FED-DUET: DUAL EXPERT-ORCHESTRATED FRAME-WORK FOR CONTINUAL FEDERATED VISION-LANGUAGE LEARNING

Anonymous authorsPaper under double-blind review

000

001

002

003

004

006

008 009 010

011

013

014

015

016

017

018

019

021

023

025

026

027

028

029

031

032 033 034

035

037

038

040

041

042

043

044

046

047

048

051

052

ABSTRACT

Pretrained vision-language models (VLMs), such as CLIP, have shown promise in federated learning (FL) by bringing strong multimodal representations to edge devices. However, continual adaptation remains a core challenge in practical federated settings, where task distributions evolve over time and data remain non-IID across clients. In this emerging area, recent works adopt parameter-efficient fine-tuning (PEFT) as a lightweight way to reduce communication overhead, yet they fail to preserve satisfactory performance under continual learning conditions. Meanwhile, traditional federated continual learning (FCL) methods lack the capacity to maintain cross-modal alignment crucial to VLM performance. We introduce **Fed-Duet**, a novel **Du**al-channel **E**xpert-orchestrated framework for efficient federated continual learning in vision-language models. Fed-Duet features a dualchannel adaptation mechanism, combining server-coordinated semantic prompts with client-personalized modular adapters. These channels are dynamically fused via a cross-attention mechanism, enabling effective knowledge transfer while preserving multimodal alignment and mitigating forgetting. We evaluate Fed-Duet across multiple challenging continual learning tasks in federated vision-language settings and demonstrate that it achieves superior performance and stability compared to existing approaches. Our work highlights the importance of coordinated expert composition in enabling scalable and robust multimodal continual learning. The code is available at https://anonymous.4open.science/r/ FedDuet-0426/.

1 Introduction

Federated Learning (FL) has emerged as a key paradigm for learning from sensitive and siloed data by enabling collaborative model training across decentralized data sources while preserving privacy (McMahan et al., 2017; Kairouz et al., 2021; Li et al., 2020; Zhao et al., 2018). Meanwhile, Large-scale Vision-Language Models (VLMs), pretrained on web-scale data, such as CLIP (Radford et al., 2021), have revolutionized multimodal AI by demonstrating unprecedented capabilities in zero-shot generalization and cross-modal understanding (Jia et al., 2021; Li et al., 2022; Du et al., 2022). Integrating VLMs into the FL paradigm offers great potential for empowering edge devices with advanced multimodal intelligence (Ren et al., 2025), while also introducing new challenges in downstream learning and adaptation.

The most immediate challenge in realizing this paradigm is the large-scale of modern VLMs. Their massive size makes the full-model fine-tuning required in each federated round prohibitively expensive due to immense communication costs. To surmount this bottleneck, Parameter-Efficient Fine-Tuning (PEFT) techniques (Houlsby et al., 2019; Lester et al., 2021; Feng et al., 2023; Li et al., 2024; Yu et al., 2024), such as adapter-based and prompt-based methods, have emerged as an essential solution. By freezing the vast VLM backbone and only training and communicating a small fraction of parameters, these methods drastically reduces the communication overhead (Lu et al., 2023; Guo et al., 2023b; 2024).

Despite recent advances in efficiency, real-world edge environments remain highly challenging: tasks evolve continuously and client data exhibit non-IID distributions. (Kirkpatrick et al., 2017;

Aljundi et al., 2019; Belouadah & Popescu, 2019; Wang et al., 2024). These dynamics give rise to the paradigm of Federated Continual Learning (FCL), which seeks to continually adapt models to new knowledge without catastrophic forgetting, while also addressing the inherent data heterogeneity of federated networks (Lopez-Paz & Ranzato, 2017; Zhao et al., 2020; Dong et al., 2022; So et al., 2022). A solution is therefore needed to address these challenges effectively.

However, directly applying traditional FCL methods to this new context is problematic because they are fundamentally ill-suited for VLMs. First, their reliance on a full-model update paradigm is computationally prohibitive for large-scale models and risks catastrophically overwriting pretrained knowledge. Second, their uni-model perspective ignores the specialized, dual-stream nature of VLMs, thereby threatening the delicate cross-modal alignment that is central to their capabilities.

Furthermore, in the emerging field of federated vision-language model learning, most existing works have focused on applying PEFT techniques as a lightweight solution to reduce communication overhead, yet these methods introduce new challenges when applied under continual learning settings. 1) First, applying a singular PEFT strategy leads to an adaptation imbalance. Relying solely on high-level prompts may fail to capture client-specific nuances, while using only low-level adapters can weaken global semantic consistency. 2) Second, aggregating sparse and heterogeneous PEFT updates across clients risks disrupting the VLMs' intrinsic cross-modal alignment, which is essential for unified vision-language understanding. These challenges highlight the need for a coordinated orchestration mechanism that can effectively integrate diverse client adaptations while preserving the model's semantic integrity. This leads to our core research question:

How can we design an orchestrated framework that efficiently addresses adaptation imbalance and cross-modal misalignment in federated vision-language learning?

To overcome the above challenges, we propose **Fed-Duet**, a novel **Dual**-channel **Expert**-orchestrated framework for FCL in VLMs. Fed-Duet employs a dual-channel architecture that orchestrates two complementary experts—'guiding prompts' for high-level semantic alignment and 'modular adapters' for fine-grained task specialization—whose outputs are dynamically integrated to maintain coherence and adaptability. Our main contributions are summarized as follows:

- We pioneer a paradigm shift in federated VLM learning by identifying a critical gap between the latest PEFT-based approaches and the need for continual adaptation as tasks evolve. By filling this gap, we unlock greater potential for retaining old knowledge while acquiring new insights, paving the way for more sustainable federated VLM learning.
- We propose Fed-Duet, a novel framework designed specifically for federated continual learning in CLIP-like VLMs under data heterogeneity. Fed-Duet enables effective knowledge transfer by synergizing two complementary expert channels, while simultaneously maintaining multimodal alignment and mitigating forgetting.
- We validate the effectiveness of Fed-Duet through extensive experiments on challenging federated VLMs and FCL benchmarks. The superiority of our results offering a scalable solution for lifelong adaptation in federated vision-language settings.

2 Related Work

2.1 CLIP-LIKE VLMs in Federated Learning

Powerful, pre-trained Vision-Language Models (VLMs) like CLIP (Radford et al., 2021) offer a promising foundation for Federated Learning. However, their massive size makes full-model aggregation infeasible due to prohibitive communication costs. To overcome this bottleneck, the community has rapidly adopted PEFT techniques, prominently featuring prompt-tuning and adapter-tuning (Houlsby et al., 2019; Lester et al., 2021; Li & Liang, 2021; Hu et al., 2022).

Prompt-tuning has emerged as the dominant PEFT paradigm for VLMs. The concept was popularized in centralized learning by methods like CoOp (Zhou et al., 2022b;a), which demonstrated that optimizing a few learnable prompt vectors could effectively steer a frozen VLM. The communication efficiency of this approach made it a natural fit for FL. Building on this idea, subsequent works such as PromptFL (Guo et al., 2023b) and pFedPrompt (Guo et al., 2023a) have proposed federated prompt-tuning strategies, leveraging global aggregation and personalization, respectively.

In parallel, adapter-tuning has shown promising results in federated settings, with methods like Fed-CLIP (Lu et al., 2023) demonstrating that communicating small, trainable adapter modules can also successfully adapt VLMs in a federated manner.

While these PEFT paradigms are effective in alleviating training or communicating costs, they have primarily focused on static FL scenarios. *This remains an open and challenging question of how to leverage and potentially synergize them in a continually evolving, non-stationary FCL environment.*

2.2 FEDERATED CONTINUAL LEARNING

Early efforts in Federated Continual Learning (FCL) often involved directly adapting centralized Continual Learning (CL) strategies, such as rehearsal-based iCaRL (Rebuffi et al., 2017) or regularization-based LwF (Li & Hoiem, 2017) and EWC (Lee et al., 2017). However, as these methods were not designed to handle the statistical heterogeneity inherent to federated networks, they typically resulted in compromised global model performance.

This limitation motivated the development of native FCL frameworks with more sophisticated mechanisms. Foundational works like FedWeIT (Yoon et al., 2021) focused on parameter-level transfer, introducing a decomposition scheme where clients selectively re-weight sparse parameters from others to mitigate negative interference. Alternatively, FedKNOW (Luopan et al., 2023) proposed a client-centric, gradient-level solution that extracts critical model weights and integrates restored gradients to guide learning without interference. More recently, the paradigm has shifted towards prompt-based adaptation. Methods like Fed-CPrompt (Bagwe et al., 2023) demonstrate that communicating only lightweight prompts, guided by a dedicated contrastive loss to tackle heterogeneity and asynchronicity, can achieve excellent performance with high efficiency.

While these native FCL methods represent significant advances, their fundamentally uni-modal designs are ill-suited for VLMs. They risk disrupting the delicate cross-modal alignment that is essential to VLM capabilities. *This creates a critical need for a new FCL framework designed specifically to learn continually while preserving this vital cross-modal integrity.*

2.3 MIXTURE-OF-EXPERTS

The sparse Mixture-of-Experts (MoE) paradigm, matured through foundational works like the Switch Transformer (Fedus et al., 2022) and state-of-the-art models like GPT-4 (Achiam et al., 2023), has recently been applied to enhance centralized Continual Learning. For instance, the work by Yu et al. (Yu et al., 2024) successfully uses MoE-Adapters to learn sequential tasks in VLMs, employing a selector to preserve zero-shot capabilities. However, such centralized MoE-for-CL approaches are fundamentally ill-suited for the Federated Continual Learning setting. They would require communicating entire expert modules or complex selector mechanisms, leading to prohibitive communication costs and failing to address client data heterogeneity.

While some reseraches have attempted to incorporate MoE paradigm into FCL setting, to bridge the gap, such as MoAFCL (Zhang & Liu, 2025), they focus solely on applying MoE to server-side adapters, overlooking the crucial role of semantic guidance, limiting their adaptability in Non-IID scenarios. This presents a critical research opportunity for a unified architecture that synergizes both semantic guidance and parametric specialization to create a robust, communication-efficient FCL framework, which our work aims to address.

3 METHODOLOGY

Preliminaries and Problem Formulation. We consider a Federated Continual Learning setting composed of a central server and C clients, $\mathcal{C} = \{c_1, \dots, c_C\}$. Each client c observes a private, sequential stream of tasks $\{D_c^t\}_{t=1}^{T_c}$. The objective in this challenging setting is for each client to learn an effective model, that performs well on its cumulative history of tasks by leveraging both shared knowledge from the federation and its own evolving, local experience.

Framework Overview. We show the overview of **Fed-Duet** in Figure 1, which consists of two principal modules. 1) The server-side *Federated Knowledge Orchestrator* addresses the knowledge orchestration challenge by leveraging a shared knowledge repository and an adaptive gate to dispatch

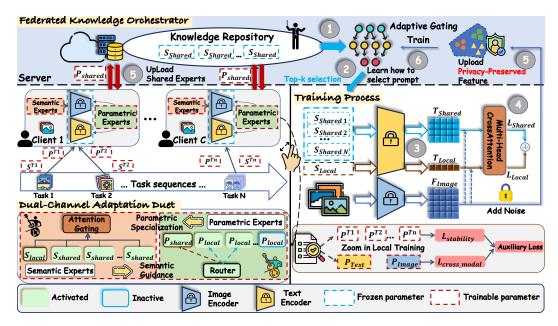


Figure 1: The architecture of **Fed-Duet** illustrates the interaction between the server-side *Federated Knowledge Orchestrator* and the client-side *Dual-Channel Adaptation Duet*. The server employs an adaptive gate to dispatch shared semantic experts. Meanwhile, the client fuses the shared ones with local semantic experts via a cross-attention module, with the other channel adapting parametric experts complementarily. The server aggregates only the shared parametric expert and updates its gate using client features. To further stabilize local training, two auxiliary losses are employed to preserve cross-modal alignment and mitigate forgetting.

tailored semantic guidance. Concurrently, 2) the client-side *Dual-Channel Expert Duet* uses its parallel pathways to resolve adaptation imbalance: a semantic channel employs a Cross-Attention Gating mechanism to fuse prompt-based experts, while a parametric channel adapts adapter-based experts. This architecture is explicitly designed to preserve cross-modal alignment and is further stabilized against catastrophic forgetting by two auxiliary losses, $\mathcal{L}_{cross_modal}$ and $\mathcal{L}_{stability}$.

3.1 FEDERATED KNOWLEDGE ORCHESTRATOR

A key innovation of **Fed-Duet** lies in redefining the server's role from a simple aggregator to an intelligent knowledge orchestrator.

The Orchestrated Knowledge Repository. To ensure the Global Prompt Pool, $\mathcal{P} = \{p_1, \dots, p_K\}$, is a semantically diverse repository from the outset, we forgo random initialization. Instead, we first derive K conceptual anchors by performing K-Means clustering on the word embeddings of a large vocabulary (e.g., ImageNet-1k class names), yielding a set of centroids $\{c_1, \dots, c_K\}$, where each $c_k \in \mathbb{R}^D$.

Each prompt $p_k \in \mathbb{R}^{L \times D}$ is then constructed using the textual template "a photo of <code>[CLS]</code>". The embeddings for the static context tokens ("a", "photo", "of") use their pre-trained values, while the embedding for the learnable <code>[CLS]</code> token is initialized directly with its corresponding centroid c_k . This approach grounds each prompt in a meaningful semantic concept while providing a natural linguistic structure.

The Adaptive Dispatch Mechanism. To efficiently dispatch shared knowledge without broadcasting the entire repository, we introduce an Adaptive Gating, g_{θ} . This gating learns an optimal dispatch policy: it takes a client's privacy-preserving feature summary, \tilde{f}_c , to select the Top-k experts and is subsequently trained using the client's final training loss, $\mathcal{L}_c^{\text{final}}$, as a direct supervisory signal. This feedback mechanism is formalized via a loss-weighted binary cross-entropy (BCE) objective:

$$\mathcal{L}_{\text{gate}} = \sum_{c \in S_r} w_c \cdot \ell_{BCE}(g_{\theta}(\tilde{\boldsymbol{f}}_c), \boldsymbol{y}_c), \tag{1}$$

where the weight $w_c = 1/(\mathcal{L}_c^{\text{final}} + \epsilon)$ gives higher importance to prompt selections that lead to better client performance (i.e., lower loss), and ϵ is a small constant added to ensure numerical stability.

3.2 DUAL-CHANNEL ADAPTATION DUET

The client-side module embodies our core **Du**al-channel **E**xpert principle, which is composed of two expert channels, each leveraging a distinct family of experts to intelligently fuse local and shared knowledge.

♦ Semantic Experts Channel. This channel performs high-level semantic adaptation using learnable prompts, which function as semantic experts. Their knowledge is fused by a Dual-Stream Cross-Attention architecture. In this mechanism, the input image feature acts as a dynamic query to simultaneously interrogate two distinct expert sources: a private Local Semantic Expert that captures client-specific nuances, and a set of server-provided Shared Semantic Experts that embody general knowledge. This process generates separate alignment scores from each source (logits_{local} and logits_{shared}), which are then combined into a final prediction via a weighted sum:

$$Logits_{final} = \lambda \cdot logits_{local} + (1 - \lambda) \cdot logits_{shared}. \tag{2}$$

This design enables the model to dynamically arbitrate between the guidance of local and shared semantic experts on a per-sample basis (detailed in Appendix 6.4).

 \Diamond **Parametric Experts Channel.** While semantic experts provide high-level guidance, they are insufficient for adapting the model's deep feature representations to each client's unique data. We address this limitation with a Channel of Parametric Experts that uses modular adapters. The channel's core innovation is an asymmetric activation strategy: a server-aggregated Shared Expert is always active to provide a stable, foundational representation, while a set of private Local Experts are activated sparsely via a Top-k router for efficient, on-demand specialization. This design creates dedicated parametric pathways for both robust generalization and precise personalization, directly resolving the adaptation imbalance.

Synergy of Dual Expert Channels. The synergy between the two expert channels is a bidirectional reinforcement loop. The establishment of a stable feature foundation by the parametric channel enables more precise high-level guidance from the semantic channel. In turn, this refined semantic guidance provides clearer contextual signals that improve the parametric channel's expert routing. This continuous interplay between foundational stability and refined guidance is the core of our framework's effectiveness.

3.3 Overall Training Strategy

Progressive Decoupled Optimization. To resolve optimization conflicts and actualize the bidirectional synergy between our two channels, we introduce a Progressive Decoupled Optimization strategy. This sequential protocol first exclusively trains parametric experts to establish a stable and robust feature foundation. With this foundation in place, the strategy then freezes the parametric experts and trains the semantic experts to perform precise, high-level refinement. This progressive approach is crucial as it directly enables the virtuous cycle we described previously: by first building a stable parametric base, we empower the semantic experts to provide more effective guidance, which in turn improves subsequent parametric specialization.

3.4 Synergistic Multi-Objective Loss

The client-side training is guided by a composite loss function, \mathcal{L}_{client} , designed to pursue multiple objectives simultaneously. The overall objective is defined as:

$$\mathcal{L}_{\text{client}} = \mathcal{L}_{CE} + \alpha \mathcal{L}_{\text{moe}} + \beta \mathcal{L}_{\text{cross_modal}} + \gamma \mathcal{L}_{\text{stability}}. \tag{3}$$

where the primary objective is standard cross-entropy loss (\mathcal{L}_{CE}), supplemented by a load-balancing loss (\mathcal{L}_{moe}) for the experts, as well as two auxiliary losses, $\mathcal{L}_{cross_modal}$ and $\mathcal{L}_{stability}$ to address key FCL challenges. Specifically, we detail the two loss components as follows:

Routing Consistency Loss ($\mathcal{L}_{cross_modal}$) for Alignment. To counteract the tendency of standard MoE layers to disrupt a VLM's inherent vision-language alignment, we introduce a routing consistency loss inspired by CLIP's contrastive objective. This loss enforces multimodal alignment by

ensuring the expert routing for an image is consistent with that of its corresponding text. The loss is formulated as:

 $\mathcal{L}_{\text{cross.modal}} = \frac{1}{2} \left(\text{CE}(\mathbf{S}/\tau, \mathbf{y}) + \text{CE}(\mathbf{S}^{\top}/\tau, \mathbf{y}) \right), \tag{4}$

where S is a similarity matrix computed from the routing distributions of image-text pairs within a batch, τ is a temperature hyperparamete that controls the sharpness of the distribution, and y represents the ground-truth labels indicating the correct image-text pairings. The symmetric crossentropy objective guides the experts toward learning modality-invariant representations.

Expert Stability Loss ($\mathcal{L}_{\text{stability}}$) for Anti-Forgetting. To mitigate catastrophic forgetting, we employ an expert stability loss that acts as a form of knowledge distillation on the routing policy. This loss preserves learned routing behaviors by regularizing the current policy to remain close to historical ones. It is defined using KL divergence as:

$$\mathcal{L}_{\text{stability}} = D_{KL}(\mathbf{p}^{(t)}||\bar{\mathbf{p}}^{(t-1)}), \tag{5}$$

where $\mathbf{p}^{(t)}$ is the expert routing distribution for the current task and $\bar{\mathbf{p}}^{(t-1)}$ is the historical policy from previous tasks, maintained on a per-layer basis using an exponential moving average. This regularization provides a robust mechanism against forgetting at the expert level.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETTINGS

Datasets and Task Settings. We evaluate our framework on three benchmarks: CIFAR-100 (Krizhevsky & Hinton, 2009) and Tiny-ImageNet (Le & Yang, 2015) for Class-Incremental Learning (partitioned into 5 and 10 tasks), and DomainNet (Peng et al., 2019) for Domain-Incremental Learning. The DomainNet setup, where each domain constitutes a new task, rigorously tests the model's robustness against distribution shifts.

Evaluation Metrics. Model evaluation is conducted using six key continual learning metrics. Overall performance is measured by *Average Accuracy* and *Last Accuracy*. Knowledge retention is quantified by the *Forgetting Measure*, where lower values are better. To analyze the core stability-plasticity dilemma, we also measure *Stability* and *Plasticity*. The trade-off between these is assessed using *Continual Utility*, their weighted sum. Detailed formulations for all evaluation metrics are provided in Appendix 6.3.1.

Baselines. To comprehensively evaluate our framework, we compare it against a range of benchmarks from three main categories: 1) Standard FL Methods, foundational approaches not designed for continual learning, used to establish a performance baseline (e.g., FedAvg (McMahan et al., 2017) and FedProx (Li et al., 2020)); 2) Representative FCL Benchmarks, which include regularization-based approaches (Fed-EWC (Lee et al., 2017), Fed-LwF (Li & Hoiem, 2017)), a parameter-transfer method (FedWeIT (Yoon et al., 2021)), and a gradient-integration approach (Fed-KNOW (Luopan et al., 2023)); 3) SOTA PEFT-based FCL Approaches for VLMs, methods that also leverage Parameter-Efficient Fine-Tuning method for FCL with large-scale models. This category includes prompt-based methods like Fed-CPrompt (Bagwe et al., 2023), pFedMoAP (Luo et al., 2024), a Mixture-of-Experts framework where clients download multiple prompts as experts, and Powder (Piao et al., 2024), which facilitates dual knowledge transfer across clients and tasks using a task correlation matrix to guide prompt aggregation. This group is rounded out by the Mixture-of-Adapter based MoAFCL (Zhang & Liu, 2025). In addition to the FCL methods, our benchmark includes FedCLIP (Lu et al., 2023), a strong PEFT-based method for static Federated VLMs setting.

Furthermore, to ensure a fair comparison against our VLM-native approach, we enhance key baselines by uniformly equipping them with a CLIP backbone, following the protocol of MoAFCL. This includes methods like FedWeIT, FedKNOW, and other ViT-based approaches such as Fed-CPrompt and Powder. This creates a level playing field for all comparisons.

Implementation Details. Following previous research, we employ a pre-trained CLIP model as the VLM backbone across all experiments. Our federated system consists of one central server and five clients. Each experiment is repeated three times with different random seeds (42, 2005, 2026) and we report the averaged outcomes. We use the Adam optimizer with a learning rate of 3e-5. All experiments were conducted on a single NVIDIA RTX 4090 GPU with 24GB of VRAM.

Data partition			II	D]	Non-IID	$\beta = 0.5$)]	Non-IID	$(\beta = 0.1$)
Tasks		T:	=5	T=	:10	T:	=5	T=10		T=5		T=10	
Dataset	Method	Avg↑	Last ↑	Avg↑	Last ↑	Avg ↑	Last ↑						
	FedAvg FedProx	46.44 48.55	21.53 23.32	32.83 37.11	10.56 14.06	35.67 38.76	15.75 17.49	20.71 29.07	5.45 12.70	16.60 19.48	5.50 10.03	10.05 12.58	1.96 2.48
	Fed-EWC Fed-LwF	49.00 53.57	23.45 27.01	38.08 41.11	13.11 17.76	38.44 39.61	18.30 19.87	26.83 30.90	10.02 13.03	28.23 21.09	22.36 12.12	11.60 16.00	2.98 6.85
CIFAR- 100	FedWeIT FedKNOW FedCLIP Fed-CPrompt Powder pFedMoAP MoAFCL	72.52 78.52 78.00 73.18 73.17 76.08 77.97	56.92 74.45 70.11 65.85 65.71 67.8 69.00	72.55 79.61 78.21 74.28 74.30 76.63 77.82	62.64 73.82 68.74 65.59 65.72 66.76 65.91	71.51 77.97 76.37 73.65 73.37 71.02 73.55	56.27 73.27 68.38 66.03 65.56 63.79 64.59	72.25 78.27 77.52 74.13 76.15 71.10 75.54	62.05 72.49 68.94 65.52 67.35 60.50 65.47	70.41 77.83 74.14 73.16 73.56 51.64 65.79	54.43 73.10 65.38 66.10 65.02 46.07 59.65	71.69 78.97 75.93 74.34 73.51 58.94 68.32	62.32 72.67 66.37 65.34 65.51 50.15 60.70
	Fed-Duet (Ours)	86.21	79.10	86.05	78.76	84.85	78.84	84.20	75.97	84.59	76.01	84.20	75.97
	FedAvg FedProx	40.38 43.10	17.82 20.25	27.92 30.72	9.12 11.09	29.88 33.61	12.43 15.25	19.81 22.77	5.86 6.84	14.12 17.42	5.38 8.37	9.22 12.51	2.59 3.72
	Fed-EWC Fed-LwF	42.87 42.73	20.95 17.46	29.97 51.24	9.48 27.10	31.88 32.19	13.44 14.16	27.34 23.90	9.95 7.64	22.26 19.77	10.73 9.30	16.60 15.13	4.26 5.62
Tiny- ImageNet	FedWeIT FedKNOW FedCLIP Fed-CPrompt Powder pFedMoAP MoAFCL	72.04 76.71 76.21 73.04 72.96 74.90 74.20	62.97 72.05 68.52 65.71 66.01 67.51 65.30	72.01 77.94 76.66 74.23 74.18 75.46 74.19	62.81 71.88 67.59 65.75 65.89 66.3 64.00	71.26 75.52 74.51 72.54 72.81 70.90 70.44	62.20 70.80 66.55 65.92 65.58 62.86 62.32	71.52 76.89 75.50 74.30 74.47 70.92 71.32	62.77 70.44 66.12 65.78 65.78 61.93 61.03	71.31 74.71 71.99 72.98 72.60 55.12 62.38	62.37 70.04 63.77 65.89 65.45 47.62 59.38	70.45 75.60 73.57 74.06 74.07 57.37 66.66	61.82 70.15 63.66 65.73 65.20 52.13 59.43
	Fed-Duet (Ours)	82.56	77.64	83.63	75.94	81.00	75.45	81.83	74.86	80.43	74.09	81.65	73.76

Table 1: The performance of our method against SOTA benchmarks on representative CIFAR-100 and Tiny-ImageNet datasets under various FCL settings. Non-IID (β) indicates the Dirichlet parameter is set to β , Avg denotes the average accuracy over all learned tasks, while Last denotes the accuracy on the final task. Best results are in bold. The results show the superiority of our method.

4.2 Overall Performance

Performance Comparison. As shown in Table 1, Fed-Duet demonstrates clear superiority across three key dimensions. 1) *Absolute Accuracy:* It achieves the highest final accuracy, outperforming the strongest baseline (FedKNOW Luopan et al. (2023)) by a significant 5.23% in the challenging CIFAR-100 scenario ($\beta=0.1$, T=10). 2) *Data Heterogeneity Stability:* Fed-Duet is exceptionally robust to Non-IID data. While the performance of advanced methods like pFedMoAP (Luo et al., 2024) degrades by 24% under severe heterogeneity, our method's accuracy drops by a mere 2%. 3) *Continual Learning Stability:* As visualized in Figure 2, Fed-Duet also shows superior resistance to catastrophic forgetting, maintaining a high and stable performance trajectory as new tasks are introduced. These advantages validate that our dual-channel architecture effectively balances client specialization with robust knowledge retention.

Evaluation on Domain Shift. To assess robustness against severe domain shifts, we evaluate our method on DomainNet in domain-incremental setting. As shown in Table 2, Fed-Duet achieves new SOTA performance, with an average accuracy of 68.47%. This marks a substantial improvement of 5.64% over the strongest baseline. This result underscores the effectiveness of our dual-channel adaptation, which preserves general knowledge while specializing in new visual domains to achieve strong generalization.

Table 2: Performance comparison on DomainNet. Best results are in bold. The blue values (Δ) indicate a score lower than our method. Fed-Duet demonstrates strong domain generalization capabilities.

Method	Avg Acc↑	Δ	Last Acc↑	Δ
FedWeIT	62.20	-6.27	58.54	-7.51
FedKNOW	56.09	-12.38	54.25	-11.80
FedCLIP	62.83	-5.64	60.04	-6.01
Fed-CPrompt	56.16	-12.31	54.38	-11.67
Powder	56.14	-12.33	54.35	-11.70
pFedMoAP	59.98	-8.49	56.35	-9.70
MoAFCL	60.92	-7.55	52.52	-13.53
Fed-Duet (Ours)	68.47	-	66.05	-

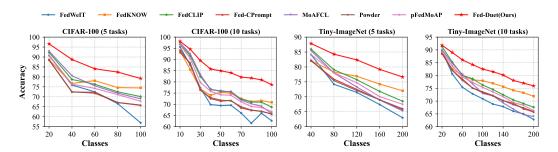


Figure 2: Performance trajectory of our method compared to SOTA benchmarks as the number of classes increases. The plots show accuracy as a function of the number of learned classes under 5 or 10 task scenarios. Our method, Fed-Duet, consistently outperforms all baselines, demonstrating superior resistance to catastrophic forgetting as tasks accumulate. Additional experimental results under various settings are presented in Appendix 6.5 with details.

4.3 ABLATION STUDY

Analysis on Core Compo**nents.** We conduct detailed ablation studies of different components on our framework, as shown in Table 3, from two central innovations of our framework: the dualchannel architecture and the composite loss function. 1) First, we evaluate the architecture by decoupling its pathways, which confirms our adaptation imbalance premise. The semantic-only pathway (Base-w/o PE) is unstable and suffers from high forgetting, while the parametric-only pathway

Table 3: We analyze the performance contribution of Fed-Duet's core components separately. Best results are in bold. The Δ values denote the performance difference from the full model (blue for lower, red for higher).

Method	Avg Acc ↑	Δ	Forget ↓	Δ
Fed-CPrompt	72.98	-7.45	7.95	+0.13
Powder	72.60	-7.83	8.51	+0.69
pFedMoAP	55.12	-25.31	8.28	+0.46
Fed-Duet (Base-w/o PE)	64.34	-16.09	11.89	+4.07
Fed-Duet (Base-w/o SE)	70.64	-9.79	8.89	+1.07
Fed-Duet (Base)	77.96	-2.47	9.22	+1.40
Fed-Duet (Base + $\mathcal{L}_{cross_modal}$)	79.09	-1.34	8.96	+1.14
Fed-Duet (Base + $\mathcal{L}_{stability}$)	79.46	-0.97	8.02	+0.20
Fed-Duet (Full)	80.43	_	7.82	_

(Base-w/o SE), though more stable, remains suboptimal. In contrast, our dual-channel Fed-Duet (Base) model substantially boosts accuracy, proving the synergistic benefit of combining semantic and parametric pathways. 2) Next, We analyze our auxiliary losses against advanced prompt-based FCL methods. While our Fed-Duet (Base) model is highly accurate, it exhibits a higher forgetting score than prompt-based FCL methods. This trade-off is systematically resolved by our two complementary auxiliary losses: the expert stability loss primarily mitigates forgetting, while the routing consistency loss enhances accuracy by enforcing multimodal alignment in the expert representations. The final Fed-Duet (Full) model, benefiting from their synergy, achieves both the highest accuracy and the lowest forgetting score, outperforming the baseline and validating the contribution of our composite loss design.

Analysis of the Stability-Plasticity. Our framework effectively resolves the core stability-plasticity dilemma. As shown in Figure 3a, Fed-Duet simultaneously achieves both the highest stability and plasticity scores, a superior balance that contrasts sharply with leading baselines that are forced to compromise one for the other. To formally quantify this trade-off, we introduce the Continual Utility score, a unified metric representing the weighted aggregation of Stability and Plasticity. This robust balance is further validated by our framework's continual utility score, as shown in Figure 3b. Our method consistently dominates the top score, maintaining a significant performance lead across the entire range of the trade-off hyperparameter ζ . This result shows that Fed-Duet not only establishes a superior balance but also demonstrates remarkable robustness to hyperparameter choices.

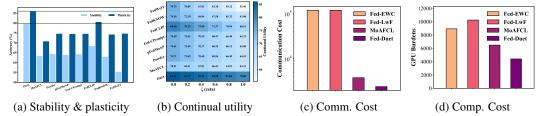


Figure 3: Analysis of the stability-plasticity trade-off and efficiency on CIFAR-100 with 5 incremental tasks. (a) Stability and plasticity scores compared with different methods. Our method excels in both aspects. (b) The continual utility trade-off for our method across varying values of the hyperparameter ζ . A larger ζ indicates more emphasis on stability. Darker colors indicate better utility. (c) Total communication cost (MB, log scale) required by each method. (d) Peak GPU memory footprint (MiB) during local client training. Lower values correspond to higher efficiency.

Efficiency Analysis. Our framework offers significant advantages in both communication and computational efficiency. As shown in Figure 3c, Fed-Duet is not only orders of magnitude more efficient than full-parameter tuning methods like Fed-EWC and FedLwF but also reduces overhead by approximately another order of magnitude relative to the PEFT-based MoAFCL. A rigorous baseline selection is crucial, as existing methods are largely incompatible with our research context. For instance, pFedMoAP (Luo et al., 2024) is not designed for continual learning, while Powder (Piao et al., 2024) utilizes a standard ViT backbone. We therefore selected MoAFCL (Zhang & Liu, 2025) as our primary baseline since it is the only framework specifically engineered for the federated continual learning of VLMs like CLIP. This efficiency extends to computation, measured by GPU memory burdens in Figure 3d, where Fed-Duet again achieves the lowest footprint. This demonstrates that our dual-channel architecture provides a highly resource-efficient solution for FCL, enhancing continual learning while significantly reducing both communication and computation burdens.

Analysis of Large-Scale Scalability. To rigorously assess our framework's scalability, we extended the experimental setup to a more demanding 20-client environment. The performance trajectories, visualized in Figure 4, reveal two critical insights into Fed-Duet's robustness. Firstly, Fed-Duet occupies the highest accuracy plane across all 10 tasks, consistently surpassing all baseline methods. Secondly, its performance slope is notably flatter, demonstrating superior resilience against the performance degradation that hinders the baselines. This visual evidence substantiates that Fed-Duet's core mechanism for mitigating catastrophic forgetting remains potent at scale, ensuring sustained high performance in larger and more complex federated systems.

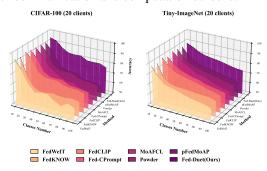


Figure 4: Performance trajectory with large-scale (20-client) setting. Accuracy (z-axis) is plotted across 10 sequential tasks on CIFAR-100 and Tiny-ImageNet. Fed-Duet remains robust under large-scale scenario.

5 CONCLUSION

In this work, we introduced Fed-Duet, a novel framework that addresses the critical challenges of catastrophic forgetting, data heterogeneity, and multimodal alignment degradation in Federated Continual Learning for large-scale Vision-Language Models. Fed-Duet's core innovation is a dual-channel architecture that synergizes two complementary PEFT mechanisms: guiding prompts for robust semantic alignment and modular adapters for fine-grained task specialization. Extensive experiments demonstrate that Fed-Duet significantly outperforms state-of-the-art methods on challenging FCL benchmarks, validating its potential as an effective and practical solution for deploying foundation models in dynamic, decentralized environments.

Ethics Statement. This paper adheres to the ICLR Code of Ethics. The research presented in this paper does not involve human participants, animals, or any other ethical considerations. All data used in this study were obtained from publicly available sources. Our method does not involve discrimination, bias, or fairness concerns.

Reproducibility Statement. To ensure reproducibility, we provide our source code, running scripts, and hyperparameter configurations. We also include detailed algorithm pseudocode, dataset information, and additional experimental results in Appendix 6.5 to facilitate the verification of our findings.

LLM Usage. We employed a large language model (LLM) exclusively for polishing the language and correcting typos. All ideas and methodologies were developed independently.

REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. GPT-4 Technical Report. arXiv preprint arXiv:2303.08774, 2023.
- Rahaf Aljundi, Klaas Kelchtermans, and Tinne Tuytelaars. Task-free continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11254–11263, 2019.
- Gaurav Bagwe, Xiaoyong Yuan, Miao Pan, and Lan Zhang. Fed-cprompt: Contrastive prompt for rehearsal-free federated continual learning. *arXiv preprint arXiv:2307.04869*, 2023.
- Eden Belouadah and Adrian Popescu. Il2m: Class incremental learning with dual memory. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 583–592, 2019.
- 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*, pp. 10164–10173, 2022.
- Yifan Du, Zikang Liu, Junyi Li, and Wayne Xin Zhao. A survey of vision-language pre-trained models. *arXiv preprint arXiv:2202.10936*, 2022.
- William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39, 2022.
- Chun-Mei Feng, Bangjun Li, Xinxing Xu, Yong Liu, Huazhu Fu, and Wangmeng Zuo. Learning federated visual prompt in null space for mri reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8064–8073, 2023.
- Tao Guo, Song Guo, and Junxiao Wang. Pfedprompt: Learning personalized prompt for vision-language models in federated learning. In *Proceedings of the ACM Web Conference 2023*, pp. 1364–1374, 2023a.
- Tao Guo, Song Guo, Junxiao Wang, Xueyang Tang, and Wenchao Xu. Promptfl: Let federated participants cooperatively learn prompts instead of models–federated learning in age of foundation model. *IEEE Transactions on Mobile Computing*, 23(5):5179–5194, 2023b.
- Tao Guo, Song Guo, and Junxiao Wang. Explore and cure: Unveiling sample effectiveness with context-aware federated prompt tuning. *IEEE Transactions on Mobile Computing*, 2024.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International conference on machine learning*, pp. 2790–2799. PMLR, 2019.
 - Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.

- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pp. 4904–4916. PMLR, 2021.
 - Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, and et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021.
 - 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.
 - Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. Technical Report.
 - Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. Stanford University CS231N Course Project Report, 2015.
 - Sang-Woo Lee, Jin-Hwa Kim, Jaehyun Jun, Jung-Woo Ha, and Byoung-Tak Zhang. Overcoming catastrophic forgetting by incremental moment matching. *Advances in neural information processing systems*, 30, 2017.
 - Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv* preprint arXiv:2104.08691, 2021.
 - Hongxia Li, Wei Huang, Jingya Wang, and Ye Shi. Global and local prompts cooperation via optimal transport for federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12151–12161, 2024.
 - Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pretraining for unified vision-language understanding and generation. In *International conference on machine learning*, pp. 12888–12900. PMLR, 2022.
 - Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020.
 - Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv* preprint arXiv:2101.00190, 2021.
 - Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.
 - David Lopez-Paz and Marc' Aurelio Ranzato. Gradient episodic memory for continual learning. In *Advances in neural information processing systems*, volume 30, 2017.
 - Wang Lu, Xixu Hu, Jindong Wang, and Xing Xie. Fedclip: Fast generalization and personalization for clip in federated learning. *arXiv preprint arXiv:2302.13485*, 2023.
 - Jun Luo, Chen Chen, and Shandong Wu. Mixture of experts made personalized: Federated prompt learning for vision-language models. *arXiv preprint arXiv:2410.10114*, 2024.
 - Yaxin Luopan, Rui Han, Qinglong Zhang, Chi Harold Liu, Guoren Wang, and Lydia Y Chen. Fedknow: Federated continual learning with signature task knowledge integration at edge. In 2023 IEEE 39th International Conference on Data Engineering (ICDE), pp. 341–354. IEEE, 2023.
 - 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*, pp. 1273–1282. PMLR, 2017.
 - Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 1406–1415, 2019.

- Hongming Piao, Yichen Wu, Dapeng Wu, and Ying Wei. Federated continual learning via prompt-based dual knowledge transfer. In *Forty-first International Conference on Machine Learning*, 2024.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PmLR, 2021.
- Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 2001–2010, 2017.
- Chao Ren, Han Yu, Hongyi Peng, Xiaoli Tang, Bo Zhao, Liping Yi, Alysa Ziying Tan, Yulan Gao, Anran Li, Xiaoxiao Li, et al. Advances and open challenges in federated foundation models. *IEEE Communications Surveys & Tutorials*, 2025.
- Jinhyun So, Kevin Hsieh, Behnaz Arzani, Shadi Noghabi, Salman Avestimehr, and Ranveer Chandra. Fedspace: An efficient federated learning framework at satellites and ground stations. *arXiv* preprint arXiv:2202.01267, 2022.
- Liyuan Wang, Xingxing Zhang, Hang Su, and Jun Zhu. A comprehensive survey of continual learning: Theory, method and application. *IEEE transactions on pattern analysis and machine intelligence*, 46(8):5362–5383, 2024.
- 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*, pp. 12073–12086. PMLR, 2021.
- Jiazuo Yu, Yunzhi Zhuge, Lu Zhang, Ping Hu, Dong Wang, Huchuan Lu, and You He. Boosting continual learning of vision-language models via mixture-of-experts adapters. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 23219–23230, 2024.
- Dian Zhang and Bingyan Liu. Moafcl: Feature-aware mixture-of-adapter for federated continual learning. In *Proceedings of the 2025 International Conference on Multimedia Retrieval*, pp. 1804–1813, 2025.
- Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and ShuTao Xia. Maintaining discrimination and fairness in class incremental learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 13208–13217, 2020.
- 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.
- Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16816–16825, 2022a.
- Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022b.

6 APPENDIX

648

649 650

651 652

653

654 655

656

679 680

682

683

684

696 697

699

700

6.1 ALGORITHM PSEUDOCODE

We provide the detailed pseudocode for the Fed-Duet framework. The overall process is divided into the server-side orchestration (Algorithm 1) and the client-side dual-channel training (Algorithm 2). This describes the federated learning process for a single continual learning task.

Algorithm 1 Fed-Duet: Server-Side Procedure

```
657
                1: Initialize: Global Prompt Pool \mathcal{P}, Adaptive Gating Net g_{\phi}.
658
                2: Initialize: Global Shared Expert E_{\text{shared}}.
659
                3: for each communication round r = 1, 2, ..., R do
660
                         Select a subset of clients S_r.
                4:
                         Initialize GateNet buffer \mathcal{B} \leftarrow \emptyset.
661
                5:
                6:
                         for each client c \in \mathcal{S}_r in parallel do
662
                7:
                             Receive feature summary f_c from client c.
663
                8:
                             if f_c is available then
                9:
                                 Select prompt indices idx_c \leftarrow TopK_i(g_{\phi}(\mathbf{f}_c)).
              10:
                             else
666
                                   Cold Start }
              11:
667
                                 Select random prompt indices idx_c.
              12:
668
              13:
669
                             Send selected prompts \{\mathcal{P}[i]\}_{i\in\mathbf{idx}_c} and E_{\mathrm{shared}} to client c. Receive updated shared expert E_c^{\mathrm{shared}} and final loss \mathcal{L}_c.
              14:
670
              15:
              16:
                             Add tuple (f_c, idx_c, \mathcal{L}_c) to buffer \mathcal{B}.
671
              17:
                         end for
672
                         Aggregate Shared Experts:
              18:
673
                         m{E}_{	ext{shared}} \leftarrow \sum_{c \in \mathcal{S}_r} \omega_c m{E}_c^{	ext{shared}} \left\{ \omega_c \text{ is client weight} \right\}
Train Gating Net:
              19:
674
              20:
675
                         Compute loss \mathcal{L}_{\text{gate}} = \sum_{(\mathbf{f}_c, \mathbf{idx}_c, \mathcal{L}_c) \in \mathcal{B}} \frac{1}{\mathcal{L}_c + \epsilon} \cdot \text{BCE}(g_{\phi}(\mathbf{f}_c), \mathbf{1}_{\mathbf{idx}_c}).
              21:
676
              22:
                         Update g_{\phi} by descending the gradient of \mathcal{L}_{\text{gate}}.
677
              23: end for
678
```

6.2 Dataset and Task Settings

We evaluate our framework under three distinct continual learning paradigms, each designed to test different facets of model performance. The statistics for all datasets are summarized in Table 4.

Table 4: Statistics of the datasets used in our experiments.

Dataset	Classes	Train	Test
CIFAR-100	100	50,000	10,000
Tiny-ImageNet	200	100,000	10,000
DomainNet	345	586,576	(Total)
Flowers102	102	2,040	6,149
Oxford Pets	37	3,680	3,669
Food101	101	75,750	25,250
Caltech-101	101	3,000	6,000
DTD	47	3,760	1,880

Class-Incremental Learning on CIFAR-100 and Tiny-ImageNet. For this setting, the model must learn new classes over time without forgetting old ones.

• Task Construction: For CIFAR-100 (100 classes), we create scenarios with 5 sequential tasks of 20 classes each, and 10 tasks of 10 classes each. For Tiny-ImageNet (200 classes),

Algorithm 2 Fed-Duet: Client-Side Procedure

702

727 728 729

730

731

732

733

734 735

736

737 738

739

740

741 742

743

744

745

746

747748749

750

751

752 753

754

```
703
           1: Input: Shared prompts \mathcal{P}_s, shared expert \boldsymbol{E}_{\text{shared}}, current round r.
704
           2: Initialize: Local prompt p^{local}, local experts E^{local}.
705
           3: Compute privacy-preserving feature summary f_c from local data D_c^t.
706
           4: if r < R/2 then
                   {Stage 1: Adapt Parametric Experts}
           5:
708
                  for local epoch e = 1, ..., E do
           6:
709
           7:
                     for batch (x,y) from D_c^t do
710
                        Compute features using MoE-Adapter (combining E_{\rm shared} and E^{\rm local}).
           8:
           9:
                        Compute loss \mathcal{L} = \mathcal{L}_{CE} + \alpha \mathcal{L}_{moe} + \beta \mathcal{L}_{cross} + \gamma \mathcal{L}_{stab}.
711
          10:
                        Update parametric experts parameters.
712
          11:
                     end for
713
          12:
                  end for
714
          13: else
715
          14:
                  Stage 2: Refine Semantic Guidance
716
          15:
                  Freeze MoE-Adapter parameters.
717
                  for local epoch e = 1, ..., E do
          16:
718
                     for batch (x, y) from D_c^t do
          17:
719
                        Compute logits using Dual-Stream Cross-Attention.
          18:
720
          19:
                        Compute loss \mathcal{L} = \mathcal{L}_{CE}.
721
          20:
                        Update semantic-experts-related parameters.
                     end for
          21:
722
          22:
                  end for
723
          23: end if
724
          24: Let \mathcal{L}_{final} be the average loss from the last local epoch.
725
          25: Send updated E^{\text{shared}}, \mathcal{L}_{\text{final}}, and f_c to server.
```

we create scenarios with 5 tasks (40 classes/task) and 10 tasks (20 classes/task). The class order is fixed for reproducibility.

• Data Heterogeneity: To simulate federated environments, we partition data among clients using a Dirichlet distribution over class labels, controlled by β . We test under moderate ($\beta = 0.5$), severe ($\beta = 0.1$), and IID settings.

Domain-Incremental Learning on DomainNet. This challenging scenario evaluates the model's ability to adapt to severe distribution shifts, where each task is a new visual domain.

- Task Construction: We use all six domains from the dataset (Clipart, Infograph, Painting, Quickdraw, Real, and Sketch) as a sequence of six tasks. The model learns these domains in alphabetical order.
- Evaluation Protocol: This scenario uses a specific protocol to measure cumulative knowledge. After training on each new domain, the model is evaluated on a dynamically expanding test set comprising the test splits of all previously learned domains plus the current one (e.g., after training on *Sketch*, the model is tested on the combined test sets of *Real* and *Sketch*). The evaluation uses a zero-shot classification head spanning all 345 classes of DomainNet.

Multi-Domain Task-Incremental Learning on Fine-grained Datasets. To further validate the generalizability of our framework, we construct a task-incremental scenario where the model must learn a sequence of entirely different datasets, assuming the task identity is known at inference time.

• Task Construction: We use a sequence of five distinct, fine-grained visual classification datasets, learned in the following order: Flowers102 → OxfordPets → Food101 → Caltech101 → DTD. Each dataset constitutes a single, separate task, testing the model's ability to acquire and retain diverse knowledge without interference.

6.3 FORMAL DEFINITIONS OF METRICS AND LOSSES

6.3.1 EVALUATION METRICS

Here, we provide the formal definitions for the standard continual learning metrics used to evaluate our framework. Let T be the total number of tasks, and let $a_{i,j}$ be the accuracy of the model on task j after it has been trained on task i.

Average Accuracy (Avg Acc) This metric measures the overall performance across all tasks after training is complete.

$$\operatorname{Avg} \operatorname{Acc} = \frac{1}{T} \sum_{j=1}^{T} a_{T,j} \tag{6}$$

Forgetting Measure (Forget) This metric quantifies how much the model forgets about past tasks.

Forget =
$$\frac{1}{T-1} \sum_{j=1}^{T-1} \left(\max_{k < T} a_{k,j} - a_{T,j} \right)$$
 (7)

Stability This measures the ability to retain knowledge of past tasks. After training on task i, it is the average accuracy on all prior tasks:

Stability_i =
$$\frac{1}{i-1} \sum_{j=1}^{i-1} a_{i,j}$$
 (8)

Plasticity This measures the ability to learn the current new task. After training on task i, it is the accuracy on that task:

$$Plasticity_i = a_{i,i} (9)$$

Continual Utility This metric provides a score to evaluate the trade-off between stability and plasticity, defined as their weighted linear combination. The hyperparameter $\zeta \in [0,1]$ controls the balance, assigning a weight of ζ to stability and $(1-\zeta)$ to plasticity. A larger ζ places more emphasis on stability.

$$Utility_i = \zeta \cdot Stability_i + (1 - \zeta) \cdot Plasticity_i$$
 (10)

6.3.2 DETAILS OF PRIMARY AND REGULARIZATION LOSSES

Cross-Entropy Loss (\mathcal{L}_{CE}) The primary objective for the main classification task is the standard cross-entropy loss. It is designed to minimize the dissimilarity between the model's predicted probability distribution and the ground-truth label distribution. The formula is:

$$\mathcal{L}_{CE} = -\sum_{i=1}^{C} y_i \log(\hat{y}_i), \tag{11}$$

where C is the total number of classes, y is the one-hot encoded ground-truth label vector, and \hat{y} is the predicted probability distribution produced by applying the Softmax function to the model's final logits. Minimizing this loss guides the model to assign the highest possible probability to the correct class.

MoE Load Balancing Loss (\mathcal{L}_{moe}) To prevent the MoE router from consistently favoring a small subset of experts (a phenomenon known as expert collapse), we employ an auxiliary load-balancing loss. This loss encourages the router to distribute the computational load as evenly as possible across all available experts, ensuring that all experts are sufficiently trained. Following common practice, we define this loss as:

$$\mathcal{L}_{\text{moe}} = \alpha \cdot N \sum_{i=1}^{N} f_i \cdot P_i, \tag{12}$$

where N is the number of experts and α is a scaling hyperparameter. For a batch of tokens, f_i is the fraction of tokens dispatched to expert i, and P_i is the average router probability (the gating value) assigned to expert i over the batch. This objective incentivizes the router to assign higher probabilities (P_i) to experts that are utilized less frequently (low f_i), thus promoting balanced expert utilization throughout training.

6.4 IMPLEMENTATION AND ARCHITECTURAL DETAILS

6.4.1 DUAL-STREAM CROSS-ATTENTION MECHANISM

To create a dynamic, content-aware fusion of knowledge that is native to the CLIP architecture, we introduce a Dual-Stream Cross-Attention mechanism. This design moves beyond static feature combination by using the visual feature of an input image, F_{img} , as a dynamic Query to actively interrogate the available textual knowledge. This process unfolds in two parallel streams.

Local Knowledge Stream. First, to capture client-specific nuances, each client learns a private, local prompt p_c^{local} . This prompt is prepended to the class token embeddings and processed by the text encoder to produce the local text feature set, z_c^{local} . The local knowledge stream then attends to this private local prompt feature, which serves as both Key and Value. This produces an image-aware local representation and its corresponding alignment score, or logits:

$$logits_{local} = Cross-Attention(F_{img}, z_c^{local}, z_c^{local})$$
(13)

Shared Knowledge Stream. Concurrently, to evaluate the relevance of the generalized knowledge from the federation, the shared knowledge stream attends the same image query to the set of k received shared prompt features, Z_s . This computes the shared logits:

$$logits_{shared} = Cross-Attention(F_{img}, Z_s, Z_s)$$
(14)

Final Fusion. Finally, to create a prediction that flexibly balances local expertise with global consensus, the logits from both streams are dynamically fused via a weighted sum, controlled by a hyperparameter λ :

$$Logits_{final} = \lambda \cdot logits_{local} + (1 - \lambda) \cdot logits_{shared}$$
 (15)

This mechanism allows the model to dynamically weigh the importance of local versus global semantic guidance based on the task and data.

6.5 ADDITIONAL RESULTS AND VISUALIZATIONS

To further substantiate the claims made in the main paper, this section provides additional visualizations and analyses of our framework's performance.

6.5.1 Detailed Performance Trajectory Analysis

Figure 5 presents a comprehensive visualization of the accuracy trajectory for Fed-Duet against baseline methods across Non-IID scenarios. These plots complement the summary results in Table 1 of the main paper by showing the performance degradation as more tasks are learned. Across all settings, our method consistently establishes the highest accuracy. More importantly, its performance curve exhibits a significantly flatter slope compared to all baselines, visually confirming its superior resistance to catastrophic forgetting, especially in the Non-IID scenarios.

6.5.2 MULTI-DOMAIN TASK-INCREMENTAL LEARNING

In the main body of our paper, we extensively evaluated our framework under class-incremental and domain-incremental learning scenarios. To further demonstrate the robustness and transferability of our proposed Fed-Duet framework, we present an auxiliary experiment in this section under the **Multi-Domain Task-Incremental Learning** setting.

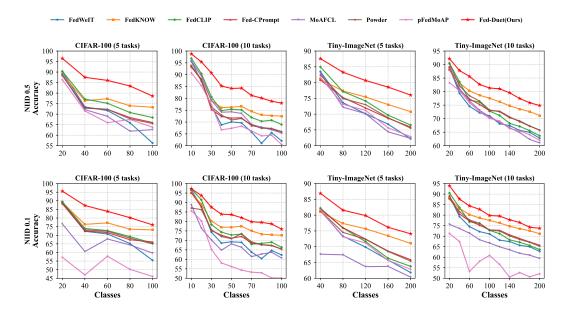


Figure 5: Detailed performance trajectories on CIFAR-100 and Tiny-ImageNet. **Rows**: The top row displays results under the IID setting, while the bottom row shows the more challenging Non-IID ($\beta=0.1$) setting. **Columns**: From left to right, the columns correspond to CIFAR-100 (5 tasks), CIFAR-100 (10 tasks), Tiny-ImageNet (5 tasks), and Tiny-ImageNet (10 tasks).

Table 5: Multi-domain Task-incremental learning (MTIL) performance on five benchmark datasets. The final columns show the average accuracy and the Δ (blue for lower) compared to our method.

Method	Flowers102	OxfordPets	Food101	Caltech101	DTD	Average	Δ
FedWeIT	62.03	53.97	66.35	69.50	71.68	64.71	-8.61
FedKNOW	63.69	53.98	64.81	70.02	72.31	64.96	-8.36
FedCLIP	68.97	64.67	71.83	76.31	79.84	72.32	-1.00
Fed-CPrompt	63.96	54.10	64.95	70.21	72.46	65.14	-8.18
Powder	65.48	58.37	67.87	72.32	76.47	68.10	-5.22
pFedMoAP	66.56	59.58	69.82	73.22	77.77	69.39	-3.93
MoAFCL	63.18	54.09	66.32	71.45	74.85	65.98	-7.34
Fed-Duet (Ours)	70.35	64.97	73.33	77.77	80.17	73.32	-

Analysis of Multi-Domain Task-Incremental. The performance of all methods on this Multi-domain Task-incremental benchmark is presented in Table 5. The results unequivocally show that our Fed-Duet framework achieves the highest accuracy across all five datasets in the sequence. This consistent state-of-the-art performance in a third, distinct continual learning paradigm provides strong evidence for the generalizability of our approach. It confirms that the superiority of Fed-Duet is not confined to a specific problem formulation but is instead a robust attribute stemming from its core design, further validating its advanced capabilities in diverse continual learning environments.

Table 6: Performance of FedDuet with Varying DP Noise on CIFAR-100 and Tiny-ImageNet.

Dataset	DP Noise Multiplier (σ)	Last Acc. (%)	Avg Acc. (%)
	0.1	72.16	80.77
CIFAR-100	0.5	69.58	79.79
	1.0	68.78	78.54
	0.1	68.96	76.34
Tiny-ImageNet	0.5	67.22	75.28
-	1.0	66.69	74.58

Table 7: Performance comparison on CIFAR-100 and Tiny-ImageNet under compression strategy.

(a) CIFAR-100

(b) Tiny-ImageNet

Method	T:	=5	T=10		
11201104	Avg↑	Last ↑	Avg ↑	Last ↑	
FedWeIT	38.71	25.21	39.81	24.52	
FedKNOW	42.08	34.60	46.69	35.86	
FedCLIP	49.39	34.27	49.39	32.87	
Fed-CPrompt	34.75	24.10	37.69	23.88	
Powder	34.65	23.71	37.51	23.55	
pFedMoAP	37.99	25.28	39.73	25.15	
MoAFCL	40.69	27.47	42.29	27.17	
Fed-Duet (Ours)	61.31	49.04	63.29	46.45	

Method	T:	=5	T=10		
112011011	Avg ↑	Last ↑	Avg↑	Last ↑	
FedWeIT	23.67	16.90	25.64	15.58	
FedKNOW	28.60	22.86	32.01	24.02	
FedCLIP	35.15	22.82	35.25	21.18	
Fed-CPrompt	23.80	15.98	25.36	15.86	
Powder	24.03	15.86	25.46	15.82	
pFedMoAP	27.75	18.26	28.63	17.29	
MoAFCL	24.81	16.91	25.81	15.68	
Fed-Duet (Ours)	46.62	34.96	47.58	33.29	

6.5.3 ROBUSTNESS ANALYSIS UNDER NOISY CONDITIONS

To further investigate the robustness of our proposed Fed-Duet framework, we conduct a sensitivity analysis with respect to varying levels of data noise. This experiment simulates a common real-world challenge where client data quality may be inconsistent. Table 6 illustrates the performance trajectory of our method as the noise intensity increases.

It is evident that Fed-Duet exhibits remarkable resilience. Across both CIFAR-100 and Tiny-ImageNet, the accuracy degradation is notably graceful and marginal, even under significant noise levels. This stability suggests that our framework's dual-stream architecture and guidance mechanisms are not easily perturbed by data corruption. The learned representations can effectively distill the core semantic information while filtering out irrelevant noise. This result strongly validates the inherent robustness of our approach and its transferability to non-ideal, practical federated learning scenarios.

6.5.4 AUXILIARY EXPERIMENT: PERFORMANCE WITH A COMPRESSED BACKBONE

In our main experiments, we utilized the full CLIP ViT-B/16 model to ensure a fair comparison with prior works. However, practical federated learning scenarios often involve clients with **limited computational resources**, necessitating more lightweight models. To evaluate the robustness and transferability of our framework under such constraints, we conduct an auxiliary experiment where all methods are built upon a compressed CLIP backbone.

Pruning the CLIP Backbone. We adopt a simple yet effective layer-dropping strategy to compress the CLIP backbone. Specifically, we prune the model by removing the first two and last two transformer layers from both the vision and text encoders. While more sophisticated pruning techniques could be explored in future work, this method serves to effectively reduce the model's computational footprint for this analysis.

Sustained Performance with a Compressed Model. As shown in Table 7, Fed-Duet maintains its state-of-the-art performance even with the compressed backbone, consistently outperforming all baselines. This result demonstrates that the superiority of our framework is intrinsic to its architectural design, not merely a byproduct of a large-scale model. Retaining this performance advantage validates the robustness and practical applicability of Fed-Duet in resource-constrained federated environments.

6.5.5 Hyperparameter Analysis

We conduct a hyperparameter analysis, with results in Table 8. The findings demonstrate that our method exhibits strong robustness, showing low sensitivity to parameters such as the number of K-means clusters across a wide range of values.

Table 8: Hyperparameter analysis for Fed-Duet on CIFAR-100.

(a) Impact of Experts (E) and Top-K (k).

Experts (E)]	Гор-К (k	:)	
-F ()	k=2	k=4	k=8	
2	84.97	_	_	
4	85.77	_	_	
8	86.21	86.42	_	
16	87.12	86.98	86.75	

Clusters	Accuracy (%)
8	86.23
16	86.17
32	86.10
64	86.21
128	86.34

Table 9: Ablation study on our Progressive Decoupled Optimization strategy. Results are average accuracy. The Δ indicates the performance drop relative to our method, highlighted in blue.

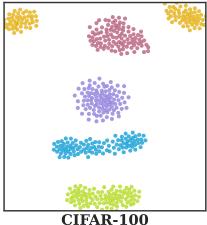
Training Strategy	CIFAR-100	Tiny-ImageNet	Average	Δ
Semantic Experts First	85.53	81.28	83.41	-0.98
Joint Training	86.02	81.88	83.95	-0.44
Parametric Experts First (Ours)	86.21	82.56	84.39	_

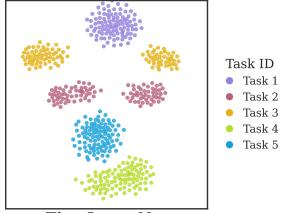
6.5.6 ABLATION ON OPTIMIZATION STRATEGY

We conduct an ablation study on our optimization strategy in Table 9. The results clearly show that our approach significantly outperforms both "Semantic Experts First" and "Joint Training" alternatives, which validates the effectiveness of our Progressive Decoupled Optimization Strategy.

6.5.7 FEATURE SPACE VISUALIZATION VIA T-SNE

To qualitatively verify the efficacy of our framework in mitigating catastrophic forgetting, we employ t-SNE to visualize the learned feature space after the model has trained on all sequential tasks. As depicted in Figure 6, we project the feature embeddings of test samples into a 2D space. The visualization reveals that Fed-Duet maintains a highly structured and separable feature space, evidenced by the dense and distinct clusters formed by each task's samples. This clear spatial separation serves as strong qualitative evidence that our framework preserves the integrity of feature representations for earlier tasks, thereby effectively mitigating catastrophic forgetting at the feature level.





CIFAF

Tiny-ImageNet

Figure 6: t-SNE visualization of feature representations for test samples from all tasks on CIFAR-100 (Left) and Tiny-ImageNet (Right) after the final task. Each color corresponds to a different continual learning task.

6.5.8 Analysis of Expert Utilization

The expert activation patterns, visualized in Figure 7, reveal both effective load balancing and a high degree of specialization. The heatmaps show a distributed activation strategy, confirming that our model avoids expert collapse by engaging a diverse set of experts. Furthermore, we observe a strong tendency for specific experts to be responsible for distinct task categories, particularly in the deeper layers. This dual observation of balance and specialization mutually corroborates the findings from our t-SNE visualization (Figure 6), which similarly indicates that experts form distinct, specialized clusters.

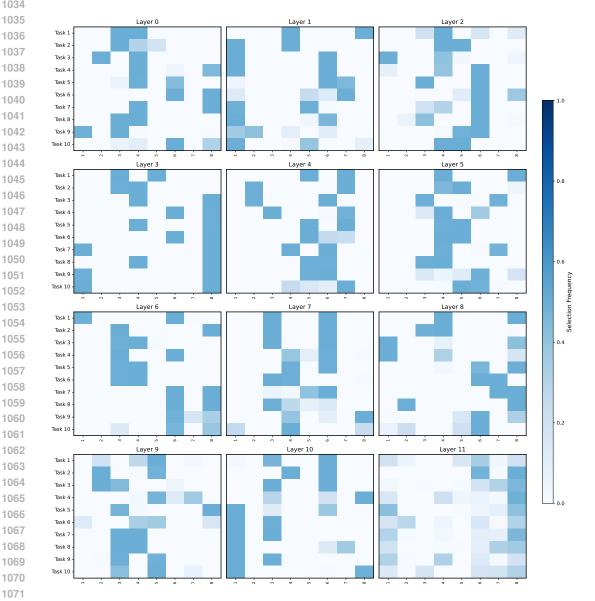


Figure 7: Visualization of expert selection frequency per task in our proposed FedDuet framework. The heatmaps illustrate which of the 8 experts are activated for each of the 10 sequential tasks across 12 transformer layers. The color intensity corresponds to the selection frequency, highlighting how expert utilization evolves for different tasks and layers.