

Handling Spatial-Temporal Data Heterogeneity for Federated Continual Learning via Tail Anchor

Hao Yu
Southwestern University of
Finance and Economics
yuhao2033@163.com

Hanlin Gu
WeBank
allengu@webank.com

Xin Yang*
Southwestern University of
Finance and Economics
yangxin@swufe.edu.cn

Tianrui Li
Southwest Jiaotong University
trli@swjtu.edu.cn

Qiang Yang
Hong Kong University of
Science and Technology
qyang@cse.ust.hk

Le Zhang
University of Electronic Science
and Technology of China
lezhang@uestc.edu.cn

Lixin Fan
WeBank
lixinfan@webank.com

Abstract

*Federated continual learning (FCL) allows each client to continually update its knowledge from task streams, enhancing the applicability of federated learning in real-world scenarios. However, FCL needs to address not only spatial data heterogeneity between clients but also temporal data heterogeneity between tasks. In this paper, empirical experiments demonstrate that such input-level heterogeneity significantly affects the model's internal parameters and outputs, leading to severe spatial-temporal catastrophic forgetting of local and previous knowledge. To this end, we propose Federated Tail Anchor (FedTA) to mix trainable **Tail Anchor** with the frozen output features to adjust their position in the feature space, thereby overcoming parameter-forgetting and output-forgetting. Three novel components are also included: **Input Enhancement** for improving the performance of pre-trained models on downstream tasks; **Selective Input Knowledge Fusion** for fusion of heterogeneous local knowledge on the server; and **Best Global Prototype Selection** for finding the best anchor point for each class in the feature space. Extensive experiments demonstrate that FedTA not only outperforms existing FCL methods but also effectively preserves the relative positions of features. Code is available at: https://github.com/SkyOfBeginning/FedTA_CVPR2025.*

*Xin Yang is the corresponding author.

1. Introduction

Data heterogeneity across different clients (Non-IID) is one of the most important challenges in traditional Federated Learning (FL), which greatly hinders the integration of knowledge, leading to the aggregated global model underperforming on local tasks. Many studies have attempted to address this issue and have made some progress [6, 37, 45]. However, they are based on an unrealistic static assumption that the training data of all clients will remain unchanged. Federated Continual Learning (FCL) breaks the static limits by allowing clients to continually accumulate knowledge from task sequences [1, 7, 19, 38]. While FCL expands the applicability of FL in real-world scenarios, it also introduces a more challenging issue, i.e., *spatial-temporal data heterogeneity*. Not only is the data heterogeneous across different clients (spatial), but also the data within different tasks of the same client is heterogeneous (temporal), as shown on the left side of Fig. 1.

The most direct negative impact of spatial-temporal data heterogeneity is *spatial-temporal catastrophic forgetting* (ST-CF) [36], as illustrated on the middle side of Fig. 1. Catastrophic Forgetting (CF) is a term from Continual Learning (CL), used to describe the phenomenon where a deep model, after learning multiple tasks, tends to forget the knowledge of previous tasks, resulting in a decrease in accuracy [8, 13, 23]. In FCL, clients face *temporal catastrophic forgetting* as local models would continually learn different tasks over time. Additionally, Non-IID data would lead the aggregated global model to *spatial catastrophic forgetting*

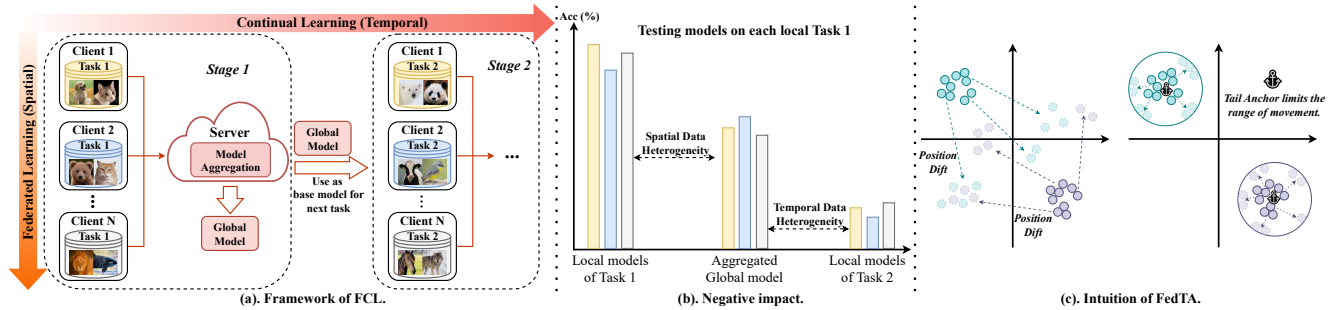


Figure 1: Illustration of FCL, the negative impact of spatial-temporal data heterogeneity and the intuition of Tail Anchor.

(i.e., a decline in the performance of the global model on local test sets). Moreover, *spatial forgetting will interact with temporal forgetting*, as clients use the global model as the base model to learn the next task [36].

Spatial-temporal data heterogeneity, which manifests as differences at the *model input*, leads to corresponding changes in *model parameters and outputs*. We believe this is the fundamental cause of forgetting. To be specific, as data changes over time and space, feature extractor and classification head of the model will adapt to the most recent inputs, leading to forgetting of previous and client-specific knowledge (a detailed analysis of the effect of forgetting is presented in Fig. 1(a-b) and Sec. 3.2). Fortunately, the use of pre-trained large models can effectively mitigate catastrophic forgetting at the parameter level, as they have sufficient capacity to extract features without changing internal parameters [5, 30, 32]. However, frozen pre-trained models often perform poorly in downstream tasks, making them unsuitable for direct application [11]. Additionally, they cannot handle forgetting at the output since the classification head is trainable and will fit to the most recent task data.

To fully leverage the power of pre-trained models and address spatial-temporal forgetting from both the parameter and the output aspects, we first define Tail Anchors and mix them with frozen output features to fix the position, as shown in Fig. 1(c). Based on this concept, we propose **Federated Tail Anchor (FedTA)**, which leverages the tail anchor to keep the positions of each class invariant.

Firstly, each client shares a pre-trained Vision Transformer (ViT) [4] and a two-stage training strategy is designed to enhance the performance of pre-trained models at the input level and alleviate forgetting at the output level. By adding *tail anchors* to the output features, the features of samples that experience forgetting can quickly return to their original positions in the feature space, thereby avoiding forgetting caused by spatial-temporal changes. After completing local training, the server will separately process the parameters added during input enhancement by the client and the local prototype. On the one hand, we design a *selective input knowledge fusion* mechanism to selectively integrate the knowledge used for input enhancement from

different clients; on the other hand, the server will calculate the similarity between local prototypes to form a similarity adjacency matrix. In each iteration, the server will select the local prototype with the lowest average similarity for each class as the global prototype. If the average similarity falls below a threshold, the global prototype will be fixed to prevent forgetting.

Currently, there is little research on spatial-temporal heterogeneity in FCL [6]. To our knowledge, we are the first to attempt to address both temporal and spatial data heterogeneity from the perspective of forgetting. The main contributions can be summarized as:

1. Empirical experiments are conducted to show that spatial-temporal data heterogeneity can cause significant changes in the important features extracted by the model for the same samples, and it also causes shifts in their positions within the feature space. This leads to severe ST-CF of previous knowledge and local knowledge. We refer to these two phenomena as “parameter-forgetting” and “output-forgetting”, respectively.
2. FedTA leverages a pre-trained ViT along with four novel components, aiming to prevent features from shifting their positions in the feature space due to spatial-temporal data heterogeneity. FedTA efficiently overcomes both parameter-forgetting and output-forgetting in FCL caused by spatial-temporal data heterogeneity.
3. Extensive experiments not only demonstrate the state-of-the-art performance of FedTA but also show its remarkable ability to resist spatial-temporal forgetting. Moreover, visualization results indicate that FedTA effectively preserves the relative positions of features, preventing positional shifts due to spatial-temporal variations.

2. Related Work

Spatial data heterogeneity, as known as the Non-IID problem, has attracted much attention [2, 12, 28, 43]. Existing methods tackle data heterogeneity by either incorporating more effective local training or devising more comprehensive aggregation mechanisms [16].

Although these studies have made progress in overcom-

ing spatial data heterogeneity, they are unable to cope with more realistic and dynamic scenarios where each client continually learns on their own task stream.

FCL has indeed greatly enhanced the practical value of FL in real-world scenarios, especially on the edge computing side [35, 42, 44]. It allows each client to rapidly learn knowledge from the current task without forgetting previously knowledge, thus avoiding the need to retrain from scratch and greatly saving computational resources.

In a survey paper on FCL, the authors identified a key issue that existing FCL articles have overlooked: the interaction between spatial heterogeneity and temporal heterogeneity, which leads to a unique challenge: spatial-temporal catastrophic forgetting (ST-CF) [36]. It means that models not only forget previous knowledge due to continual learning but also forget local knowledge due to federated aggregation. Existing FCL methods do not realize that spatial heterogeneity can exacerbate the temporal forgetting, so when the spatial heterogeneity becomes stronger, the performance is not as expected. Besides, effective FCL methods currently heavily rely on replaying or generating pseudo data to mitigate the effects of spatial-temporal data heterogeneity [1, 18, 20, 22, 24, 41]. However, this may pose certain privacy risks and incur high computational cost.

Only a very small portion of work has attempted to address data heterogeneity from time and space simultaneously now [15, 39]. However, none of them have delved into how this heterogeneity leads to forgetting, nor have they ensured sufficiently strong spatial-temporal heterogeneity in their experimental settings. To our best knowledge, we are the first to deeply analyze how the heterogeneity of inputs affects model parameters and outputs.

3. Spatial-Temporal Data Heterogeneity

3.1. Problem Definition

The purpose of Spatial-Temporal Data Heterogeneity is to continually integrate knowledge from different clients and different time periods. We extend the traditional FL to FCL with strong spatial-temporal data heterogeneity.

For spatial heterogeneity, given a clients (denoted as $\mathcal{A} = \{A_1, A_2, \dots, A_a\}$), and a central server S , each client's data is composed of private classes C_v^i and public classes C_p , where private classes refer to the class of data that can only be seen by the client itself. We ensure that the data of C_p is non-overlapping between clients. Further, we can set $|C_p| = 0$ to ensure extreme spatial heterogeneity. We run experiments with $|C_p| = 0$ on Imagenet-R dataset.

For temporal heterogeneity, the task sequence of client A_i is denoted as $\mathcal{T}_i = \{T_i^1, T_i^2, \dots, T_i^{n_i}\}$, where n_i represents the total number of tasks on client A_i . Each task consists of the same number but entirely different classes.

During the training of task r , the global model on the

server already possesses the knowledge of T_i^1 to T_i^{r-1} from client $\{A_i, 1 \leq i \leq a\}$. The server S then distributes it back to clients. After personalizing the received global model θ_g^{r-1} , the client A_i continually trains it on T_i^r as the initial model to get the new local model θ_i^r . The local model θ_i^r should perform well in classifying classes of $\{\mathcal{T}_1^1 \cup \mathcal{T}_1^{r-1}, \dots, \mathcal{T}_i^1 \cup \mathcal{T}_i^r \cup \dots, \mathcal{T}_a^1 \cup \mathcal{T}_a^{r-1}\}$.

3.2. Negative Impact

Spatial-temporal data heterogeneity is a type of heterogeneity in model inputs. Due to the back-propagation mechanism, it would significantly affect the internal parameters of the model and the outputs [29]. It not only introduces differences between models of different clients but also causes the features output by the same sample to undergo significant changes, thereby causing spatial-temporal forgetting of previous knowledge and local knowledge.

For the feature extractor, as shown on *the left side* of Fig. 2, traditional centralized single-task training can accurately extract beneficial features. However, after continual learning of four other tasks, for the same image, the extracted features are completely unrelated to the cat. Similarly, for FL, the aggregated global model also fails to extract features related to the cat itself. More critically, when the model encounters spatial-temporal data heterogeneity, the crucial features extracted by the feature extraction layer are predominantly concentrated at the edges of the images, which is highly negative to classification task. We term this phenomenon as **parameter-forgetting**.

For the output (feature space), as shown on *the right side* of Fig. 2, the features extracted by the initial model have clear classification boundaries, and the features of each class are relatively concentrated in the same area. However, after CL or FL, the features extracted for the same samples no longer possess clear boundaries, especially in FCL. Moreover, the positions of the features gradually deviate from their original locations with the spatial-temporal variation, leading to the forgetting of old samples. We term this phenomenon as **output-forgetting**.

Let's delve even further into the effects of spatial-temporal data heterogeneity. For deep neural networks, changes at the input level directly affect model parameters and corresponding outputs, *thereby causing continual variations of the feature space*. For spatial data heterogeneity, the absence of a common feature space among clients makes it challenging to share heterogeneous knowledge. For temporal data heterogeneity, changes in the feature space over time lead to variations in the locations of features of the same samples. If we can address the issues mentioned above simultaneously, then spatial-temporal catastrophic forgetting will be resolved.

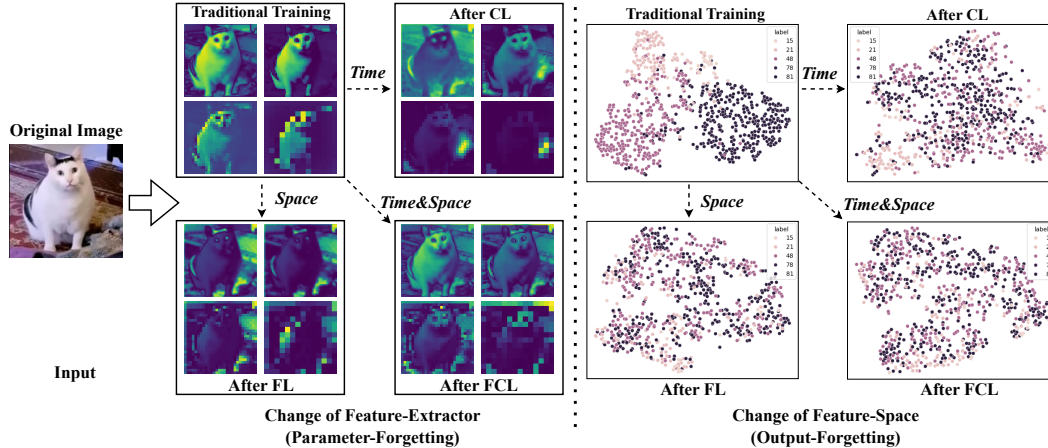


Figure 2: Illustrations of the negative impact of spatial-temporal data heterogeneity on the feature extractor and feature space. [Left Side] illustrates the variation of significant features extracted by the feature extractor for the same input sample, where brighter colors indicate more important features. As spatial-temporal changes occur, the extracted features gradually shift away from “cat”, even extracting features near the image edges. [Right Side] depicts the changes in the positions of the features in the feature space after undergoing spatial-temporal transformations for the same batch.

3.3. Motivation

It is evident that spatial-temporal data heterogeneity leads to both parameter-forgetting and output-forgetting. Therefore, methods need to possess the following three capabilities: (1) Ensure that the model extracts nearly identical features for the same sample; (2) Fix the positions of extracted features. (3) Allow clients to have a common feature space to better utilize heterogeneous knowledge.

However, due to the training method of deep networks and the large number of parameters involved, parameter updates are uncontrollable, *making it nearly impossible to mitigate parameter-forgetting*. Similarly, since ensuring consistency within the parameters is impossible, it is also hard to guarantee the invariance of outputs in the feature space.

Pre-trained large models have attracted attention due to their powerful representation capabilities. There are already articles attempting to apply pre-trained ViT to overcome forgetting in CL [21, 27]. Inspired by this, we find that freezing the feature extractor of pre-trained ViT can effectively **eliminate parameter-forgetting**. In FCL, clients share the same pre-trained model, ensuring that they have **the same knowledge/feature space**, which makes knowledge transfer between clients easier. Furthermore, by mixing learnable parameters (referred to “tail anchor” in this paper) with frozen features, we can effectively control their positions in the feature space, thus **addressing output-forgetting**. The server selects anchor points with the lowest similarity to other classes’ anchor points as the global anchor points in the feature space. During client-side training, the tail anchor converges towards these class-specific anchor points. Therefore, we mitigate the performance degradation caused by parameter-forgetting and output-forgetting.

4. Methodology: FedTA

To address the spatial-temporal catastrophic forgetting, which includes parameter-forgetting in the feature extractor and output-forgetting in the feature space, we propose FedTA. Its aim is to leverage a frozen pre-trained model and cross-mix learnable parameters after the output features, ensuring that the position of features in the feature space remains fixed and unaffected by spatial-temporal changes.

The overall framework of FedTA is shown in Fig. 3. and the algorithm is summarized in Appendix A.

4.1. Input Enhancement

We assign each client with the same pre-trained ViT model as a foundation model. With its parameters frozen, clients learn common knowledge that operate at the input level. The purpose is to extract knowledge into a common space through the same model and enhance ViT’s performance.

Knowledge Base. We devise a knowledge base for storing and selecting the input enhancement parameters. The knowledge base of client i is defined as

$$\mathcal{KB}_i = \{IE_1, IE_2, \dots, IE_M\}, \quad (1)$$

where M is the base size and IE is a set of learnable parameters. Then, let x and $E = f_e(x)$ be the input and its corresponding embedding feature, respectively. $f_e(\cdot)$ refers to the embedding function of ViT. Denoting $\{s_i\}_1^N$ be the indices of selected N sets, then we can modify the embedding feature as follows:

$$E' = [IE_{s_1}, \dots, IE_{s_N}; E], 1 \leq N \leq M, \quad (2)$$

where $[:]$ represents concatenation along the token length dimension. Each IE has a corresponding key, denoted as

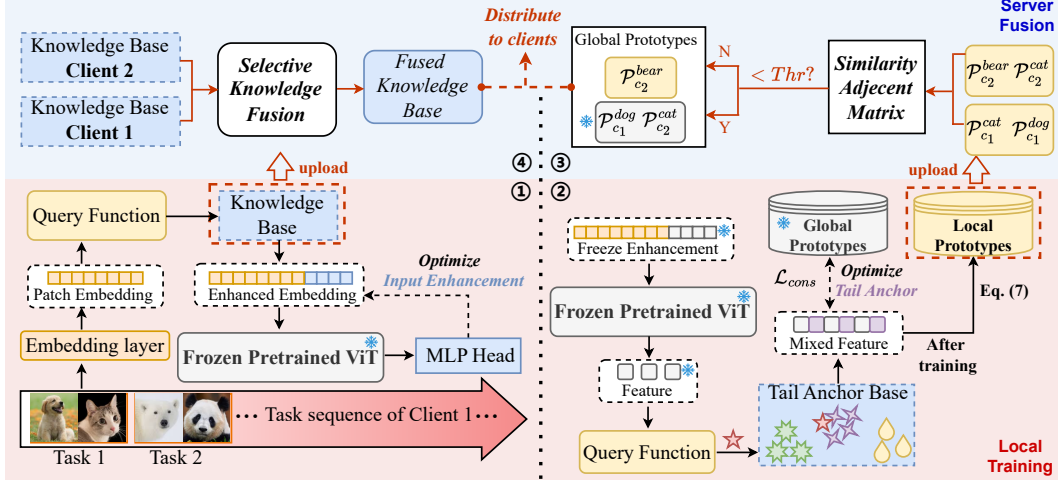


Figure 3: An overview of FedTA. Local training is a two-stage training process. The first stage involves adding *input enhancement* to the image embeddings to fully utilize ViT (see 1). In the second stage, the extracted features are fixed, and the corresponding *tail anchor* is mixed with them to adjust the similarity between classes by applying contrastive learning with global prototypes (see 2). Then, the local knowledge base of input enhancements and the local prototypes of each class are uploaded to the server, where *selective input knowledge fusion* for the knowledge base (see 4) and *global best prototype selection* for the local prototypes (see 3) are performed, respectively.

K^{ie} , to facilitate the selection of the IE based on the similarity of keys.

Optimization for the input enhancement. Each client has a classification head used for training input enhancement parameters, denoted as H_e^i . At the beginning of training, it is necessary to load the pre-trained model with H_e^i to enable it to perform the classification task, and we denote the model with H_e^i as \mathcal{V}_e^i . Overall, the training loss function is as follows:

$$\min \mathcal{L}(\mathcal{V}_e^i(E'), y) + \lambda_1 \sum_{K_s^{ie}} \text{dis}(K_{in}^{ie}, K_{s_i}^{ie}), \quad (3)$$

where λ_1 is a hyperparameter, K_s^{ie} and K_{in}^{ie} are used to find the best IE . The initial term comprises the softmax cross-entropy loss to optimize selected IE s, while the subsequent term serves as a surrogate loss aimed at bringing selected keys closer to their corresponding query features. Cosine similarity is used as the distance function.

In a nutshell, IE is a set of learnable parameters that can be concatenated to the image embedding E to enhance the performance of ViT. Each IE has a corresponding key K^{ie} . E is first processed by ViT to obtain features, which are then used as a query key K_{in}^{ie} to calculate the similarity with the key of each IE . Then best IE_s is selected to form an **enhance embedding** E' .

4.2. Tail Anchor

Query function. Once the input enhancement parameters are well trained, they will be frozen, including their corresponding keys, until the next task training. The enhanced

input embedding E' would be processed by the frozen ViT [4] again to get the features, denoted as \mathcal{F}_{out} . Then it will be used as the key to find the corresponding tail anchor based on the cosine similarity.

Tail Anchor is defined as key-value pairs for m classes: $\mathcal{TA} = \{(K_1^{ta}, TA_1), (K_2^{ta}, TA_2), \dots, (K_m^{ta}, TA_m)\}$. Specifically, K_s^{ta} , $s \in [m]$ is obtained as:

$$K_s^{ta} = \underset{\mathcal{K}^{ta}}{\text{argmin}} \text{dis}(\mathcal{F}_{out}, K_i^{ta}), \quad (4)$$

where K_s^{ta} denotes the chosen tail anchor's key, and \mathcal{K}^{ta} represents the set of keys for all tail anchors.

In a nutshell, the tail anchor of each class acts as additional parameters to manage the distance to a fixed position (global prototype of each class) in the feature space. Its main advantage is its relative position, which remains consistent regardless of changes in space and time. This consistency ensures that the tail anchor can effectively guide the positioning of general output features.

Optimization for the tail anchor. Once the tail anchor is chosen, it will be mixed with \mathcal{F}_{out} to form a new feature \mathcal{F}_{TA} . If a client has global prototypes (i.e., not the first round), then contrastive learning is utilized to unify the features across clients through the following unified representation loss function:

$$\mathcal{L}_{cons}(\mathcal{F}_{TA}) = -\log \frac{\exp(\mathcal{F}_{TA} \cdot \mathcal{G}^y / \tau)}{\sum_{y_a \in \mathcal{Y}^t} \exp(\mathcal{F}_{TA} \cdot \mathcal{G}^{y_a} / \tau)}, \quad (5)$$

where \mathcal{Y}^t represents the global available classes up to task t and \mathcal{G}^y represents the global prototypes of class y . τ

denotes the temperature that controls the tolerance of difference between extracted features and the corresponding global prototype. The overall loss function to optimize the tail anchor can be formulated as follows:

$$\mathcal{L}_{ta} = \mathcal{L}_{CE}(\mathcal{F}_{TA}) + \lambda_2 \mathcal{L}_{cons}(\mathcal{F}_{TA}) + \lambda_3 \text{dis}(\mathcal{F}_{TA}, K_s^{ta}), \quad (6)$$

where \mathcal{L}_{CE} is the standard cross-entropy loss, the second term is to adjust its position in the feature space through contrastive loss with the global prototypes (fixed positions). The last term is to enhance the similarity between the selected key and the query key.

Local prototypes. Once the training process of the tail anchor is done, the tail anchors will be frozen and remain unchanged. The local prototype is obtained by averaging features with tail anchors belonging to the same class, computed through

$$P_i^y = \frac{1}{|\mathcal{D}_a^y|} \sum_{(x,y) \in \mathcal{D}_a^y} \mathcal{F}_{TA}^x, \quad (7)$$

where \mathcal{D}_a^y denotes the subset of private dataset of client a of class y . Each client forms a local set of prototypes, which is then uploaded to the server. The server iteratively selects the prototype with the lowest average similarity as the global prototype for that class.

4.3. Selective Input Knowledge Fusion

We follow a common setting, which allows the server to possess a small-scale surrogate dataset, denoted as \mathcal{D}_s . $\{x_s, y_s\}$ are the samples and corresponding labels from \mathcal{D}_s for the distillation process. Assuming the total number of \mathcal{KB} is n , \mathcal{KB}_i is randomly selected as the target for distillation. E'_i represents the enhanced embedding formed by concatenating the x_s 's embedding with the IE from \mathcal{KB}_i , where $\mathcal{V}(\cdot)$ denotes the ViT's feature extraction process. Therefore, the distillation loss can be formulated as follows:

$$\mathcal{L}_{KD} = \frac{1}{n-1} \sum_{\substack{j=1 \\ j \neq i}}^n \text{MSE}(\mathcal{V}(E'_i), \mathcal{V}(E'_j)). \quad (8)$$

4.4. Best Global Prototype Selection

When the server receives local prototype sets from different clients, it reorders them to form a new set P_G according to the class. Specifically, when two clients both have prototypes related to class q , denoted as P_i^q and P_j^q , they will be adjacent to each other in the reordered prototype set. Then, the server computes the similarity between each pair of sets in the collection, forming an adjacency matrix \mathcal{M} . The element of \mathcal{M} is computed through:

$$\mathcal{M}_{ij} = \text{dis}(P_G^i, P_G^j), 0 < i \leq j \leq |P_G|. \quad (9)$$

Notice that if P_G^i and P_G^j belong to same class, then $\mathcal{M}_{ij} = 1$. In each round, the server selects the prototype with the lowest average similarity with all local prototypes as the global prototype for one class. The selection process of the global prototype \mathcal{G}^y for class y can be expressed as follows:

$$\mathcal{G}^y = P_g^s = \underset{y_{low} \leq i \leq y_{high}}{\text{argmin}} \bar{\mathcal{M}}_i = \frac{1}{|P_G|} \sum_{j=1}^{|P_G|} A_{ij}, \quad (10)$$

where y_{low} and y_{high} are the start index and end index of the local prototypes of class y in P_G . P_g^s is the local prototype who has lowest similarity for class y . If the average similarity $\bar{\mathcal{M}}_i$ falls below the threshold Thr during the iteration process, then that prototype is fixed as the global prototype for its class and will not be altered further. As a result, this global prototype will serve as a fixed anchor point for that class in the feature space.

5. Experiments

5.1. Setup

Continual Setting. To ensure that the impact of spatial-temporal data heterogeneity is adequately reflected, we partition the data as follows: Each client continually learn from a task sequence of 5 tasks, and there are 5 clients in the experimental setting. For *CIFAR-100* [14], we allow each client to have access to 15 private classes exclusive to itself, resulting in 25 public classes. Thus, each client has data for a total of 40 classes, with each task consisting of only 8 classes. For *ImageNet-R* [9], we exacerbate spatial data heterogeneity by assigning 40 private classes to each client, with no public classes across clients. Similarly, each task consists of 8 classes. Notice that we use the Dirichlet distribution for the public classes to assign data, ensuring there is no overlap between different clients.

Surrogate Data. We follow the common setting [10, 25] where the server possesses a small surrogate dataset. For *CIFAR-100*, each class has only 20 samples, while for *ImageNet-R*, each class has only 5 samples.

Brief introductions of baseline methods and implementation details are shown in Appendix B.

5.2. Metrics

To verify whether the method can effectively address the challenges brought by spatial-temporal data heterogeneity, we use two new metrics from [36] to evaluate the performance of mitigating forgetting.

Definition 1. (Temporal Knowledge Retention):

$$KR_t = \frac{1}{a} \sum_{i=1}^a \frac{\text{Acc}(\theta_i^r; T_i^0)}{\text{Acc}(\theta_i^0; T_i^0)}, \quad (11)$$

where $\text{Acc}(\theta_i^r; T_i^0)$ denotes the test accuracy of client A_i 's local model at r -th round on the 0-th task and $\text{Acc}(\theta_i^0; T_i^0)$

Table 1: Accuracy of the aggregated global model on local test sets with 5 class-incremental tasks.

Algorithm	Type	CIFAR-100 Task ID					Imagenet-R Task ID				
		1	2	3	4	5	1	2	3	4	5
FedAvg [26]	FL	43.9	50.6	57.3	55.5	61.2	37.7	35.4	35.5	35.8	36.7
FedProx [17]		23.7	22.8	26.0	22.1	23.6	20.2	19.7	19.7	18.9	17.8
FedNova [31]		13.7	18.8	20.1	16.1	15.1	2.0	4.7	4.6	7.8	7.7
FedLwF [23]	FL+CL	36.9	12.5	17.1	13.6	9.7	5.9	7.0	2.6	4.0	3.8
FedViT [4]		70.2	70.0	71.4	66.0	67.3	68.2	59.8	57.3	59.8	57.9
FedL2P [32]		28.4	29.9	29.3	25.4	25.7	27.1	27.6	24.8	25.1	26.5
FedDualP [33]		31.7	42.8	52.8	39.0	46.3	23.5	26.6	26.4	30.2	32.0
GLFC [3]	FCL	82.0	63.1	73.4	64.2	64.8	61.9	67.1	67.0	71.7	57.2
TARGET [41]		54.0	41.4	32.2	13.9	15.9	39.9	15.0	16.0	17.5	16.1
MFCL [1]		46.7	16.2	10.6	14.6	13.5	28.8	14.5	16.2	13.3	13.8
FedMGP [40]		90.2	85.3	90.7	89.2	82.2	77.3	76.8	78.0	75.6	75.4
Ours (FedTA)	FCL	96.1	94.0	94.6	94.4	89.4	81.5	78.8	79.2	80.6	85.0
Ours-w/o TA		78.7	75.5	73.4	69.3	72.3	79.6	72.3	72.3	74.1	63.2
Ours-w/o SIKF		90.7	88.8	91.1	91.4	89.1	80.0	80.5	81.1	82.9	81.7
Ours-w/o BGPS		90.8	92.5	88.4	91.4	88.6	80.5	78.7	78.3	81.7	79.1

denotes the accuracy of client A_i 's local model at the initial round on the 0-th task.

Definition 2. (Spatial Knowledge Retention):

$$KR_s = \frac{1}{a} \sum_{i=1}^a \frac{Acc(\theta_g^r; T_i^r)}{Acc(\theta_i^r; T_i^r)}, \quad (12)$$

where $Acc(\theta_g^r; T_i^r)$ denotes the accuracy of the global model θ_g^r on the current local task T_i^r at client A_i and $Acc(\theta_i^r; T_i^r)$ denotes the accuracy of the local model θ_i^r on its current local task T_i^r .

5.3. Results & Ablation Study

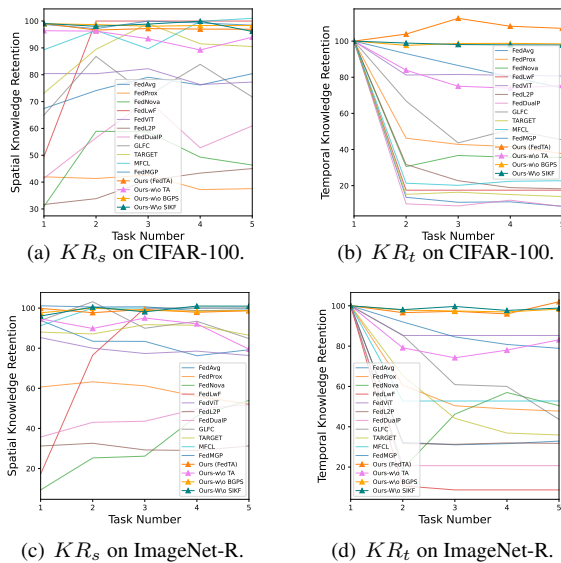


Figure 4: Knowledge retention on different dataset.

Table 1 illustrates the average accuracy of the aggregated global model on local test sets. The performance of *FedViT* is acceptable because all the parameters of its feature extractor are frozen, and only the classification head is involved in training and aggregation. However, it still experiences a certain degree of forgetting. *FedL2P* and *FedDualP*, which introduce trainable parameters on the input side and within the model, perform very well on the local side, achieving around 90% accuracy. However, as we concluded in Sec. 3.2, almost all trainable parameters are directly affected by the data. Consequently, after aggregation, there is significant forgetting on the local test sets.

Surprisingly, the performance of *TARGET*, *FedLwF* and *MFCL*, the three baseline methods that use replay data to mitigate forgetting, is extremely poor. We speculate that the large size of data ($3 \times 224 \times 224$) results in the low quality of replayed pseudo-samples. Moreover, replay-based methods pose a certain risk of privacy leakage in federated learning, limiting the further application of these methods in real-world scenarios. *GLFC* also suffers significant performance degradation when faced with severe spatial-temporal data heterogeneity. However, its performance in Table 1 remains the best among the baseline methods.

FedTA demonstrates the superior performance in these two settings, indicating its successful mitigation of the impact of spatial data heterogeneity. Furthermore, ablation studies highlight the effectiveness of the proposed novel components, with the *Tail Anchor* contributing the most to the performance improvement. However, the selective input knowledge fusion at the server-side sometimes falls below the results of direct weighted averaging on ImageNet-R. We believe this is due to insufficient surrogate data, which prevents adequate selective knowledge fusion.

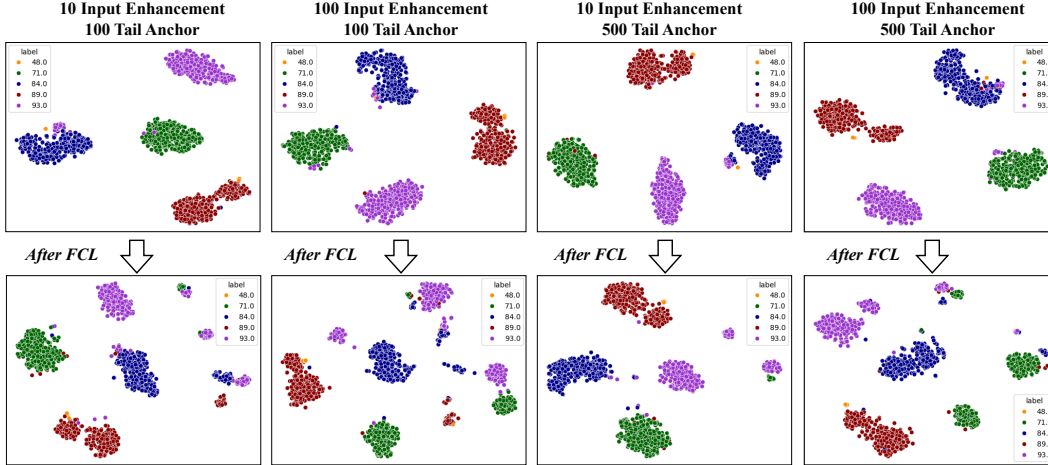


Figure 5: T-SNE for position changes of features corresponding to the same samples after FCL.

Fig. 4 further illustrates the impact of spatial-temporal data heterogeneity on the methods using Equ. 12 and Equ. 11. Only FedTA maintains a high level of temporal and spatial knowledge retention, both at around 98%. While other baseline methods, especially in KR_t , perform extremely poorly. *Tail Anchor* also has been verified to play a significant role in overcoming ST-CF.

Visualization & Sensitivity Analysis. Fig. 5 illustrates that FedTA can effectively control the relative distances between features. When the number of Tail Anchors is set to 100, regardless of the quantity of Input Enhancements, the features of a portion of the data still deviate from their original positions. We speculate that this is due to the insufficient number of Tail Anchors, causing samples from the same class to match with completely dissimilar Tail Anchors. When we set the number of Tail Anchors to 500, the number of shifted points significantly decreases. The combination of 10 Input Enhancements and 500 Tail Anchors shows the most satisfactory results.

Detailed analysis is in Appendix D.

5.4. Privacy & Efficiency Analysis

Computational Burden. During the local training phase, clients train both the Input Enhancement and Tail Anchor components while the ViT itself remains frozen. Therefore, the number of parameters in these two components, along with the classification head, determines the training overhead of FedTA. The size of Input Enhancement is determined by the number, length and embedding dimension, which are set to 10, 10, and 768 in our setting. The size of the Tail Anchor is set to 100×768 . The total size of keys is $(100+10) \times 768$. Therefore, the total number of trainable parameters amounts to 253,440. Compared to a ResNet-18 with 11,306,804 parameters, FedTA is efficient.

Communication Cost. Each client only needs to submit its own input enhancement and local prototypes to the server, with sizes of 76,800 and 768×2 per class, respectively.

Such small communication cost makes FedTA highly efficient, and also makes FedTA scalable for multi-clients.

Privacy Protection. For *Input Enhancement*, on the one hand, ViT is frozen, and on the other hand, due to its minimal number of parameters, it contains extremely little information. Moreover, since this method does not use replay data to alleviate forgetting, privacy protection is further strengthened. However, the uploaded local prototypes are class-specific, and employing cross-mixing might easily reveal the original features, posing a certain degree of privacy risk. If we randomly mix Tail Anchor with features, this issue will be resolved.

Please refer to Appendix E for experimental results in scenarios with 10 tasks and extra baselines.

6. Conclusion

This article extends the issue of data heterogeneity in static FL to the more realistic problem of spatial-temporal data heterogeneity in FCL. Empirical experiments are conducted to demonstrate that spatial-temporal data heterogeneity can cause parameter-forgetting and output-forgetting. Based on this finding, we first define the representative feature embedding of each class as the tail anchor. Then we propose FedTA by utilizing a frozen pre-trained ViT to mitigate parameter forgetting and combining Tail Anchors. Just as a ship in the ocean requires an anchor to hold its position, the features of samples also need a fixed location in the feature space to remain unaffected by the variations introduced by spatial-temporal data heterogeneity, thereby avoiding forgetting. Extensive experiments have verified the superiority of our method, and ablation studies demonstrate the effectiveness of each component, especially *Tail Anchor*. Visualized results demonstrate that our method effectively fixes the features' relative positions, preventing them from being affected by spatial-temporal changes.

Acknowledgment

This work was supported by the National Natural Science Foundation of China (No. 62476228), the Sichuan Science and Technology Program (No. 2024ZYD0180), the Graduate Representative Achievement Cultivation Project of Southwestern University of Finance and Economics (Nos. JGS2024065, JGS2024004) and the Student Study Abroad Exchange Funding Program of Southwestern University of Finance and Economics.

References

- [1] Sara Babakniya, Zalan Fabian, Chaoyang He, Mahdi Soltanolkotabi, and Salman Avestimehr. A data-free approach to mitigate catastrophic forgetting in federated class incremental learning for vision tasks. *Advances in Neural Information Processing Systems*, 36, 2024. 1, 3, 7
- [2] Michael Crawshaw, Yajie Bao, and Mingrui Liu. Federated learning with client subsampling, data heterogeneity, and unbounded smoothness: A new algorithm and lower bounds. *Advances in Neural Information Processing Systems*, 36, 2024. 2
- [3] Jiahua Dong, Lixu Wang, Zhen Fang, Gan Sun, Shichao Xu, Xiao Wang, and Qi Zhu. Federated class-incremental learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10164–10173, 2022. 7, 1, 3
- [4] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2, 5, 7, 1
- [5] Beyza Ermis, Giovanni Zappella, Martin Wistuba, Aditya Rawal, and Cédric Archambeau. Continual learning with transformers for image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3774–3781, 2022. 2
- [6] Tao Fan, Hanlin Gu, Xuemei Cao, Chee Seng Chan, Qian Chen, Yiqiang Chen, Yihui Feng, Yang Gu, Jiayang Geng, Bing Luo, et al. Ten challenging problems in federated foundation models. *arXiv preprint arXiv:2502.12176*, 2025. 1, 2
- [7] Xin Gao, Xin Yang, Hao Yu, Yan Kang, and Tianrui Li. Fed-prok: Trustworthy federated class-incremental learning via prototypical feature knowledge transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4205–4214, 2024. 1
- [8] Raia Hadsell, Dushyant Rao, Andrei A Rusu, and Razvan Pascanu. Embracing change: Continual learning in deep neural networks. *Trends in cognitive sciences*, 24(12):1028–1040, 2020. 1
- [9] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8340–8349, 2021. 6, 1
- [10] Wenke Huang, Mang Ye, and Bo Du. Learn from others and be yourself in heterogeneous federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10143–10153, 2022. 6
- [11] Yan Kang, Tao Fan, Hanlin Gu, Lixin Fan, and Qiang Yang. Grounding foundation models through federated transfer learning: A general framework. *arXiv preprint arXiv:2311.17431*, 2023. 2
- [12] Taehyeon Kim, Eric Lin, Junu Lee, Christian Lau, and Vaikkunth Mugunthan. Navigating data heterogeneity in federated learning: a semi-supervised federated object detection. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. 2
- [13] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. 1
- [14] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6, 1
- [15] Baoxue Li, Pengyu Song, Chunhui Zhao, and Min Xie. Facing spatiotemporal heterogeneity: A unified federated continual learning framework with self-challenge rehearsal for industrial monitoring tasks. *Knowledge-Based Systems*, 289: 111491, 2024. 3
- [16] Qinbin Li, Yiqun Diao, Quan Chen, and Bingsheng He. Federated learning on non-iid data silos: An experimental study. In *2022 IEEE 38th international conference on data engineering (ICDE)*, pages 965–978. IEEE, 2022. 2
- [17] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2:429–450, 2020. 7, 1, 3
- [18] Yichen Li, Qunwei Li, Haozhao Wang, Ruixuan Li, Wenliang Zhong, and Guannan Zhang. Towards efficient replay in federated incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12820–12829, 2024. 3
- [19] Yichen Li, Haozhao Wang, Wenchao Xu, Tianzhe Xiao, Hong Liu, Minzhu Tu, Yuying Wang, Xin Yang, Rui Zhang, Shui Yu, Song Guo, and Ruixuan Li. Unleashing the power of continual learning on non-centralized devices: A survey, 2024. 1
- [20] Yichen Li, Wenchao Xu, Haozhao Wang, Yining Qi, Ruixuan Li, and Song Guo. Sr-fdil: Synergistic replay for federated domain-incremental learning. *IEEE Transactions on Parallel and Distributed Systems*, 2024. 3
- [21] Yujie Li, Xin Yang, Hao Wang, Xiangkun Wang, and Tianrui Li. Learning to prompt knowledge transfer for open-world continual learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 13700–13708, 2024. 4
- [22] Yichen Li, Wenchao Xu, Haozhao Wang, Yining Qi, Jingcai Guo, and Ruixuan Li. Personalized federated domain-incremental learning based on adaptive knowledge match-

- ing. In *European Conference on Computer Vision*, pages 127–144. Springer, 2025. 3
- [23] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. 1, 7
- [24] Jinglin Liang, Jin Zhong, Hanlin Gu, Zhongqi Lu, Xingxing Tang, Gang Dai, Shuangping Huang, Lixin Fan, and Qiang Yang. Diffusion-driven data replay: A novel approach to combat forgetting in federated class continual learning. In *European Conference on Computer Vision*, pages 303–319. Springer, 2025. 3
- [25] Yuhang Ma, Zhongle Xie, Jue Wang, Ke Chen, and Lidan Shou. Continual federated learning based on knowledge distillation. In *IJCAI*, pages 2182–2188, 2022. 6
- [26] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017. 7, 1, 3
- [27] Francesco Pelosin, Saurav Jha, Andrea Torsello, Bogdan Raducanu, and Joost van de Weijer. Towards exemplar-free continual learning in vision transformers: an account of attention, functional and weight regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3820–3829, 2022. 4
- [28] Sara Pieri, Jose Restom, Samuel Horvath, and Hisham Cholakkal. Handling data heterogeneity via architectural design for federated visual recognition. *Advances in Neural Information Processing Systems*, 36:4115–4136, 2023. 2
- [29] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986. 3
- [30] James Seale Smith, Junjiao Tian, Shaunak Halbe, Yen-Chang Hsu, and Zsolt Kira. A closer look at rehearsal-free continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2409–2419, 2023. 2
- [31] Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H Vincent Poor. Tackling the objective inconsistency problem in heterogeneous federated optimization. *Advances in neural information processing systems*, 33:7611–7623, 2020. 7, 1
- [32] Zhen Wang, Liu Liu, Yiqun Duan, Yajing Kong, and Dacheng Tao. Continual learning with lifelong vision transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 171–181, 2022. 2, 7, 3
- [33] Zifeng Wang, Zizhao Zhang, Sayna Ebrahimi, Ruoxi Sun, Han Zhang, Chen-Yu Lee, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, et al. Dualprompt: Complementary prompting for rehearsal-free continual learning. In *European Conference on Computer Vision*, pages 631–648. Springer, 2022. 7, 1, 3
- [34] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 139–149, 2022. 1
- [35] Zichuan Xu, Lin Wang, Weifa Liang, Qiufen Xia, Wenzheng Xu, Pan Zhou, and Omer F Rana. Age-aware data selection and aggregator placement for timely federated continual learning in mobile edge computing. *IEEE Transactions on Computers*, 2023. 3
- [36] Xin Yang, Hao Yu, Xin Gao, Hao Wang, Junbo Zhang, and Tianrui Li. Federated continual learning via knowledge fusion: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 2024. 1, 2, 3, 6
- [37] Zhiqin Yang, Yonggang Zhang, Yu Zheng, Xinmei Tian, Hao Peng, Tongliang Liu, and Bo Han. Fedfed: Feature distillation against data heterogeneity in federated learning. *Advances in Neural Information Processing Systems*, 36, 2024. 1
- [38] Jaehong Yoon, Wonyong Jeong, Giwoong Lee, Eunho Yang, and Sung Ju Hwang. Federated continual learning with weighted inter-client transfer. In *International Conference on Machine Learning*, pages 12073–12086. PMLR, 2021. 1
- [39] Hao Yu, Xin Yang, Xin Gao, Yihui Feng, Hao Wang, Yan Kang, and Tianrui Li. Overcoming spatial-temporal catastrophic forgetting for federated class-incremental learning. In *ACM Multimedia 2024*, 2024. 3
- [40] Hao Yu, Xin Yang, Xin Gao, Yan Kang, Hao Wang, Junbo Zhang, and Tianrui Li. Personalized federated continual learning via multi-granularity prompt. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4023–4034, 2024. 7, 1
- [41] Jie Zhang, Chen Chen, Weiming Zhuang, and Lingjuan Lyu. Target: Federated class-continual learning via exemplar-free distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4782–4793, 2023. 3, 7, 1
- [42] Zhouyangzi Zhang, Bin Guo, Wen Sun, Yan Liu, and Zhiwen Yu. Cross-fcl: Toward a cross-edge federated continual learning framework in mobile edge computing systems. *IEEE Transactions on Mobile Computing*, 23(1):313–326, 2022. 3
- [43] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018. 2
- [44] Zhengyi Zhong, Weidong Bao, Ji Wang, Xiaomin Zhu, and Xiongtao Zhang. Flee: A hierarchical federated learning framework for distributed deep neural network over cloud, edge, and end device. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 13(5):1–24, 2022. 3
- [45] Zhengyi Zhong, Ji Wang, Weidong Bao, Jingxuan Zhou, Xiaomin Zhu, and Xiongtao Zhang. Semi-hfl: semi-supervised federated learning for heterogeneous devices. *Complex & Intelligent Systems*, 9(2):1995–2017, 2023. 1