F = 2^{10}$.
 1956

1957 Third, we focus on analyzing the contributions of the activation function in our DeepAFL, which
 1958 introduces *nonlinearity* to our DeepAFL. The corresponding ablation model is named Ablation (3).
 1959 It can be observed that the performance of this ablation model remains nearly invariant as the layer
 1960 depth T increases. Specifically, the performance fluctuation from $T = 0$ to $T = 50$ consistently
 1961 falls below 0.05% across all datasets, thereby underscoring the critical importance of the activation
 1962 function in DeepAFL. This is because, without the activation function, each residual block merely
 1963 performs a purely linear transformation on the feature representations, fundamentally limiting the
 1964 model’s representative power and preventing further performance improvement with added depth.
 1965

1966 Fourth, we study the contributions of the trainable transformation Ω_{t+1} , which imparts *learnability*
 1967 to our DeepAFL. The corresponding ablation model is denoted as ablation (4). As illustrated in
 1968 Figures 8–10, this ablation model exhibits modest performance gains followed by rapid convergence
 1969 on the simple CIFAR-10, while showing nearly no improvement on the complex CIFAR-100 and
 1970 Tiny-ImageNet. This evidence arises because, without the trainable transformation for *learnability*,
 1971 each residual block essentially functions as a fixed random nonlinear feature updater. For the simple
 1972 CIFAR-10, these updaters can fortuitously render random features progressively more discriminative
 1973 to a limited extent. For complex CIFAR-100 and Tiny-ImageNet, relying solely on random feature
 1974 updates fails to capture the high-level features, leading to negligible performance improvement.
 1975

1976 E.3.7 COMPARING DEEPAFL WITH OTHER DEEPENING STRATEGIES

1977 Here, we further compare our DeepAFL with the other four alternative deepening strategies to high-
 1978 light its superiority. Specifically, we design four distinct deepening strategies: (1) cascades with
 1979 random features, (2) cascades with random features and label encoding, (3) cascades with activated
 1980 random features, and (4) cascades with activated random features and label encoding. Notably, the
 1981 strategies (1) and (3) here are similar to the naive approaches (a) and (b) in Figure 1. However, in our
 1982 comparison, strategies (1) and (3) additionally incorporate random projection and activation after the
 1983 backbone to form the zero-layer features Φ_0 , thereby aligning more closely with our DeepAFL at
 1984 the starting point to ensure fairness in comparison. These two naive strategies are consistent with
 1985 most existing attempts in the literature to deepen analytic networks (Low et al., 2019). Meanwhile,
 1986 strategies (2) and (4) are essentially extensions of (1) and (3) through the incorporation of label
 1987 encoding, which is an established technique in analytic learning that introduces an additional linear
 1988 mapping for the labels at each layer to facilitate the training of deep analytic networks (Zhuang et al.,
 1989 2021; 2025). The detailed experimental results are presented in Tables 18–23 and Figures 11–13.

1990 From the results, our DeepAFL consistently outperforms all other deepening strategies across all
 1991 datasets, thereby underscoring its superiority. Across all datasets, strategies (1) and (2) achieve
 1992 very similar performance, exhibiting nearly no improvement with an increasing number of layers.
 1993 Next, we turn our attention to strategies (3) and (4), which differ in that strategy (4) incorporates
 1994 additional label encoding. On the CIFAR-10, strategies (3) and (4) yield comparable performance,
 1995 both showing some improvement followed by rapid convergence. On the CIFAR-100, strategy (4)
 1996 lags substantially behind strategy (3), despite both showing gains with increasing layer depth. This
 1997 disparity widens further on the Tiny-ImageNet, where strategy (4) experiences a pronounced per-
 1998 formance decline from added depth. The varying effects of label encoding across datasets may stem
 1999 from differences in the number of classes, as these datasets contain 10, 100, and 200 classes, respec-
 2000 tively. Taken together, these results suggest that existing deepening strategies are of very limited
 2001 effectiveness and fall far short of the performance achieved by our proposed DeepAFL.

1998 **F ADDITIONAL DISCUSSION**

1999

2000 Beyond the discussion in our main text, we provide additional discussion here. A key feature of
 2001 our DeepAFL is its capability to perform gradient-free representation learning after the feature ex-
 2002 traction backbone. This stands in contrast to traditional methods that typically apply representation
 2003 learning directly to the backbone model. A natural question arises: what are the differences and
 2004 advantages of our approach compared to direct representation learning on the backbone?

2005 Direct representation learning on the backbone typically involves two main strategies.

- 2006 • The first strategy is to fine-tune the backbone ***during the FL training process***. This aims to
 2007 adapt the model to the data distribution within the FL system. However, this approach re-
 2008 quires enabling gradient-based training in the FL training, which can be severely impacted
 2009 by data heterogeneity. Moreover, there is a high risk of damaging the knowledge acquired
 2010 during pre-training, known as catastrophic forgetting. As many existing works have shown
 2011 that fine-tuning the backbone with gradients during FL training may be counterproductive,
 2012 it is precisely why we developed DeepAFL for gradient-free representation learning.
- 2013 • The second strategy is to build a more complex backbone ***during the pre-training process***,
 2014 leveraging massive pre-training datasets to enhance its robustness. Our approach is orthog-
 2015 onal to this strategy, as our DeepAFL is designed to further improve the representations
 2016 during the FL training process, given any fixed, pre-trained backbone. This training-process
 2017 representation learning is both necessary and meaningful, as the pre-trained backbone will
 2018 inevitably exhibit a domain shift when applied to a new FL system.

2019 In addition, another noteworthy aspect of our work is the effectiveness of residual connections in our
 2020 DeepAFL, especially since its gradient-free nature seems to be at odds with the original gradient-
 2021 based motivation for these connections in ResNet. We believe that it is a profound theoretical ques-
 2022 tion that warrants future research. Here, we would like to offer several intuitive explanations for it:

- 2023 • First, our comprehensive ablation studies show that the model without residual connections
 2024 struggles to improve its training accuracy as the number of layers increases, let alone testing
 2025 accuracy. This situation mirrors the challenges faced by gradient-based deep networks be-
 2026 fore the introduction of ResNet. Deep networks without residual connections can similarly
 2027 suffer from an information bottleneck as they continually update feature representations.
 2028 By adding a skip connection, each layer’s representation learning is built upon the previous
 2029 layer’s foundation, enabling continuous and cumulative improvements in representations.
- 2030 • Second, as we prove in Theorems 2 and 3, our DeepAFL’s representation learning capabili-
 2031 ty is primarily reflected in the non-increasing nature of its empirical risk with added layers.
 2032 This beneficial property is ensured by the use of skip connections. When network features
 2033 converge, each residual block can be just set to zero, at a minimum, to achieve an unchang-
 2034 ing empirical risk, thereby guaranteeing that the overall empirical risk of our DeepAFL will
 2035 not increase. Without this design, the ablated model would no longer be able to guarantee
 2036 these theorems, severely compromising its representation learning capability.
- 2037 • Third, the philosophy of our DeepAFL intrinsically aligns with that of *Gradient Boosting*.
 2038 Specifically, under the chosen MSE loss function, the negative gradient of the loss function
 2039 is precisely the residual. Therefore, by continuously fitting the residuals from previous lay-
 2040 ers, our DeepAFL can be seen as learning *pseudo-gradients*. This perspective may further
 2041 help unify the understanding of both gradient-based and gradient-free learning approaches.

2042 In summary, we believe our DeepAFL represents a significant first step. There is substantial potential
 2043 for future theoretical and practical exploration based on this foundation, which we hope will further
 2044 advance various fields, including FL, analytic learning, and representation learning, etc.

2045 **G STATEMENT ON THE USAGES OF LARGE LANGUAGE MODELS**

2046

2047 In adherence to the ICLR 2026 policy, we report the use of Large Language Models (LLMs) during
 2048 the preparation of this paper. Here, the only usages of LLMs were to aid and polish the writing
 2049 and illustration. Specifically, the LLMs were utilized to improve sentence structure, enhance clarity,
 2050 ensure grammatical accuracy, and optimize the illustration. Furthermore, all the core scientific con-
 2051 tributions (including the hypothesis formulations, the theoretical analyses, the experimental designs,
 the data analyses, and the final conclusions, etc.) are the original work of the human authors.

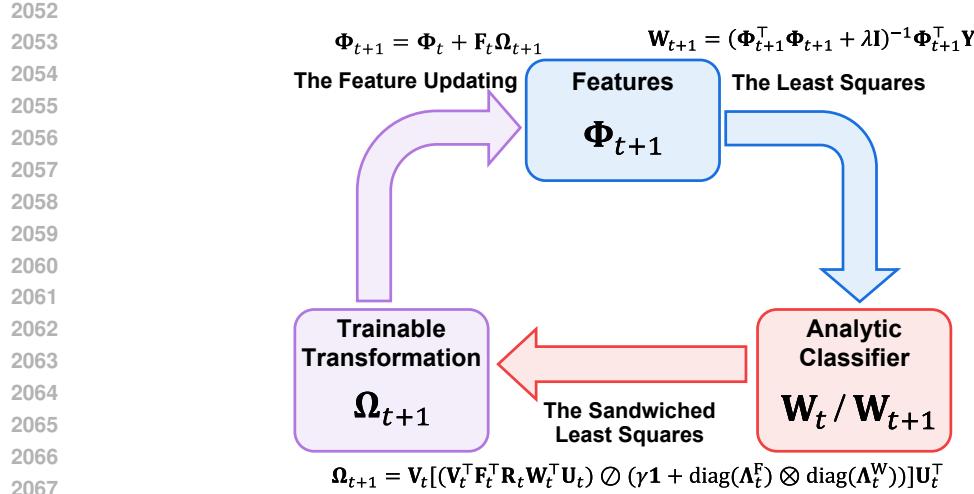


Figure 14: The training process for each layer in our DeepAFL.

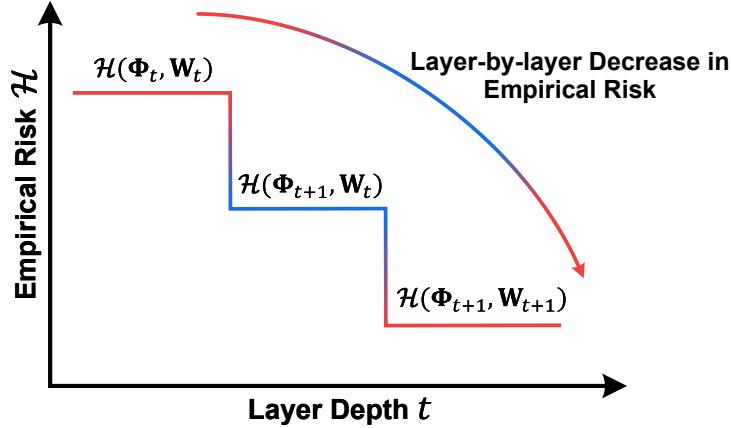


Figure 15: Layer-wise reduction of empirical risk in our DeepAFL.

The Sandwiched Least-Squares Process

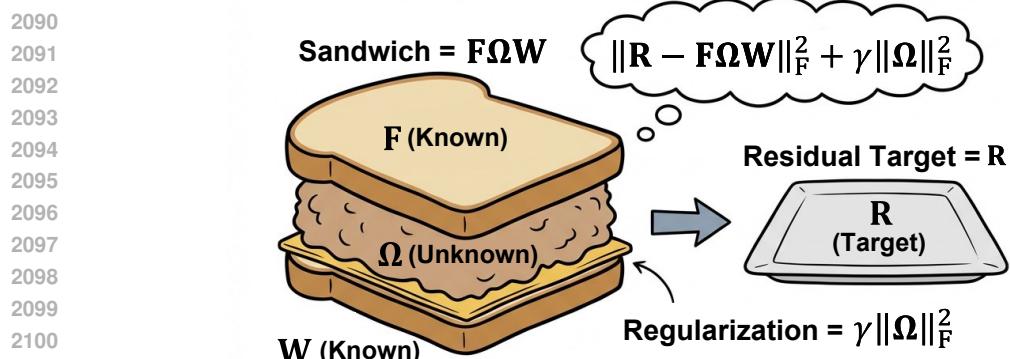


Figure 16: The sandwiched least-squares process in our DeepAFL.

2106
 2107 Table 24: Performance comparisons of the top-1 accuracy (%) among our DeepAFL and two new
 2108 baselines (i.e., FedAWA (Shi et al., 2025) and (Liu et al., 2024a)) on CIFAR-10, CIFAR-100, and
 2109 Tiny-ImageNet. The best result is highlighted in **bold**, and the second-best result is underlined.

Baseline	CIFAR-10		CIFAR-100		Tiny-ImageNet	
	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.1$	$\alpha = 0.1$	$\alpha = 0.1$	$\alpha = 0.1$
FedAWA (2025)	79.42%	76.37%	<u>58.94%</u>	50.48%	<u>59.07%</u>	52.33%
FedLPA (2024a)	49.57%	42.56%	40.04%	22.55%	39.48%	26.67%
AFL (2025b)	<u>80.75%</u>	<u>80.75%</u>	58.56%	<u>58.56%</u>	54.67%	<u>54.67%</u>
DeepAFL ($T = 5$)	85.20%	85.20%	64.72%	64.72%	60.31%	60.31%
DeepAFL ($T = 10$)	85.93%	85.93%	65.96%	65.96%	61.37%	61.37%
DeepAFL ($T = 20$)	86.43%	86.43%	66.98%	66.98%	62.35%	62.35%
Improvement \uparrow	5.68%	5.68%	8.04%	8.42%	3.28%	7.68%

H MORE EXPLANATIONS ON THE DETAILED PROCESS OF DEEP AFL

2123 Here, we provide comprehensive descriptions of our DeepAFL process to enhance clarity. Specifically, as illustrated in Figures 14–16, we elaborate on the training process of each layer, demonstrate
 2124 the layer-by-layer empirical risk reduction, and visually depict the sandwiched least-squares mech-
 2125анизm that underpins our DeepAFL. Further details are presented below.

2126 First of all, let’s focus on the training process for each layer in our DeepAFL. As illustrated in
 2127 Figure 14, each layer entails the sequential update and computation of the trainable transformation
 2128 matrix Ω_{t+1} , the feature matrix Π_{t+1} , and the analytic classifier \mathbf{W}_{t+1} . Specifically, for each $(t + 1)$ -th layer,
 2129 the trainable transformation matrix Ω_{t+1} is initially derived through the sandwiched least
 2130 squares process, utilizing the preceding analytic classifier \mathbf{W}_t alongside \mathbf{F}_t and \mathbf{R}_t . Subsequently,
 2131 based on the newly acquired Ω_{t+1} , the feature matrix is updated to yield Φ_{t+1} according to Equation
 2132 (equation 9). Finally, employing the derived feature matrix Φ_{t+1} , the analytic classifier is further
 2133 updated to obtain \mathbf{W}_{t+1} using the least squares process. Notably, as a special case, the zero-layer
 2134 feature matrix Φ_0 is obtained directly via feature extraction using Equations (1) and (2), thereby
 2135 eliminating the need to compute a trainable transformation matrix Ω_0 . The aforementioned process
 2136 is executed layer-by-layer, adhering to the update sequence $\Phi_0 \mapsto \mathbf{W}_0 \mapsto \dots \mapsto \mathbf{W}_{T-1} \mapsto \Omega_T \mapsto$
 2137 $\Phi_T \mapsto \mathbf{W}_T$ until the network construction is complete.

2138 Subsequently, we present intuitive and visual explanations of the monotonic decrease in empirical
 2139 risk as the depth of our DeepAFL increases. The empirical risk $\mathcal{H}(\Phi_t, \mathbf{W}_t)$ is determined by the
 2140 feature matrix Φ_t and the analytic classifier \mathbf{W}_t . As previously analyzed, within each layer, Deep-
 2141 AFL sequentially updates the feature matrix and the analytic classifier to yield improved Φ_{t+1} and
 2142 \mathbf{W}_{t+1} . Consequently, the empirical risk undergoes two reductions within each layer, as depicted
 2143 in Figure 15. We now detail these two reduction steps within each layer. Specifically, the fea-
 2144 ture matrix Φ_t is first updated to obtain Φ_{t+1} via the trainable transformation matrix Ω_{t+1} , which
 2145 constitutes the optimal solution to Equation (7) for minimizing the empirical risk with respect to
 2146 feature optimization. Thus, the updated empirical risk $\mathcal{H}(\Phi_{t+1}, \mathbf{W}_t)$ is lower than the initial risk
 2147 $\mathcal{H}(\Phi_t, \mathbf{W}_t)$. Subsequently, with Φ_{t+1} held fixed, the new analytic classifier \mathbf{W}_{t+1} is computed as
 2148 the optimal solution to Equation (4) for minimizing the local empirical risk. Thus, the final empirical
 2149 risk $\mathcal{H}(\Phi_{t+1}, \mathbf{W}_{t+1})$ is also lower than the risk $\mathcal{H}(\Phi_{t+1}, \mathbf{W}_t)$. By executing this procedure layer-
 2150 by-layer, the empirical risk is progressively reduced, thereby leading to enhanced performance.

2151 Finally, as we term this optimization process for Ω as the *sandwiched least squares problem*, we
 2152 would like to provide a detailed illustration of its process. Specifically, the optimization objective in
 2153 Equation (7) can be viewed as a special case of generalized Sylvester matrix equations (Wu et al.,
 2154 2008; Ding et al., 2008; Duan, 2015). This structure is characterized by the unknown variable Ω
 2155 being *sandwiched* between two known matrices, \mathbf{F} and \mathbf{W} . In this context, the “*sandwich*” structure,
 2156 therefore, refers to this three-matrix product form $\mathbf{F}\Omega\mathbf{W}$. The empirical risk minimization aims to
 2157 minimize the residual, meaning the sandwich term should be as close as possible to the residual
 2158 target \mathbf{R} . Meanwhile, the regularization term $\gamma\|\Omega\|_F^2$ is applied to constrain the magnitude of Ω and
 2159 prevent it from becoming excessively large. Consequently, this particular structure facilitates the
 derivation of a distinct analytical solution, as presented in Equation (8).

2160 Table 25: Accuracy of gradient-based baselines with varying layers on CIFAR-100.
2161

Layers	$T = 1$	$T = 2$	$T = 5$	$T = 10$	$T = 20$
FedAvg	56.62%	55.83%	<5%	<3%	<1%
FedDyn	57.55%	57.22%	56.26%	<3%	<1%

2166 I MORE EXPERIMENTAL RESULTS ON RECENT BASELINES
2167

2169 Following a reviewer’s suggestion, we implemented two new baselines (i.e., FedAWA (Shi et al.,
2170 2025) and FedLPA (Liu et al., 2024a)) across the three benchmark datasets utilized in our main ex-
2171 periment: CIFAR-10, CIFAR-100, and ImageNet-R. All experimental settings were maintained to
2172 be consistent with the other baselines established in our manuscript. Specifically, because FedLPA
2173 is classified as a one-shot communication FL method, and we empirically observed that its perfor-
2174 mance essentially converges within 50 local epochs, we set the local epochs to 50 for this baseline.
2175 This setting allows us to mitigate unnecessary computational overhead for FedLPA, thereby more
2176 fully demonstrating its potential for efficiency. We present the results of these two new baselines
2177 alongside AFL and our DeepAFL for additional comparison, as shown in Table 24.

2178 In terms of accuracy, FedAWA exhibits very strong results, approaching (on CIFAR-10) or even
2179 surpassing AFL (on CIFAR-100 and Tiny-ImageNet) when $\alpha = 0.1$. Notably, even though FedAWA
2180 outperforms AFL, our DeepAFL still consistently achieves the best performance across all scenarios.
2181 Furthermore, as a gradient-based one-shot method, FedLPA’s performance is compromised by data
2182 heterogeneity, as it lacks the inherent invariance demonstrated by our DeepAFL. Consequently,
2183 despite our best efforts to tune its hyperparameters, FedLPA still performs poorly. Furthermore, as
2184 the degree of Non-IID data increases, a pronounced performance degradation is observed for both of
2185 the newly introduced gradient-based baselines (i.e., FedAWA and FedLPA). This observation further
2186 highlights the advantage of the inherent invariance property of our DeepAFL.

2187 In terms of efficiency, using the CIFAR-100 dataset as an example, FedAWA requires approximately
2188 10 hours, while FedLPA requires approximately 2.5 hours. In sharp contrast, our DeepAFL com-
2189 pletes the task in less than 100 seconds (specifically, 91.74 seconds), achieving a speedup exceeding
2190 $90\times$. Notably, even though FedLPA is also classified as a one-shot FL approach, and we have already
2191 minimized its number of local epochs to fully reflect its efficiency potential, its overall overhead re-
2192 mains significantly higher than that of DeepAFL. This phenomenon is primarily due to FedLPA still
2193 being gradient-based and facing time-consuming backpropagation. Conversely, our DeepAFL is in
2194 a forward-only manner, further highlighting its efficiency advantage beyond its one-shot nature.

2195 J MORE EXPLANATIONS ON OUR DEEPAFL’S EXPERIMENTAL SETUP
2196

2197 Here, we would like to clarify our experimental setup. To maintain fairness in comparisons, the
2198 results for all baselines in our paper are directly employed from the benchmark provided in the
2199 original AFL paper (He et al., 2025b). Specifically, the gradient-based approaches in the original
2200 AFL paper adhered to a similar setup as AFL, where the backbone network is entirely frozen during
2201 the FL process, and only the one-layer classifier is trained.

2202 Moreover, to address potential concerns that gradient-based baselines might benefit from multiple
2203 trainable layers, we conduct additional experiments on FedAvg and FedDyn with multiple trainable
2204 layers ($T \in \{2, 5, 10, 20\}$) on CIFAR-100 under the Dirichlet distribution ($\alpha = 0.1$). For these
2205 experiments, the feature dimensions of the additional layers are aligned with those of our DeepAFL
2206 (i.e., 1024) to ensure consistency. Detailed results are presented in Table 25.

2207 Through these results, we can observe that the performance of the gradient-based methods gradually
2208 and constantly degrades as T increases. Specifically, when T increases beyond a certain threshold
2209 (e.g., $T \geq 10$), the excessive number of parameters, combined with the Non-IID data in the FL
2210 scenario, may cause the training process to collapse. These findings fully indicate that directly in-
2211 creasing the number of trainable layers after the backbone for these baselines, similar to our Deep-
2212 AFL setup, has a detrimental effect on their performance. Consequently, far from benefiting from
2213 multiple trainable layers, these gradient-based baselines exhibit significantly inferior performance
under such configurations, thereby validating the propriety of our experimental setup.

2214 K MORE ANALYSES ON DEEPAFL’S GENERALIZATION

2216 Here, we present detailed analyses of the generalization ability of our DeepAFL. To capture the full
 2217 scope of this attribute, we dissect our analyses into two key aspects: (1) generalization from training
 2218 data to test data, and (2) generalization (or transferability) from IID to Out-Of-Distribution (OOD).
 2219 The detailed analyses are presented below.

2220 First, let’s focus on the generalization from training data to test data. Specifically, we have already
 2221 established a strong theoretical guarantee in Theorems 2 and 3, which prove the monotonic decrease
 2222 of empirical risk in DeepAFL. In this case, we assume that the training and test data remain IID,
 2223 which is the common practice and prerequisite in almost all supervised statistical learning theory.
 2224 More encouragingly, these theoretical guarantees can be further extended to imply a better Gener-
 2225 alization Error Bound (GEB) in statistical learning theory, as the complexity of the linear analytic
 2226 classifier remains constant at each layer.

2227 Below, we provide detailed theoretical explanations of the GEB. Specifically, in statistical learning
 2228 theory, the GEB typically follows the core form:

$$2230 \quad \mathcal{R}(f) \leq \hat{\mathcal{R}}(f) + \text{Complexity Term}, \quad (78)$$

2231 where $\mathcal{R}(f)$ denotes the generalization risk (or generalization error) and $\hat{\mathcal{R}}(f)$ denotes the empirical
 2232 risk, which is identical to \mathcal{H} defined in Theorem 2. The complexity term (i.e., complexity gap) mea-
 2233 sures the model’s complexity, which is typically related to the sample size, the feature dimension,
 2234 and the data’s geometric structure. The size of the GEB, $\mathcal{B}(f)$, is what we are concerned with:

$$2236 \quad \mathcal{B}(f) = \hat{\mathcal{R}}(f) + \text{Complexity Term}. \quad (79)$$

2237 According to standard *VC-dimension* and *Rademacher* complexity generalization bounds (Alain
 2238 & Bengio, 2016; Vapnik, 1999), for an analytic linear hypothesis class in \mathbb{R}^d , the generalization
 2239 gap/complexity is bounded by $O(\sqrt{d_\Phi \log(N/\delta)/N})$. Since N (data size), $d_\Phi = 1024$ (feature
 2240 dimension), and $\delta \in (0, 1)$ (confidence level) remain constant in our DeepAFL, the linear analytic
 2241 classifiers at each layer of our DeepAFL can be considered to share the same complexity term.
 2242 Furthermore, the aforementioned bound implicitly relies on the assumption that the input features
 2243 are bounded. This prerequisite is rigorously satisfied in our DeepAFL, as we enforce constant L2-
 2244 normalization as a constraint throughout the training process, as detailed in Equations (4) and (7).

2245 Subsequently, considering the analytic classifiers f_t and f_{t+1} trained on the representations of layer
 2246 t and $(t+1)$ respectively, we have:

$$2248 \quad \mathcal{B}(f_t) = \hat{\mathcal{R}}(f_t) + \text{Complexity Term}_t, \quad (80)$$

$$2250 \quad \mathcal{B}(f_{t+1}) = \hat{\mathcal{R}}(f_{t+1}) + \text{Complexity Term}_{t+1}. \quad (81)$$

2251 As we have established that the complexity term is constant across layers, and our paper proves
 2252 that the empirical risk is non-increasing (i.e., $\hat{\mathcal{R}}(f_t) \geq \hat{\mathcal{R}}(f_{t+1})$), it directly follows that $\mathcal{B}(f_t) \geq$
 2253 $\mathcal{B}(f_{t+1})$. In this way, we successfully show that the theory introduced in our paper can be easily
 2254 extended to deduce that the GEB of the analytic classifier is also non-increasing. This result pro-
 2255 vides stronger evidence for the effectiveness and training-to-testing generalization of our DeepAFL’s
 2256 representation learning capabilities.

2257 Second, we further turn our attention to the discussion of IID-to-OOD generalization (i.e., transfe-
 2258 rability), which signifies the capacity of learned features to be effectively applied to cross-domain
 2259 applications. It is important to note that transferability is not a necessary prerequisite for successful
 2260 representation learning but pertains to the distinct field of transfer learning. In fact, the vast majority
 2261 of current gradient-based supervised methods also fall far short of robust transferability, as they tend
 2262 to focus narrowly on the current task and inherently suffer when faced with out-of-distribution data.
 2263 As the development of analytic-learning-based methods is still in its nascent stages, our current core
 2264 aim is to achieve direct performance improvement with marginal additional cost. Consequently,
 2265 since we make no explicit claims regarding the transferability of our DeepAFL, the absence of
 2266 immediate cross-domain capabilities does not constitute a severe weakness for our current work.
 2267 Furthermore, we believe that exploring how to improve the transferability of the underlying model
 2268 is a promising and interesting direction for future work.

2268 **L MORE ANALYSES ON DEEPAFL’S SCALABILITY WITH ViT BACKBONES**
22692270 Table 26: Analyses of DeepAFL’s scalability with the ViT-B-16-I-1K backbone on CIFAR-100.
2271

Layers	Testing Acc	$\Delta_{\text{Testing Acc}}$	$\Delta_{\text{AFL Acc}}$	Time Cost	$\Delta_{\text{Time Cost}}$	$\Delta_{\text{AFL Cost}}$
AFL	75.45%	/	/	107.19s	/	/
DeepAFL ($T = 5$)	78.05%	2.60%	2.60%	118.97s	11.78s	11.78s
DeepAFL ($T = 10$)	79.01%	0.96%	3.56%	126.55s	7.58s	19.37s
DeepAFL ($T = 15$)	79.73%	0.72%	4.28%	133.79s	7.24s	26.60s
DeepAFL ($T = 20$)	80.13%	0.40%	4.68%	141.02s	7.23s	33.83s
DeepAFL ($T = 25$)	80.33%	0.20%	4.88%	148.56s	7.54s	41.37s
DeepAFL ($T = 30$)	80.51%	0.18%	5.06%	156.21s	7.65s	49.02s

2279 Table 27: Analyses of DeepAFL’s scalability with the ViT-B-16-I-21K backbone on CIFAR-100.
2280

Layers	Testing Acc	$\Delta_{\text{Testing Acc}}$	$\Delta_{\text{AFL Acc}}$	Time Cost	$\Delta_{\text{Time Cost}}$	$\Delta_{\text{AFL Time}}$
AFL	86.35%	/	/	107.19s	/	/
DeepAFL ($T = 5$)	87.71%	1.36%	1.36%	118.97s	11.78s	11.78s
DeepAFL ($T = 10$)	88.15%	0.44%	1.80%	126.55s	7.58s	19.37s
DeepAFL ($T = 15$)	88.49%	0.34%	2.14%	133.79s	7.24s	26.60s
DeepAFL ($T = 20$)	88.80%	0.31%	2.45%	141.02s	7.23s	33.83s
DeepAFL ($T = 25$)	88.95%	0.15%	2.60%	148.56s	7.54s	41.37s
DeepAFL ($T = 30$)	89.08%	0.13%	2.73%	156.21s	7.65s	49.02s

2289 Here, we present detailed analyses of DeepAFL’s practical scalability with ViT backbones. Given
2290 that our DeepAFL is built upon AFL (He et al., 2025b), which uses ResNet-18 as its primary backbone,
2291 we also employ the aligned backbone in our main experiments. This alignment facilitates a
2292 clearer and more transparent comparison between the two methods. Here, we further extend our
2293 evaluation by adopting larger ViT models to investigate our DeepAFL’s practical scalability.

2294 Specifically, we select two versions of the ViT-B-16 backbone: **ViT-B-16-I-1K** (pre-trained on the
2295 ImageNet-1K with approximately 1.28 million image) and **ViT-B-16-I-21K** (pre-trained on the
2296 larger ImageNet-21K dataset with approximately 14 million images). Owing to the significantly
2297 broader pre-training scale, ViT-B-16-I-21K typically yields better performance than ViT-B-16-I-1K.
2298 Since the backbone sizes of both ViT models are identical (Base-16), the runtime for our DeepAFL
2299 on both models remains essentially the same. Furthermore, all experiments are conducted on the
2300 CIFAR-100 dataset with 100 clients using a single NVIDIA RTX 4090 GPU.

2301 The detailed results are presented in Tables 26–27, where “Testing Acc” and “Time Cost” represent
2302 the test accuracy and total wall-clock runtime. The Δ columns denote the step-wise change relative
2303 to the preceding row, while the Δ_{AFL} columns indicate the cumulative difference from AFL.

2304 From these results, it is evident that applying DeepAFL to both ViT backbones leads to significant
2305 performance improvements with minimal time cost, thereby demonstrating its robust scalability.
2306 Specifically, at $T = 30$, our DeepAFL achieves an accuracy gain of 5.06% over AFL with the ViT-B-
2307 16-I-1K backbone, and an improvement of 2.73% with the ViT-B-16-I-21K backbone. Remarkably,
2308 these substantial improvements are achieved with a runtime of merely 156.21 s, which is only 49.02 s
2309 higher than AFL. It is important to note that this reported runtime encompasses the cumulative time
2310 consumed by all 100 clients to complete the training process across all layers on a single 4090 GPU.
2311 Thus, the average cost per client is less than 1.56 s. In stark contrast, even the simplest and most
2312 efficient FL baseline (e.g., FedAvg) requires a Total Wall-Clock Runtime of over 33,000 s.

2313 Furthermore, although we observe that the performance improvement of DeepAFL exhibits diminish-
2314 ing marginal returns, this phenomenon is entirely expected and normal. First, as the backbone
2315 capability strengthens, the feature representations are already highly optimized, making it increas-
2316 ingly challenging to extract further performance gains beyond such a high baseline. Second, while
2317 increasing the number of layers enhances the fitting capacity of DeepAFL, the fixed volume of
2318 training data inevitably leads to a saturation limit in model performance. Notably, such diminish-
2319 ing returns are also common in gradient-based methods. Therefore, the ability of DeepAFL to achieve
2320 stable and significant improvements across various backbones without increasing any training data
2321 volume is highly noteworthy. Especially when considering the negligible runtime cost, the per-
2322 formance gains delivered by DeepAFL can be regarded as incurring virtually no time overhead.

2322 **M MORE ANALYSES ON PARTIAL CLIENT PARTICIPATION**

Table 28: Performance evaluation of DeepAFL on CIFAR-100 under partial client participation.

Layers	Consistent Participation						Inconsistent Participation					
	$\eta = 100\%$	$\eta = 90\%$	$\eta = 80\%$	$\eta = 70\%$	$\eta = 60\%$	$\eta = 50\%$	$\eta = 100\%$	$\eta = 90\%$	$\eta = 80\%$	$\eta = 70\%$	$\eta = 60\%$	$\eta = 50\%$
$T = 5$	64.72%	64.53%	64.38%	63.92%	62.85%	61.54%	64.72%	64.21%	64.03%	64.05%	62.19%	61.36%
$T = 10$	65.96%	65.46%	65.26%	64.49%	64.76%	63.17%	65.96%	65.36%	65.05%	64.45%	64.09%	62.93%
$T = 15$	66.56%	66.39%	66.20%	65.38%	65.17%	65.08%	66.56%	66.22%	66.17%	65.64%	64.60%	64.98%
$T = 20$	66.98%	66.76%	66.66%	66.47%	66.27%	65.14%	66.98%	66.57%	66.51%	66.32%	65.50%	64.81%

Here, we provide detailed analyses of our DeepAFL’s performance in scenarios involving partial client participation, thereby further highlighting its robustness. The detailed analyses are as follows.

First, let’s focus on the minor mechanism adjustments required to accommodate partial client participation. As detailed in Section 3.2, constructing each new layer in DeepAFL entails two rounds of communication, corresponding to the computation of the classifier weights \mathbf{W}_t and the trainable transformation matrix Ω_{t+1} . For notational convenience, let \mathcal{S} denote the complete set of clients, where the total number of clients is $|\mathcal{S}| = K$. Consequently, the aggregation processes for \mathbf{W}_t (i.e., Equation (12)) and Ω_{t+1} (i.e., Equation (16)) can be expressed as:

$$\mathbf{G}_t^{1:K} = \mathbf{G}_t^{\mathcal{S}} = \sum_{k \in \mathcal{S}} \mathbf{G}_t^k, \quad \mathbf{H}_t^{1:K} = \mathbf{H}_t^{\mathcal{S}} = \sum_{k \in \mathcal{S}} \mathbf{H}_t^k, \quad (82)$$

$$\boldsymbol{\Pi}_t^{1:K} = \boldsymbol{\Pi}_t^{\mathcal{S}} = \sum_{k \in \mathcal{S}} \boldsymbol{\Pi}_t^k, \quad \boldsymbol{\Upsilon}_t^{1:K} = \boldsymbol{\Upsilon}_t^{\mathcal{S}} = \sum_{k \in \mathcal{S}} \boldsymbol{\Upsilon}_t^k. \quad (83)$$

When only a partial set of clients participates in the aggregation for DeepAFL, we denote the subset of clients contributing to the classifier construction at layer t as $\mathcal{S}_t^{\mathbf{W}}$, and the subset contributing to the trainable transformation matrix construction as \mathcal{S}_t^{Ω} . Consequently, the aggregated matrices will be constructed using contributions from only these participating subsets:

$$\mathbf{G}_t^{\mathcal{S}_t^{\mathbf{W}}} = \sum_{k \in \mathcal{S}_t^{\mathbf{W}}} \mathbf{G}_t^k, \quad \mathbf{H}_t^{\mathcal{S}_t^{\mathbf{W}}} = \sum_{k \in \mathcal{S}_t^{\mathbf{W}}} \mathbf{H}_t^k, \quad (84)$$

$$\boldsymbol{\Pi}_t^{\mathcal{S}_t^{\Omega}} = \sum_{k \in \mathcal{S}_t^{\Omega}} \boldsymbol{\Pi}_t^k, \quad \boldsymbol{\Upsilon}_t^{\mathcal{S}_t^{\Omega}} = \sum_{k \in \mathcal{S}_t^{\Omega}} \boldsymbol{\Upsilon}_t^k. \quad (85)$$

Subsequently, the server utilizes these aggregated matrices from the partial participants to compute \mathbf{W}_t and Ω_{t+1} as before. Notably, since only a subset of clients contributed data, the effective full dataset should be redefined as the union of data held by all participating clients. Thus, the invariance property of our DeepAFL needs to be redefined as being identical to the centralized analytical solution over the effective full dataset comprising the data of the participating clients. In summary, our DeepAFL can handle partial participation easily and effectively without substantial procedural adjustments, while maintaining its analytical advantage of being invariant to data heterogeneity.

Second, we conducted experiments to evaluate the performance of our DeepAFL under varying client participation rates η . For simplicity, we assume that the participation rate remains consistent across aggregation processes within the same layer, i.e., $\eta = |\mathcal{S}_t^{\mathbf{W}}|/|\mathcal{S}| = |\mathcal{S}_t^{\Omega}|/|\mathcal{S}|$. Furthermore, the client subsets $\mathcal{S}_t^{\mathbf{W}}$ and \mathcal{S}_t^{Ω} corresponding to η are randomly sampled from \mathcal{S} . Specifically, to ensure a comprehensive evaluation, we consider two distinct cases: (1) **Consistent Participation** and (2) **Inconsistent Participation**. These cases dictate whether the client subsets for the two aggregation processes within the same layer are identical or distinct, i.e., $\forall t, \mathcal{S}_t^{\mathbf{W}} = \mathcal{S}_t^{\Omega}$ or $\forall t, \mathcal{S}_t^{\mathbf{W}} \neq \mathcal{S}_t^{\Omega}$. Intuitively, the case (2) presents a greater challenge for DeepAFL, as the inconsistency within a single layer may heighten the risk of model instability.

Building upon these two cases, and using full participation ($\eta = 100\%$) as the control group, we conducted a broad range of analyses with η ranging from 90% to 50%, to thoroughly assess the performance variation of our DeepAFL on CIFAR-100. As shown in Table 28, our DeepAFL exhibits substantial robustness to partial client participation in both cases. Specifically, for a high participation rate of $\eta \geq 70\%$, the maximum accuracy degradation observed at $T = 20$ remains below 0.7% across both cases. Furthermore, even under the extremely challenging scenario of $\eta = 50\%$, where up to half of the clients drop out, DeepAFL still maintains high accuracy at $T = 20$, achieving 65.14% and 64.81% for both cases. Even at a low layer count ($T = 5$) and the most severe dropout rate ($\eta = 50\%$), the performance of our DeepAFL (61.54% and 61.36%) still significantly outperforms the AFL’s accuracy (58.56%) achieved under the condition of 100% client participation. This underscores the strong robustness of DeepAFL in handling partial client participation scenarios.

2376 N MORE ANALYSES ON IMPERFECT BACKBONE

2378 Here, we analyze the theoretical guarantees of our DeepAFL under the imperfect backbone, which
 2379 include our DeepAFL’s invariance to data heterogeneity and non-increasing empirical risk. Further-
 2380 more, we also investigate the impact of imperfect backbones on our DeepAFL’s performance. The
 2381 detailed analyses are presented below.

2382 First, let’s focus on analyzing the invariance of our DeepAFL under the imperfect backbone. Specif-
 2383 ically, as demonstrated in Theorem 1, the invariance of our DeepAFL essentially states that the
 2384 training result of our DeepAFL is identical to the centralized analytical solutions obtained on the
 2385 full dataset $\mathcal{D} = \{\mathbf{X}, \mathbf{Y}\}$. The backbone’s primary function is to obtain the local zero-layer feature
 2386 matrix $\Phi_0^k \in \mathbb{R}^{N_k \times d_\Phi}$ for each client k , given by:

$$2387 \Phi_0^k = \sigma(\text{Backbone}(\mathbf{X}^k, \Theta)\mathbf{A}). \quad (86)$$

2389 When the shared backbone is imperfect (e.g., noisy) and affects feature extraction for each client,
 2390 the same effect also applies to the centralized feature extraction, as follows:

$$2392 \Phi_0 = \sigma(\text{Backbone}(\mathbf{X}, \Theta)\mathbf{A}). \quad (87)$$

2393 Therefore, regardless of the backbone’s quality (even if it were an identity mapping), the core con-
 2394 clusion that the results obtained by our DeepAFL are identical to the centralized analytical solutions
 2395 remains unchanged. Thus, our DeepAFL’s invariance to data heterogeneity always holds.

2396 Second, we analyze the non-increasing empirical risk of our DeepAFL under the imperfect back-
 2397 bone. Specifically, as demonstrated in Theorems 2 and 3, the empirical risk of the analytic classifier
 2398 is monotonically non-increasing as the number of layers increases. This theoretical guarantee of our
 2399 DeepAFL stems from the convex optimization result provided by least squares for the objectives (4)
 2400 and (7) in our paper. Since the derived closed-form solution for our DeepAFL is always the optimal
 2401 one for objectives (4) and (7), the worst-case scenario is that the empirical risk equals that of the
 2402 previous layer (i.e., convergence), and it is impossible for the empirical risk of a given layer to be
 2403 greater than that of the previous layer. Thus, the non-increasing empirical risk of our DeepAFL
 2404 always holds no matter how the backbone changes.

2405 Third, we further analyze the impact of imperfect backbones on our DeepAFL’s performance.
 2406 Specifically, our DeepAFL does experience performance degradation when the backbone deterio-
 2407 rates, which stems from the reduced expressiveness of the features extracted by the imperfect back-
 2408 bone. A poor backbone will reduce the representational capacity of the extracted zero-layer features
 2409 Φ_0 . In particular, the mutual information between the features Φ_0 and the labels \mathbf{Y} may become
 2410 very low. Consequently, while our DeepAFL can still perform further optimization and enhance-
 2411 ment, its performance ceiling is ultimately limited by the low information content provided by the
 2412 weaker features Φ_0 . Furthermore, while the performance dependency on the quality of the backbone
 2413 is indeed a limitation of DeepAFL, it does not diminish the primary advantage of DeepAFL: grant-
 2414 ing the analytic classifier the capability of representation learning while preserving its gradient-free
 2415 manner. This allows it to move beyond merely learning linear mappings on the extracted features,
 2416 enabling the learning of more complex non-linear relationships. In fact, our experiments have al-
 2417 ready validated DeepAFL’s direct and effective improvement, where DeepAFL achieves superior
 2418 performance while reducing computational costs by more than 99% compared to gradient-based
 2419 methods, further highlighting its unique benefits.

2420 Furthermore, in fact, the analytic learning community has already initiated efforts to address the
 2421 dependency on frozen backbones (He et al., 2025a). We believe that integrating such techniques
 2422 with our DeepAFL to achieve even stronger performance represents a promising avenue for future
 2423 research. Nevertheless, this current attempt to unfreeze the backbone fundamentally adopts a mixed
 2424 gradient-based and gradient-free approach, which inevitably compromises the invariance property.
 2425 In contrast, the core philosophy of DeepAFL is distinctly different: we aim to provide direct and
 2426 efficient performance enhancements for AFL while strictly preserving its crucial theoretical advan-
 2427 tages in the FL setting. These advantages allow our DeepAFL to transcend mere linear mappings
 2428 on extracted features, enabling the learning of more complex non-linear relationships from features
 2429 with high efficiency. In summary, while the dependency on the quality of the backbone is indeed a
 limitation, it does not diminish the primary advantage of our DeepAFL, and addressing this limita-
 tion presents an interesting avenue for future work.

2430 O MORE DISCUSSIONS ON DEEPAFL’S ROBUSTNESS TO NOISES

2432 Table 29: Performance evaluation of DeepAFL on CIFAR-100 under noisy training environments.

2434 Layers	2435 $\tau = 0\%$	2436 $\tau = 10\%$	2437 $\tau = 20\%$	2438 $\tau = 30\%$	2439 $\tau = 40\%$	2440 $\tau = 50\%$
DeepAFL ($T = 5$)	64.72%	64.20%	63.63%	63.25%	62.51%	60.91%
DeepAFL ($T = 10$)	65.96%	65.43%	65.09%	63.96%	62.98%	60.99%
DeepAFL ($T = 15$)	66.56%	66.11%	65.66%	64.29%	63.15%	60.97%
DeepAFL ($T = 20$)	66.98%	66.38%	66.11%	64.69%	63.35%	60.45%

2441 Here, we further demonstrate that our DeepAFL, utilizing the MSE loss, is less sensitive to noise
 2442 compared to gradient-based methods employing the CE loss. Specifically, we elaborate on this point
 2443 from both logical and experimental perspectives, as detailed below.

2444 From the logical perspective, a crucial factor that makes a method sensitive to noise is overfitting.
 2445 Specifically, if a method tends to overfit, then when there is even slight noise in the input data, it
 2446 will tend to fit this noise. In contrast, the analytic learning methods using MSE loss are inherently
 2447 less susceptible to overfitting, typically exhibiting a tendency toward underfitting instead. Indeed,
 2448 addressing this underfitting limitation by augmenting the model’s fitting capacity through deep rep-
 2449 resentation learning is the primary motivation for proposing our DeepAFL.

2450 Furthermore, Tables 9–11 in Appendix E.1 list the training and testing accuracies for AFL and our
 2451 DeepAFL across various datasets, corroborating that analytic learning methods are not prone to
 2452 overfitting. Taking the CIFAR-100 dataset as an example, the training set accuracy for AFL is only
 2453 61.55%, which is quite low, especially compared to its testing set accuracy of 58.56%. This phe-
 2454 nomenon suggests that AFL suffers from severe underfitting. Moreover, when we gradually increase
 2455 the depth T from 0 to 50 in our DeepAFL, both training and testing accuracies rise steadily. The
 2456 fact that testing accuracy improves commensurately with training accuracy confirms that DeepAFL
 2457 avoids overfitting, even at depths of up to 50 layers.

2458 More specifically, DeepAFL’s ability to maintain superior representation learning without overfitting
 2459 stems from the fact that its analytic classifier with MSE loss is linear. Because of this, its model
 2460 complexity is quite low, making it very difficult to overfit and thus robust to noise. In contrast, CE
 2461 loss is designed to induce steeper gradients to facilitate backpropagation optimization. While this
 2462 characteristic accelerates training, it simultaneously renders gradient-based methods more prone to
 2463 overfitting and hypersensitive to data noise. Consequently, noise robustness constitutes the inherent
 2464 advantage of our DeepAFL (with MSE loss) compared to gradient-based methods (with CE loss).

2465 From the experimental perspective, we further verify the noise robustness of our DeepAFL by ex-
 2466 plicitly simulating noisy training environments. Specifically, we simulate the noise in the training
 2467 data by randomly flipping the labels of a proportion τ of the training data. The selected samples have
 2468 their labels randomly changed to another class label. Using $\tau = 0\%$ as the control group, we con-
 2469 duct extensive analyses with $\tau \in \{10\%, 20\%, 30\%, 40\%, 50\%\}$ to thoroughly assess performance
 2470 variations of our DeepAFL on CIFAR-100. The detailed results are reported in Table 29.

2471 These results demonstrate that DeepAFL exhibits substantial robustness to data noise. Specifically,
 2472 under moderate noise levels ($\tau \leq 20\%$), the maximum accuracy drop for our DeepAFL (at $T = 20$)
 2473 is less than 1%. Furthermore, even under the extremely challenging scenario of $\tau = 50\%$, where
 2474 half of the training labels are corrupted, DeepAFL maintains high accuracy. Crucially, even at a low
 2475 layer count ($T = 5$) and the most severe noise rate ($\tau = 50\%$), the performance of our DeepAFL
 2476 (i.e., 60.91%) still significantly outperforms the AFL’s accuracy (i.e., 58.56%) achieved under the
 2477 condition of 100% accurate data labels with no noise. Since AFL represents the state-of-the-art
 2478 baseline on Non-IID-1 for CIFAR-100, these results also imply that our DeepAFL, even with 50%
 2479 label corruption, outperforms all gradient-based baselines under ideal label accuracy.

2480 Additionally, as detailed in Appendix M, we also examine the robustness of DeepAFL under partial
 2481 client participation. Similar experiments conducted with client dropout rates ranging from 10% to
 2482 50% show that DeepAFL maintains highly stable and superior performance. Notably, DeepAFL’s
 2483 performance under label flipping is slightly weaker than under client dropouts. This is expected, as
 2484 client dropouts simply reduce training data volume, whereas label flipping introduces the additional
 2485 negative effect of data poisoning. Moreover, the performance gap between these two scenarios
 2486 remains small, which further substantiates DeepAFL’s superior robustness.

2484 **P MORE DISCUSSIONS ON REDUCING REDUNDANT REPRESENTATIONS**
 2485

2486 Table 30: Analyses of reducing representation redundancy within our DeepAFL on the CIFAR-100
 2487 dataset. “Train” and “Testing” denote the accuracy (%) on the training set and testing set. “Orig.”
 2488 and “Mod.” refer to the original and modified DeepAFL after reducing redundant representations.
 2489 The symbol Δ indicates the performance difference between the original and modified DeepAFL.

Layers	Train (Orig.)	Train (Mod.)	Δ_{Train}	Test (Orig.)	Test (Mod.)	Δ_{Test}
AFL	61.55%	61.55%	\	58.56%	58.56%	\
DeepAFL ($T = 0$)	65.71%	65.71%	0.00%	60.81%	60.81%	0.00%
DeepAFL ($T = 5$)	74.93%	77.00%	2.07%	64.62%	65.11%	0.49%
DeepAFL ($T = 10$)	79.41%	82.21%	2.80%	65.98%	66.66%	0.68%
DeepAFL ($T = 15$)	82.69%	85.65%	2.96%	66.59%	67.60%	1.01%
DeepAFL ($T = 20$)	85.15%	88.05%	2.90%	66.95%	68.02%	1.07%
DeepAFL ($T = 25$)	87.21%	89.94%	2.73%	67.40%	68.23%	0.83%
DeepAFL ($T = 30$)	88.84%	91.72%	2.88%	67.71%	68.60%	0.89%

2499 Here, we provide detailed discussions on reducing representation redundancy within our DeepAFL
 2500 to enhance its performance. Specifically, in the current design of DeepAFL, each new layer’s residual
 2501 block $g_{t+1}(\Phi_t)$ is built upon the feature from the immediately preceding layer Φ_t as shown in
 2502 Equations (6) and (9). Consequently, if the feature Φ_t already exhibits strong linear separability
 2503 or fitting capabilities, subsequent layers may indeed have little room to learn additional useful in-
 2504 formation, potentially leading to representational redundancy. To address this, we propose the idea
 2505 of adding a reconstruction loss during the feature optimization process to promote feature diversity.
 2506 Furthermore, we introduce a modified version of DeepAFL to verify whether reducing representa-
 2507 tion redundancy contributes to performance improvement. Detailed discussions are presented below.

2508 First, let’s focus on discussing the idea of adding a reconstruction loss to promote feature diversity.
 2509 Recall that our original feature optimization is guided by the sandwiched least squares problem in
 2510 Equation (7). In fact, we can augment this objective with a reconstruction-based term, transforming
 2511 it into a multi-objective optimization problem, as follows:

$$\Omega_{t+1} = \arg \min_{\Omega} \rho \|\mathbf{Y} - (\Phi_t + \mathbf{F}_t \Omega) \mathbf{W}_t\|_F^2 + (1 - \rho) \|\Phi_0 - (\Phi_t + \mathbf{F}_t \Omega) \mathbf{Q}_t\|_F^2 + \gamma \|\Omega\|_F^2. \quad (88)$$

2512 Here, the first term is the original empirical risk minimization, the second term is the reconstruction-
 2513 based objective (similar to an auto-encoder with MSE), and the third term is the regularization.
 2514 The parameter ρ serves as a weighting factor to balance the prediction loss and the reconstruction
 2515 loss. Specifically, the term $\Phi_{t+1} = (\Phi_t + \mathbf{F}_t \Omega)$ can be considered the encoded feature, and the
 2516 projection matrix \mathbf{Q}_t can be viewed as a decoder to map the feature back to the initial feature
 2517 Φ_0 . Theoretically, this approach would ensure that the newly constructed feature Φ_{t+1} at each
 2518 layer contains information necessary to reconstruct the initial feature Φ_0 . Nevertheless, deriving
 2519 the closed-form solution for this new and coupled objective is significantly more complex than for
 2520 our original sandwiched least squares problem. Consequently, we confine our discussion to the
 2521 high-level idea and omit a formal derivation for brevity, positioning it as a direction for future work.

2522 Second, we further introduce a simple modification to DeepAFL to verify whether reducing rep-
 2523 resentation redundancy indeed helps boost performance. Specifically, we can readily extend DeepAFL
 2524 by directly altering the input of the residual block $g_{t+1}(\cdot)$. For instance, the simplest modification is
 2525 to decouple the residual block at each layer from the immediately preceding features Φ_t and instead
 2526 link it to the initial feature Φ_0 . Since the initial feature Φ_0 has not yet been modified or adjusted by
 2527 DeepAFL, it is expected to retain a richer set of information from the original input. The modified
 2528 formula is described as follows:

$$\Phi_{t+1} = \Phi_t + g_{t+1}(\Phi_0), \quad g_{t+1}(\Phi_0) = \sigma(\Phi_0 \mathbf{B}_t) \Omega_{t+1}. \quad (89)$$

2529 Based on this modification, we experiment with the modified DeepAFL and report the results in Ta-
 2530 ble 30. The results show that leveraging the initial features Φ_0 consistently improves both training
 2531 and testing accuracy, yielding gains of over 2% and approximately 0.5% to 1%, respectively. Cru-
 2532 cially, at $T = 30$, the test accuracy of the modified version is already comparable to the performance
 2533 of the original version at $T = 50$, effectively saving the computational cost of 20 layers. Thus, we
 2534 confirm that reducing representational redundancy does enhance DeepAFL’s performance. Despite
 2535 potential for redundant representations, we retain the current version of DeepAFL in the main text
 2536 based on considerations of clarity and accessibility for readers.

2538 **Q MORE DISCUSSIONS ON DEEP AFL’S USE CASES**
2539

2540 Here, we provide detailed discussions on DeepAFL’s use cases. Specifically, our DeepAFL requires
 2541 a given backbone to serve as an effective feature extractor for training, but it does not mandate how
 2542 this backbone is obtained. Indeed, leveraging pretrained models or foundation models as the back-
 2543 bone in FL has emerged as a very common and popular research practice in recent years (Nguyen
 2544 et al., 2023; Piao et al., 2024; Yu et al., 2024; 2025). Accordingly, we elaborate below on several
 2545 common ways for obtaining such pretrained backbones, corresponding to different use cases of our
 2546 DeepAFL, as well as the pervasiveness of pre-trained backbones in FL and the necessity of FL even
 2547 with such backbones. The detailed discussions are presented as follows:

2548 First, we want to highlight that our DeepAFL requires a pretrained backbone for feature extraction,
 2549 but it does not mandate how this backbone is obtained. For the sake of brevity, we do not discuss
 2550 this in detail within the main text. Herein, we introduce three common ways for obtaining pretrained
 2551 backbones, which collectively illustrate DeepAFL’s various use cases:

- 2552 • **Supervised Pre-training:** This is the most prevalent case and constitutes the primary set-
 2553 ting for the experimental comparisons presented in our paper. In this scenario, our Deep-
 2554 AFL can be viewed as a collaborative fine-tuning approach for FL, where clients adapt a
 2555 shared public backbone to private data with different feature distributions.
- 2556 • **Self-Supervised Pre-training:** With the emergence of large-scale pre-trained models,
 2557 there is frequently insufficient labeled data for supervised pre-training. In such cases, self-
 2558 supervised backbones have gained prominence (such as auto-encoders via reconstruction,
 2559 and DINO series via contrastive learning). These self-supervised models, however, inher-
 2560 ently lack an integrated classifier or regressor, possessing only robust generalizable feature
 2561 extraction capabilities without direct predictive functionality. Therefore, deploying such
 2562 backbones in FL necessitates training a matching classifier or regressor. In this context,
 2563 our DeepAFL serves as a direct and highly efficient FL training approach to obtain this
 2564 classifier or regressor, effectively leveraging the backbone’s established capabilities.
- 2565 • **Domain Adaptation:** If the pre-trained backbone presents a domain shift relative to the
 2566 FL system’s target data, the server can also perform initial feature domain adaptation by
 2567 fine-tuning the backbone (using a set of domain-specific or compliant public data). Once
 2568 this adaptation is complete, the backbone is then frozen, and our DeepAFL can be imple-
 2569 mented for subsequent FL training. This case can be viewed as an engineering strategy to
 2570 mitigate the issues arising from a cross-domain backbone. Given that domain adaptation
 2571 can be performed centrally by the server and DeepAFL incurs only negligible time costs,
 2572 the aggregate overhead across both stages is still considered to be very small.

2573 In summary, DeepAFL can not only leverage backbones obtained through supervised pre-training,
 2574 as demonstrated in our experiments, but also effectively utilize other types of backbones for differ-
 2575 ent use cases. Specifically, when using backbones from self-supervised pre-training, our DeepAFL
 2576 effectively solves the meaningful problem of constructing task-specific classifier in FL. Further-
 2577 more, when faced with the potential for domain gaps, this problem can be mitigated through domain
 2578 adaptation as described above, which can be viewed as an engineering strategy.

2579 Second, we further discuss the pervasiveness of pre-trained backbones in FL and the necessity of
 2580 FL even with such backbones. Indeed, incorporating pre-trained backbones is a widely embraced
 2581 practice in recent research to introduce prior knowledge and stabilize FL (Nguyen et al., 2023; Piao
 2582 et al., 2024; Yu et al., 2024; 2025). Specifically, a key commonality between our work and these
 2583 studies is the focus on using foundation models for conducting FL, rather than using FL for building
 2584 foundation models. Notably, many existing pre-trained model-based FL works are mechanistically
 2585 constrained to necessitate large foundation models. In contrast, DeepAFL accommodates backbones
 2586 of varying sizes and capabilities, which renders it particularly suitable for resource-constrained edge
 2587 environments compared to related works. Moreover, we demonstrate that a pre-trained backbone
 2588 does not negate the need for FL. Leveraging results from the original AFL paper on CIFAR-100
 2589 ($\alpha = 0.1$ and $K = 100$) (He et al., 2025b), we observe that purely local training without global ag-
 2590 gregation yields maximum and average test accuracies of merely 16.36% and 12.04%. Conversely,
 2591 traditional FL method FedAvg and AFL achieve 56.62% and 58.56%, while our DeepAFL ($T = 20$)
 2592 further elevates performance to 66.98%. These results conclusively demonstrate that despite the
 2593 availability of a pre-trained backbone, FL remains indispensable for achieving high performance.

2592 R MORE DISCUSSIONS ON DEEPAFL’S REPRESENTATION LEARNING

2593
 2594 In this section, we present detailed discussions on the representation learning capabilities of Deep-
 2595 AFL. Specifically, we first outline how the capacity of DeepAFL aligns with fundamental character-
 2596 istics of representation learning. We then analyze the intermediate features produced by DeepAFL
 2597 as empirical evidence supporting these capabilities. The specific discussions are provided below.
 2598

2599 R.1 ANALYSES ON THE CONCEPT OF REPRESENTATION LEARNING

2600 Here, we demonstrate that the capabilities of DeepAFL align with the fundamental characteristics
 2601 of representation learning. To this end, we first outline these fundamental characteristics and then
 2602 conduct a point-by-point analysis showing how our DeepAFL satisfies each one. Subsequently, we
 2603 further substantiate our claim by analyzing the architectural similarities between DeepAFL and the
 2604 Multi-Layer Perceptron (MLP). The detailed analyses are presented below.
 2605

2606 First of all, drawing upon the seminal work of Bengio et al. (Bengio et al., 2013), we can concep-
 2607 tually distill two fundamental characteristics of representation learning: **(1) Automated Feature**
 2608 **Extraction:** This entails learning feature transformations (which are often non-linear) directly from
 2609 the data, thereby avoiding the complexity and reliance on expert knowledge associated with man-
 2610 ual feature engineering. **(2) Utility for Downstream Tasks:** The extracted features are beneficial
 2611 and useful for adapting to specific downstream tasks, meaning they provide utility when building
 2612 predictors (e.g., classifiers), ultimately leading to enhanced model performance.
 2613

2614 In fact, our DeepAFL directly satisfies these two foundational characteristics: (1) The feature trans-
 2615 formations involve trainable parameters Ω , which are automatically learned from the data via the
 2616 “sandwiched” least squares. Moreover, the use of the activation function $\sigma(\cdot)$ ensures necessary
 2617 non-linearity in the layer-wise transformations. (2) The feature transformations show evident and
 2618 significant benefits for our downstream task (i.e., classification). Specifically, when coupled with
 2619 an analytic classifier, the accuracy of our DeepAFL consistently improves on the training and test
 2620 sets as T increases (as shown in Tables 9–11 of our paper). In summary, based on these conceptual
 2621 analyses, our DeepAFL indeed aligns with the fundamental concept of representation learning.
 2622

2623 Subsequently, to further clarify our DeepAFL’s alignment with fundamental representation learning
 2624 concepts, we further draw an analogy between our DeepAFL and the MLP. Given that the MLP is
 2625 recognized as the most fundamental DNN, which indisputably achieves representation learning, this
 2626 analogy serves to demonstrate that our DeepAFL similarly achieves these capabilities.
 2627

2628 Specifically, since our DeepAFL incorporates a residual block structure, we also introduce skip
 2629 connections into the MLP, which does not impair its representation learning capability (but rather
 2630 makes it easier to train). At this point, what the residual block of the MLP learns is nothing more
 2631 than subjecting the features from the previous layer to an affine transformation \mathbf{W}_{t+1} followed by an
 2632 activation function $\sigma(\cdot)$. Here, the affine transformation \mathbf{W}_{t+1} constitutes the learnable parameters
 2633 of the MLP, as follows:

$$g_{t+1}(\Phi_t) = \sigma(\Phi_t \mathbf{W}_{t+1}). \quad (90)$$

2634 Consequently, when multiple layers are stacked, the MLP functions as a continuous alternation
 2635 of affine transformations and non-linear activation functions. Crucially, a clear parallel emerges
 2636 between our DeepAFL and the MLP, as DeepAFL similarly involves an alternating stack of affine
 2637 transformations and non-linear activation functions:

$$g_{t+1}(\Phi_t) = \sigma(\Phi_t \mathbf{B}_t) \Omega_{t+1}. \quad (91)$$

2638 The primary distinction is that the learnable parameters \mathbf{W}_{t+1} in the MLP are applied prior to the
 2639 activation function, while the parameters Ω_{t+1} in our DeepAFL are applied after the activation. This
 2640 localized difference is effectively diminished when multiple layers are stacked. Furthermore, more
 2641 complex DNNs also share this philosophy of feature re-weighting via affine transformations and
 2642 non-linear activation functions. For instance, a CNN can be broadly viewed as an MLP with weight
 2643 sharing and local connectivity. Therefore, the inherent structural similarity between DeepAFL and
 2644 the MLP confirms that DeepAFL is aligned with the core concepts of representation learning.
 2645

Based on the analyses above, our DeepAFL not only satisfies the fundamental characteristics of rep-
 2646 resentation learning but also exhibits a strong structural similarity to the MLP. This dual validation
 2647 confirms that DeepAFL aligns with the core concepts of representation learning.

2646 R.2 ANALYSES ON INTERMEDIATE FEATURES
26472648 Table 31: Separability analyses of DeepAFL’s intermediate features on CIFAR-100.
2649

Layer	Training CSR	Testing CSR	Training IFS	Testing IFS	Training DM	Testing DM
DeepAFL ($T = 0$)	1.583	1.529	3.25	3.13	3.21	3.09
DeepAFL ($T = 5$)	1.565	1.520	3.20	3.11	3.16	3.07
DeepAFL ($T = 10$)	1.555	1.516	3.17	3.10	3.14	3.06
DeepAFL ($T = 15$)	1.546	1.513	3.15	3.09	3.12	3.05
DeepAFL ($T = 20$)	1.539	1.510	3.13	3.08	3.11	3.04

2655 Here, we present experimental analyses of the intermediate features generated by DeepAFL to empirically validate its representation learning capabilities. Specifically, we first conduct additional experiments designed to assess the separability of the intermediate features produced by our DeepAFL, thereby substantiating the model’s advanced learning capacity. Furthermore, we analyze DeepAFL’s utility for the downstream task to further validate its representation learning capabilities.

2661 First, let’s focus on the additional experiments for investigating the separability of intermediate features produced by our DeepAFL, which are conducted using the following three metrics:

- **Compactness-Separation- Ratio (CSR):** This metric evaluates the ratio of the average intra-class squared distance (i.e., compactness) to the average inter-class squared distance (i.e., separation). A lower CSR value indicates better separability.
- **Inverse Fisher Score (IFS):** It is defined as the reciprocal of the Fisher Score, which is a widely adopted metric to distinguish different classes. The IFS is calculated by the ratio of the intra-class variance to inter-class variance. A lower FS value implies better separability.
- **Discriminative Measure (DM):** It is a metric derived from the core principles of the Linear Discriminant Analysis (LDA), quantifying the ratio of within-class scatter to between-class scatter. A lower DM value signifies better separability.

2673 In summary, all the additional metrics are unified such that a lower value indicates better representation separability. Using these metrics, we conducted additional experiments on the CIFAR-100 dataset to analyze the intermediate features generated by our DeepAFL. The detailed experimental results are presented in Table 31. From the results, it is evident that all these metrics show better separability as the layer count of our DeepAFL increases on both the training and test sets, which further substantiates our DeepAFL’s representation learning capabilities. Furthermore, we observe that the metrics on the test set generally perform better (i.e., yield lower values) than those on the training set. This is likely because the training set contains a much more samples, leading to greater inherent noise and, consequently, slightly inflated metric values.

2682 Meanwhile, we can also observe that DeepAFL exhibits a diminishing marginal returns phenomenon as the layer count increases, but this is entirely expected and normal since we did not increase any training data volume. This phenomenon is also common in gradient-based methods. For instance, simply increasing the depth of a DNN can not lead to continuous accuracy improvement up to 100% without encountering a bottleneck. Even the high performance achieved by LLMs is a result of scaling the model alongside vast increases in training data volume (i.e., the scaling law).

2688 Second, as discussed in Appendix R.1, the utility for the downstream task (i.e., classification accuracy) is the most direct and fundamental metric to validate the representation learning capabilities of DeepAFL. This perspective also aligns with the classical work by Yoshua Bengio et al. (Alain & Bengio, 2016), which asserts that better linear separability is indicative of a more meaningful representation. In fact, in Tables 9–11, we have reported the training and testing accuracies achieved by our DeepAFL when utilizing intermediate features from various depths (i.e., DeepAFL with different numbers of layers). Let’s take the results on the CIFAR-100 dataset as an example. The baseline AFL struggles with severe underfitting, achieving a training accuracy of only 61.55%. In stark contrast, our DeepAFL (at $T = 20$) propels the training accuracy to 85.15%, representing a massive absolute improvement of 23.6%. This improved representation translates directly to generalization, yielding a substantial 8.39% increase in test accuracy. Crucially, these gains are achieved with high efficiency. The total runtime for our DeepAFL is only 91.74 s, which is a marginal increase of less than 42 s compared to AFL (50.05 s). These results show that DeepAFL learns significantly more discriminative and useful representations than AFL.