

SecureGate: Learning When to Reveal PII Safely via Token-Gated Dual-Adapters for Federated LLMs

Mohamed Shaaban and Mohamed Elmahallawy*

Washington State University

{mohamed.shaaban, mohamed.elmahallawy}@wsu.edu

Abstract

Federated learning (FL) enables collaborative training across organizational silos without sharing raw data, making it attractive for privacy-sensitive applications. With the rapid adoption of large language models (LLMs), federated fine-tuning of generative LLMs has gained attention as a way to leverage distributed data while preserving confidentiality. However, this setting introduces fundamental challenges: (i) *privacy leakage* of personally identifiable information (PII) due to LLM memorization, and (ii) a persistent *tension between global generalization and local utility* under heterogeneous data. Existing defenses, such as data sanitization and differential privacy, reduce leakage but often degrade downstream performance. We propose SECUREGATE, a *privacy-aware federated fine-tuning framework* for LLMs that provides fine-grained privacy control without sacrificing utility. SECUREGATE employs a dual-adaptor LoRA architecture: a *secure adapter* that learns sanitized, globally shareable representations, and a *revealing adapter* that captures sensitive, organization-specific knowledge. A token-controlled *gating module* selectively activates these adapters at inference time, enabling controlled information disclosure without retraining. Extensive experiments across multiple LLMs and real-world datasets show that SECUREGATE improves task utility while substantially reducing PII leakage, achieving up to a $31.66\times$ reduction in inference attack accuracy and a $17.07\times$ reduction in extraction recall for *unauthorized* requests. Additionally, it maintains 100% routing reliability to the correct adapter and incurs only minimal computational and communication overhead. Code is available at <https://github.com/wsu-cyber-security-lab-ai/SecureGate>.

1 Introduction

In the era of large language models (LLMs), organizations across industry and academia increas-

ingly rely on pre-trained models to support domain-specific applications such as scientific analysis, healthcare, and enterprise workflows (Ling et al., 2023; Lu et al., 2025). Training LLMs from scratch is often prohibitively expensive in terms of computation, data, and engineering effort; consequently, fine-tuning pre-trained models has become the dominant approach for domain and task adaptation (Pratap et al., 2025). While effective, independently fine-tuned models remain confined to local data silos, preventing organizations from benefiting from complementary knowledge distributed across institutions and often resulting in limited generalization and robustness (Lin et al., 2022).

Federated learning (FL) (Kairouz et al., 2021) provides a principled framework to address this limitation by enabling collaborative model training without sharing raw data. In a standard FL setup, each participant fine-tunes a local copy of a shared pre-trained model using private data and periodically transmits model updates to a central aggregator for global aggregation. Although FL has demonstrated success for conventional machine learning and smaller neural models, extending it to large-scale generative LLMs poses fundamental challenges. LLMs are prone to *memorizing sensitive information* (Carlini et al., 2021), and naive aggregation of model updates can *leak personally identifiable information (PII)* (Cheng et al., 2025). Moreover, aggregating updates often favors average global performance, degrading utility on participants' domain-specific data and intensifying the tension between global generalization and local personalization (Parthasarathy et al., 2024).

To reduce PII leakage during collaborative fine-tuning, organizations commonly rely on defenses such as data scrubbing (Akkus et al., 2025) and differential privacy (DP) (Hassan et al., 2019). While effective at limiting memorization-based leakage, these methods often degrade model utility by removing task-relevant context or injecting excessive

*Corresponding author.

noise, particularly as privacy constraints tighten. In practice, this trade-off can negate the benefits of cross-organization collaboration, yielding global models that underperform locally trained alternatives. Moreover, existing defenses scale poorly to large generative LLMs and rarely address personalization, communication efficiency, and client heterogeneity jointly. Thus, this work addresses the following question: *How can federated fine-tuning of generative LLMs simultaneously ensure strong privacy guarantees, high global and local utility, and practical efficiency?*

We propose SECUREGATE, a federated fine-tuning framework for organizations adapting LLMs in privacy-sensitive settings. SECUREGATE employs two lightweight (Low Rank Adapters) LoRAs (Hu et al., 2022): a *secure adapter* that learns sanitized, shareable representations, and a *revealing adapter* that retains sensitive, organization-specific knowledge. A token-controlled gating module selectively routes prompts at inference time, enabling fine-grained access control without retraining. By aligning global and local representations, SECUREGATE enables personalized adaptation while maintaining communication efficiency and scalability. We summarize our contribution as:

1. Token-controlled privacy paradigm (SECUREGATE). We introduce a token-based FL paradigm that protects private organizational knowledge by *locking sensitive personalized components*. Access to these components is granted only when a learned `[Special-Token]` is presented; unauthorized parties—including the aggregator and other clients—cannot access or infer the protected information.

2. Dual-adapter access-controlled architecture. SECUREGATE realizes model adaptation through:

- *Secure adapter*: Encodes globally shareable knowledge while preserving PII using privacy defenses such as noise injection (e.g., DP) or data scrubbing (e.g., adding `[MASK]`).
- *Revealing adapter*: Encodes highly sensitive, organization-specific knowledge and is activated only for queries from authorized organizations that include a valid `[Special-Token]`.

3. Privacy-aware gating and routing. A lightweight gating module dynamically routes prompts to the appropriate adapter based on token signals, enabling selective disclosure at inference

time: authorized queries access sensitive knowledge, while unauthorized queries rely solely on the secure adapter.

4. Empirical evaluation. We evaluate SECUREGATE on multiple LLMs and real-world datasets (e.g., ECHR, Yelp Reviews). Results show strong privacy–utility trade-offs: *authorized* performance reaches 25.20% inference accuracy and 6.32 perplexity (PPL) score, while *unauthorized* access is suppressed to a leakage floor of 4.20% and 15.89 PPL. Overall, SECUREGATE achieves up to a $31.66\times$ reduction in inference accuracy and a $17.07\times$ reduction in extraction recall for *unauthorized* requests, while preserving low computation & communication cost and 100% routing reliability.

2 Related Work

2.1 Federated Learning for LLMs

FL enables collaborative training across distributed clients without sharing raw data. Foundational methods such as FedAvg (McMahan et al., 2017), FedAdagrad, FedAdam, FedYogi (Reddi et al., 2020), and FedAvgM (Hsu et al., 2019) provide techniques for aggregating client updates and handling heterogeneous data. Early work addressed statistical heterogeneity, while recent approaches like HPFL (Shen et al.) and LG-Mix (Jiang et al., 2024) decompose models into shared backbones and personalized components.

The rise of LLMs has motivated adapting classical FL approaches to manage massive parameters and memory across clients while enabling collaborative training of globally shared LLMs. Directly transmitting full LLM weights is infeasible due to prohibitive communication and storage costs. Lightweight adaptation techniques such as LoRA mitigate this challenge: FedLoRA (Wu et al., 2024) decomposes each layer into a shared full-rank component and a client-specific low-rank adapter to reduce non-IID effects, while pFedLoRA (Yi et al., 2023) iteratively trains local adapters across heterogeneous clients’ architectures, aggregating them into a global adapter with minimal communication. These methods make federated LLM training practical, but they primarily optimize utility, *leaving models vulnerable to memorization and unintended disclosure of sensitive local data*.

2.2 Privacy Leakage in Federated LLMs

Federated LLMs are still in their infancy, but initial efforts have begun to address privacy concerns. Per-

sonalized FL (PFL) methods aim to improve generalization while limiting information leakage by combining shared knowledge with client-specific adaptations. For example, FedDPA (Long et al., 2024) introduces global–local adapter pairs dynamically fused at inference, yet it does not provide inherent access control for sensitive data.

On the security side, cryptographic techniques such as secure aggregation have been extensively explored in classical FL (Allouah et al., 2025; Zuo et al., 2024; Pu et al., 2025; Mansouri et al., 2023), but their application to federated LLMs remains limited. Most focus solely on training-time protection, leaving inference-time privacy largely unaddressed. Moreover, these methods incur substantial computation and communication overheads, which would scale poorly to LLM-sized models. FDLORA (Qi et al., 2024) takes an early step by separating global and local adapters, but it provides no privacy guarantees during inference.

Building on these insights, we propose SECURE-GATE, a framework that combines adapter-based personalization with token-controlled, per-request selective disclosure, *enabling fine-grained, secure access to federated LLM knowledge while preserving both training- and inference-time privacy and maintaining high local and global model utility.*

3 Problem Formulation & Threat Model

3.1 Problem Formulation

Personalized Federated Learning. FL enables N organizations, each with a private dataset D_n ($n = 1, \dots, N$), to collaboratively train a shared model without exchanging raw data. At each communication round t , organizations perform local training starting from the current global parameters w^t , producing updated local parameters w_n^{t+1} . The server aggregates these updates (e.g., via FedAvg (McMahan et al., 2017)) to obtain the next global model: $w^{t+1} = \sum_{n=1}^N \frac{|D_n|}{|D|} w_n^{t+1}$, and broadcasts it to all organizations until convergence.

Personalized FL (PFL) (Liu et al., 2024) extends FL to preserve per-organization utility by maintaining both shared and local parameters: $w_n^t = w^t + w_{p,n}^t$, where $w_{p,n}^t$ is the personalized component. A typical objective is

$$\min_{w^t, \{w_{p,n}^t\}} \sum_{n=1}^N \left(\mathcal{L}_n(w^t + w_{p,n}^t) + \lambda \Omega(w_{p,n}^t) \right), \quad (1)$$

with \mathcal{L}_n the local loss, Ω a regularizer, and λ a

trade-off weight. Local updates follow standard gradient steps: $w_{p,n}^{t+1} = w_{p,n}^t - \eta \nabla \mathcal{L}_n(w^t + w_{p,n}^t)$.

PFL can be realized via mixture-of-experts, meta-learning, or parameter-efficient fine-tuning, with personalized parameters typically retained locally to limit exposure. This inherent tension between global generalization and per-organization specialization motivates our work to jointly balance privacy and personalization while *maximizing both global and local model utility.*

Parameter-Efficient Adaptation of LLMs. Training an LLM from scratch involves billions of parameters and is often unnecessary when adapting to a specific domain. Full fine-tuning can also drift from pretrained knowledge and is computationally expensive. Parameter-efficient fine-tuning (PEFT) methods address this by updating only small, task-specific components while keeping the base model frozen. A common approach is LoRA, which trains compact adapter matrices instead of all model weights, reducing computation, storage, and communication costs. Formally, for a pretrained weight matrix $w_0 \in \mathbb{R}^{d \times k}$, LoRA decomposes the weight update as $w = w_0 + \Delta w$, $\Delta w = BA$, $A \in \mathbb{R}^{r \times k}$, $B \in \mathbb{R}^{d \times r}$, $r \ll \min(d, k)$, where only A and B are trainable. The forward pass becomes $h = wx = w_0x + BAx$, preserving the pretrained base while adapting to task-specific knowledge. Distinct LoRA adapters can be learned for different tasks and selectively combined with the frozen model at inference, enabling efficient, modular personalization.

When LLMs meet FL. Leveraging advances in FL, PFL, and LoRA adapters, federated personalized fine-tuning lets organizations adapt pretrained LLMs to domain data without sharing raw corpora. In the PFL-for-LLMs setting each organization n keeps a global model w and a compact personalized adapter $\Delta w_{p,n}$, often parameterized with low-rank (LoRA) updates: $\Delta w_{p,n} \approx \sum_{\ell} B_{n,\ell} A_{n,\ell}$, where ℓ is an index that iterates over the layers (or parameterized components). Thus, communicating only the adapter parameters greatly reduces storage and bandwidth. However, directly applying FL/PFL to LLMs is problematic: *large models tend to memorize sensitive data (risking PII leakage), their checkpoints and updates are communication-intensive, and organization heterogeneity and regulatory constraints undermine naive FL/PFL averaging.* These limitations motivate SECUREGATE, designed to curb leakage while preserving global and per-organization utility.

3.2 Threat Model

We consider realistic adversaries in federated LLM deployments, where organizations retain raw data locally and share only model updates or adapter weights. Adversaries aim to extract sensitive information or degrade global/local model utility.

Roles, Capabilities, and Goals. We consider three principal adversary types: (i) *honest-but-curious server* — follows the protocol but inspects updates or checkpoints; (ii) *curious collaborating organizations* — analyze exchanged parameters to extract proprietary or PII information; (iii) *external adversaries* — eavesdrop on communications to perform membership inference. Adversaries may inspect updates or checkpoints, query models via APIs, and combine auxiliary public or leaked data with observations. Their goals are to reconstruct or extract private data (memorization/extraction attacks) and perform membership or attribute inference on organizational datasets.

We also assume adversaries do *not* have direct access to raw datasets but may observe model artifacts and query outputs.

4 Methodology

This section introduces SECUREGATE, a framework that addresses privacy leakage in federated LLMs through three components: *a learnable authorization token, a dual-adapter LoRA design, and a gating module for selective adapter activation.* The framework operates as follows (see Fig. 1):

4.1 Dual-Adapter Initialization for Privacy & Personalization

In this stage, SECUREGATE allows each organization $n \in \mathcal{N}$ to initialize and fine-tune dual LoRA adapters: (i) a secure adapter that generates representations sanitized of PII, and (ii) a revealing adapter for controlled disclosure (see Fig. 1.a). Let \mathcal{D}_n denotes organization n 's local dataset, which is partitioned into two views: a *raw view* $\mathcal{D}_n^{\text{raw}}$ and a *masked view* $\mathcal{D}_n^{\text{mask}}$, where PII is automatically detected (e.g., via named-entity recognition, NER (Nadeau and Sekine, 2007)) and replaced with [MASK] using a defense scheme (e.g., scrubbing (Pilán et al., 2022)). We define two parameter-efficient LoRA updates per organization:

$$\Delta \mathbf{w}_n^{(\text{rev})} = \mathbf{B}_n^{(\text{rev})} \mathbf{A}_n^{(\text{rev})}, \quad \Delta \mathbf{w}_n^{(\text{sec})} = \mathbf{B}_n^{(\text{sec})} \mathbf{A}_n^{(\text{sec})}, \quad (2)$$

where $\Delta \mathbf{w}_n^{(\text{rev})}$ and $\Delta \mathbf{w}_n^{(\text{sec})}$ are the revealing and secure adapters trained on $\mathcal{D}_n^{\text{raw}}$ and $\mathcal{D}_n^{\text{mask}}$, respectively.

During fine-tuning, only $\Delta \mathbf{w}_n^{(\text{sec})}$ is shared and aggregated by the server, ensuring the global model encodes only sanitized knowledge. In contrast, $\Delta \mathbf{w}_n^{(\text{rev})}$ remains local and, after fusion with the final global model, is accessed exclusively via token-gated control (Section 4.4), allowing only authorized users to access sensitive information.

4.2 Protected Federated LLM Fine-Tuning

We employ a two-level optimization scheme that alternates between a server-side *outer loop* and an organization-side *inner loop* (Fig. 1.b). The goal is to learn a global LoRA adapter that captures domain-generalizable knowledge while allowing each organization to maintain a personalized variant for local utility, *ensuring that no organization's PII is exposed to others or the server.*

At communication round $t \in \{1, \dots, T\}$, the server broadcasts the current global secure adapter $\Delta \mathbf{w}^t$. Each organization $n \in \mathcal{N}$ then performs K inner-loop updates on its masked dataset $\mathcal{D}_n^{\text{mask}}$ using AdamW (Loshchilov and Hutter, 2017), producing locally updated adapters $\Delta \mathbf{w}_n^{t+1}$, which are sent to the server. Next, the server aggregates these updates via weighted averaging with Nesterov momentum (Sutskever et al., 2013):

$$\Delta \bar{\mathbf{w}}^{t+1} = \sum_{n=1}^N \frac{|\mathcal{D}_n^{\text{mask}}|}{\sum_{k=1}^N |\mathcal{D}_k^{\text{mask}}|} \Delta \mathbf{w}_n^{t+1}, \quad (3)$$

This average is combined with a scaled momentum vector $\Delta \mathbf{v}^t$ to form a lookahead position as $\Delta \mathbf{p}^t = \Delta \mathbf{w}^t + m \Delta \mathbf{v}^t$, where m controls the momentum strength. The momentum is then updated by moving towards the difference between the weighted average and lookahead scaled by the learning rate η as:

$$\Delta \mathbf{v}^{t+1} = m \Delta \mathbf{v}^t + \eta (\Delta \bar{\mathbf{w}}^{t+1} - \Delta \mathbf{p}^t) \quad (4)$$

Finally, the server updates the global model with the new momentum:

$$\Delta \mathbf{w}^{t+1} = \Delta \mathbf{w}^t + \Delta \mathbf{v}^{t+1} \quad (5)$$

By restricting communication to lightweight secure adapters, this scheme significantly reduces communication overhead while preserving strong global generalization and enabling privacy-preserving local adaptation at each organization.

4.3 Dual Personalized Adapter Fusion

After completing federated fine-tuning (Section 4.2), each organization fuses the final global adapter $\Delta \mathbf{w}^T$ with its local secure and revealing

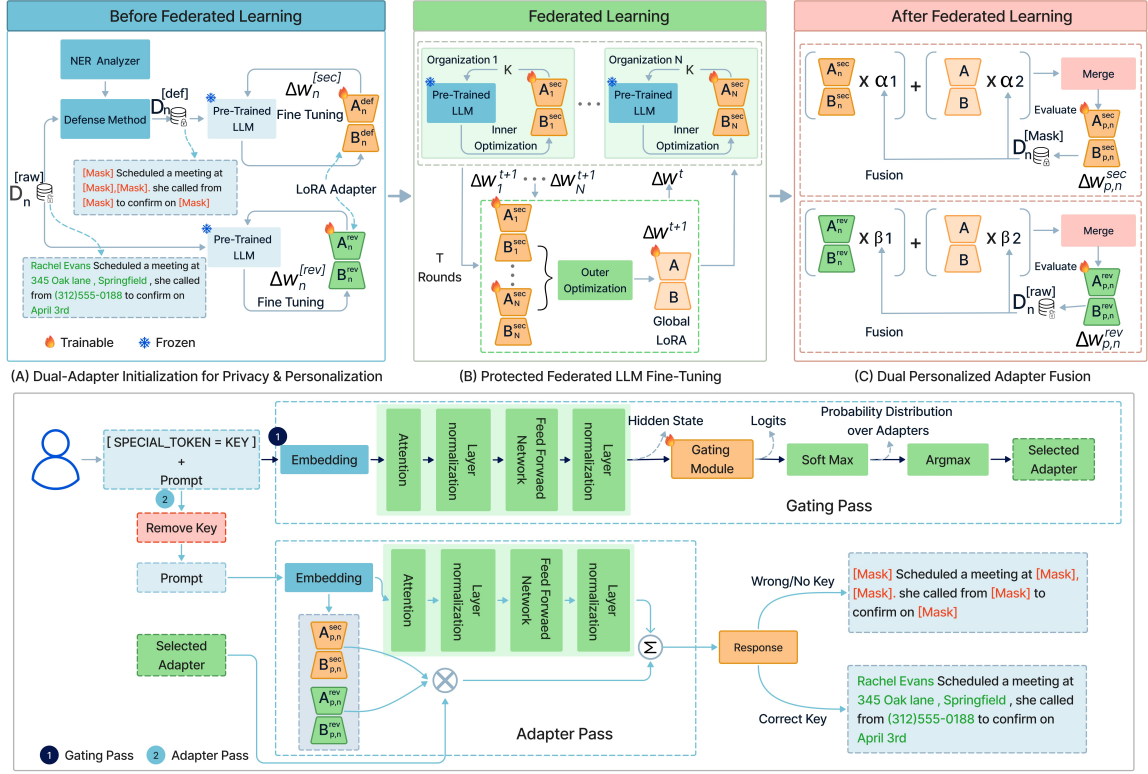


Figure 1: An Illustration of the SECUREGATE Framework.

adapters. This yields two organization-specific adapters: (i) a *secure personalized adapter* that preserves privacy while maintaining global utility, and (ii) a *revealing personalized adapter* that selectively discloses sensitive local knowledge to authorized users (see Fig. 1.c).

Secure Personalized Adapter. Let Δw_n^T denotes the locally trained secure adapter of organization n after T rounds. The secure personalized adapter is obtained via linear fusion:

$$\Delta w_{p,n}^{(\text{sec})} = \alpha_1 \Delta w^T + \alpha_2 \Delta w_n^T, \quad (6)$$

where α_1 and α_2 balance global generalization and local adaptation. Following (Huang et al., 2023), the fusion weights are learned via black-box optimization by minimizing the cross-entropy loss on a small query set $Q \subset \mathcal{D}_n^{\text{mask}}$ with ℓ_1 regularization:

$$\min_{\{\alpha_i\}} \mathcal{L}(\alpha; Q) + \psi \sum_i |\alpha_i|. \quad (7)$$

where ψ regulates the strength of regularization. Since Δw_n^T is trained on masked data (e.g., PII replaced with [MASK]), this fusion preserves privacy while integrating global knowledge.

Revealing Personalized Adapter. In parallel, each organization fuses Δw^T with its local revealing adapter $\Delta w_n^{(\text{rev})}$, trained on raw data:

$$\Delta w_{p,n}^{(\text{rev})} = \beta_1 \Delta w^T + \beta_2 \Delta w_n^{(\text{rev})}, \quad (8)$$

where β_1 and β_2 are optimized similar to α_i in Eq. (7). This adapter remains strictly local and is accessed only through token-controlled gating at inference time, enabling controlled disclosure of sensitive organization-specific knowledge.

Together, these adapters allow organizations to benefit from shared federated knowledge while preserving privacy and enabling fine-grained, authorization-based information access.

4.4 Gating Module Training

To enable controlled disclosure at inference time, we introduce a lightweight *gating module* that routes each query to either the *secure* or *revealing* personalized adapter (Section 4.3). Together with the dual adapters, the gating module forms the access-control core of SECUREGATE, enabling per-request, key-controlled disclosure while preserving shared model utility (see Fig. 1.d).

Design and Input Encoding. The gating module is a small MLP designed for minimal local overhead. Each query is prefixed with a learnable access key encoded as a special token: [SPECIAL_TOKEN=KEY] PROMPT. This key is added as an atomic token to an organization-local tokenizer, while a public tokenizer excludes it, preventing token splitting and leakage (see Appendix D). The gating module operates on the final-

layer hidden state h_{key} of the key token extracted from the frozen base model, which we find provides the most robust routing signal (see Appendix C.3). The gating logits are computed as

$$z_i = \text{MLP}(h_{\text{key}}), \quad i \in \mathcal{I}, \quad (9)$$

followed by a softmax: $p_i = \frac{\exp(z_i)}{\sum_{j \in \mathcal{I}} \exp(z_j)}$.

The adapter is selected as $a^* = \arg \max_i p_i$, with an optional threshold τ to default to the secure adapter when confidence is low. Notably, this design naturally extends to multiple adapters with distinct keys and access levels.

Gating Module Training. Each organization trains its gating module locally using a cross-entropy loss: $\mathcal{L}_{\text{gating}} = -\sum_i y_i \log p_i$, where y_i indicates the correct adapter for a given key. Training data is *synthetically* generated to cover authorized and unauthorized cases (details on synthetic data generation are provided in Appendix E). During training, the base model and adapters remain frozen, confining learning to the small gating network. The trained gating module remains local and can be efficiently retrained when keys are rotated, *without modifying the model or adapters*.

Two-Pass Inference. To prevent key contamination of outputs, SECUREGATE uses a two-pass inference procedure. In the *gating pass*, the full prompt (with key) is processed by the frozen base model to extract h_{key} and select an adapter. In the *adapter pass*, the key is removed and the cleaned prompt is re-evaluated using the selected adapter to generate the final response. This ensures reliable routing without key-induced bias in generation.

5 Performance Evaluation

5.1 Experimental Setup

Datasets. We evaluate SECUREGATE on two benchmark datasets: (i) *Yelp Reviews* (Zhang et al., 2015), a large collection of user-generated reviews with text, star ratings, and metadata such as timestamps and user/business IDs; and (ii) *ECHR* (Chalkidis et al., 2019), comprising legal case documents from the European Court of Human Rights containing personal information. We uniformly sample 10,000 instances from each dataset to create balanced subsets. Our FL setup includes 10, 20, or 30 organizations, each training on a distinct 10% partition of the sampled data.

LLM Models. We experiment with SOTA open-source LLMs Qwen/Qwen 3-1.7B (Yang et al.,

2025), Google’s Gemma 2-2b (Team, 2024) and Meta’s Llama 3.2-1B/3B (Dubey et al., 2024), fine-tuned via LoRA adapters.

Hyperparameters. For LoRA training, we set the rank $r \in \{4, 8, 12, 16\}$. Each client fine-tunes its local parameters for three epochs using AdamW with learning rate $\eta_{\text{local}} = 10^{-4}$. For global aggregation, we use $\eta_{\text{global}} = 0.01$ and $T = 20$ rounds. Remaining parameters and implementation specifics are detailed in Appendix A.

Attack Types. We evaluate privacy leakage under two threats: (i) *Extraction Attack*: The model is sampled to generate text, after which NER is applied to extract PII; leakage is measured by comparing extracted entities against a baseline model; (ii) *Inference Attack*: Given a candidate set of c PII values ($c=50$ unless stated otherwise), the attacker selects the value that minimizes model perplexity for a fixed prefix–suffix context, assuming access to an auxiliary set of 100 PII-containing statements.

Evaluation Metrics. We evaluate model performance and privacy leakage using standard metrics, including Precision, Recall, Perplexity, and computational efficiency (FLOPs); formal definitions are provided in Appendix A.2.

5.2 Experimental Results

SECUREGATE vs. Baselines. We compare SECUREGATE against standard FL baselines—FedAvg, FedAdagrad, FedAdam, FedYogi, and FedAvgM (Section 2.1)—under inference attack accuracy with 10 clients. We integrate SECUREGATE’s dual-adapter architecture and token-controlled gating into each baseline and evaluate the resulting *SecureGate-enhanced* methods (Table 1).

With valid tokens, SECUREGATE achieves an average inference accuracy of 25.20%, outperforming the strongest baseline (FedAvgM at 20.12%). Without authorization, it maintains low leakage (4.20%), comparable to restrictive baselines such as FedYogi (3.73%), while delivering substantially higher authorized utility. These results demonstrate that SECUREGATE consistently improves the privacy–utility trade-off across both classical and personalized FL settings.

Evaluating SECUREGATE’s PII-Aware Preservation. We evaluate SECUREGATE under extraction and inference attacks on ECHR and Yelp using Qwen 3-1.7B, Gemma 2-2B, and Llama 3.2-3B, comparing correct versus incorrect or missing tokens. As shown in Table 2, inference attacks with *valid tokens*, Gemma 2-2B achieves the high-

Table 1: Client-wise comparative analysis of inference attack accuracy (%): evaluating SECUREGATE in authorized vs. unauthorized scenarios across diverse FL environments on the ECHR dataset using the Llama-1B model.

Client	Authorized Access (Correct Token) ↑						Unauthorized Access (Wrong/No Token) ↓					
	FedAdagrad	FedAvg	FedAdam	FedYogi	FedAvgM	SecureGate	FedAdagrad	FedAvg	FedAdam	FedYogi	FedAvgM	SecureGate
1	14.10	20.00	26.74	30.12	24.14	35.96	4.82	4.60	8.60	6.67	3.16	3.61
2	15.00	14.10	22.08	22.50	20.48	25.64	4.65	5.81	5.75	4.88	5.75	5.56
3	9.76	11.76	20.24	21.59	20.45	26.74	2.33	1.15	8.14	1.15	1.12	4.65
4	11.39	16.46	17.65	21.25	18.82	21.79	3.75	6.82	5.75	3.49	2.17	4.94
5	20.25	21.43	16.47	13.10	19.10	29.87	7.23	4.55	6.32	2.20	2.11	2.22
6	16.46	19.32	19.54	19.51	18.29	25.00	2.67	6.98	4.35	3.45	6.74	3.80
7	7.41	13.19	10.71	18.60	15.56	28.05	4.71	7.45	3.45	2.13	3.09	5.75
8	11.76	17.86	21.33	12.50	15.12	29.87	6.25	4.55	4.55	1.11	1.05	7.79
9	11.39	16.85	22.22	20.93	20.69	32.10	6.33	5.38	5.21	4.60	7.69	1.22
10	16.46	18.99	14.81	20.00	28.57	35.90	2.44	4.65	5.56	7.61	6.25	2.44
Average	13.40	17.00	19.18	20.01	20.12	25.20	4.52	5.19	5.77	3.73	3.92	4.20

Table 2: SECUREGATE robustness: inference and extraction attacks across models, datasets, and access scenarios.

Model Client	Inference Accuracy (%)				Extraction Metrics (Precision / Recall) %								
	ECHR Dataset		Yelp Dataset		ECHR Dataset				Yelp Dataset				
	Authorized ↑ (Correct)	Unauthorized ↓ (Wrong/No)	Authorized ↑ (Correct)	Unauthorized ↓ (Wrong/No)	Authorized ↑ Prec	Unauthorized ↓ Rec	Authorized ↑ Prec	Unauthorized ↓ Rec	Authorized ↑ Prec	Unauthorized ↓ Rec	Authorized ↑ Prec	Unauthorized ↓ Rec	
Qwen 3-1.7B	1	22.08	4.94	26.74	4.35	1.78	17.32	10.00	4.25	2.85	6.80	0.00	0.00
	2	18.07	2.63	17.05	10.39	1.57	19.15	7.71	4.96	3.77	7.50	2.05	0.97
	3	20.00	2.56	29.11	3.95	2.53	12.70	5.21	2.48	4.29	9.12	0.70	0.32
	4	20.99	5.26	24.71	8.89	2.34	14.79	3.28	1.43	4.24	9.12	1.22	0.38
	5	19.78	5.68	26.83	9.72	2.38	17.21	12.14	6.82	3.36	7.78	0.85	0.38
	6	25.00	4.94	20.24	5.00	2.55	16.20	5.99	3.21	3.95	11.04	0.68	0.37
	7	22.99	2.33	21.43	4.76	2.14	18.87	6.85	4.76	3.65	7.95	1.44	0.74
	8	27.85	2.60	28.21	3.90	2.44	15.56	4.32	2.74	5.66	9.67	0.99	0.36
	9	20.48	4.40	25.30	8.00	1.74	15.38	6.40	3.78	5.17	11.51	0.93	0.36
	10	23.75	1.32	17.28	6.76	2.31	12.96	5.94	4.63	2.82	8.26	0.40	0.46
Gemma 2-2B	1	73.26	6.02	75.00	2.41	0.00	0.00	0.00	0.00	0.27	0.97	0.29	0.97
	2	72.09	4.76	72.00	2.11	0.00	0.00	0.00	0.00	0.12	0.42	0.00	0.00
	3	57.32	4.65	71.83	3.75	0.13	0.15	0.00	0.00	0.03	0.33	0.04	0.33
	4	70.89	3.75	79.47	7.95	0.00	0.00	0.00	0.00	0.03	0.35	0.00	0.00
	5	83.95	3.26	68.18	7.41	0.00	0.00	0.00	0.00	0.10	0.39	0.07	0.39
	6	61.54	4.88	75.95	3.41	0.00	0.00	0.00	0.00	0.22	0.65	0.00	0.00
	7	74.16	2.25	73.61	3.66	0.44	0.18	0.00	0.00	0.17	1.52	0.05	0.38
	8	65.38	7.23	74.65	3.61	0.00	0.00	0.00	0.00	0.08	0.37	0.17	0.37
	9	64.10	6.33	71.43	3.75	0.00	0.00	0.00	0.00	0.15	0.36	0.00	0.00
	10	74.03	5.63	77.27	2.44	0.00	0.00	0.00	0.00	0.00	0.00	0.10	0.46
Llama 3.2-3B	1	44.87	7.58	50.00	5.88	3.99	12.25	8.33	0.49	3.45	8.74	0.00	0.00
	2	45.33	7.58	46.03	6.67	3.07	12.48	3.00	0.51	2.11	11.25	4.55	0.42
	3	48.68	7.46	48.48	6.56	3.30	10.07	0.00	0.00	2.19	12.38	0.00	0.00
	4	43.75	10.34	60.56	2.94	3.66	11.92	1.82	0.16	3.41	8.77	8.82	1.05
	5	51.35	2.60	51.61	12.77	2.70	16.56	9.52	0.97	2.13	10.12	0.00	0.00
	6	46.58	7.81	49.25	8.20	3.90	13.28	4.76	0.44	3.14	11.69	8.70	0.65
	7	36.00	1.49	57.38	14.29	3.48	15.34	0.00	0.00	1.70	12.88	1.92	0.38
	8	34.38	1.89	60.61	9.09	3.77	12.82	4.44	0.68	3.04	10.04	8.33	1.12
	9	44.59	6.06	43.40	4.76	3.81	11.19	4.55	0.42	3.79	17.63	3.85	0.36
	10	43.48	4.29	48.21	6.67	3.15	11.88	4.17	0.31	1.92	9.17	3.12	0.46

est inference accuracy (57.32–83.95% on ECHR; 68.18–79.47% on Yelp), followed by Llama 3.2-3B and Qwen3-1.7B. This highlights that while Gemma resists explicit memorization, it remains susceptible to probabilistic, ranking-based inference. In *unauthorized settings*, inference leakage remains low across all models (1.32–10.34% on ECHR; 2.11–14.29% on Yelp). Notably, authorization removal causes sharp performance drops—for instance, Llama 3.2-3B (Client 5, ECHR) falls from 51.35% to 2.60% (19.75× reduction), and Gemma 2-2B (Client 10, Yelp) from 77.27% to

2.44% (31.66× reduction).

For extraction attacks with *valid tokens*, all models exhibit low precision. Qwen 3-1.7B attains precision of 1.57–2.55% (ECHR) and 2.82–5.66% (Yelp), while Llama 3.2-3B ranges from 2.70–3.99% (ECHR) and 1.70–3.79% (Yelp). Corresponding recall values remain modest, peaking at 19.15% (Qwen 3-1.7B, ECHR). Gemma 2-2B shows the lowest susceptibility to verbatim extraction, consistent with its memorization-mitigation training (Peng et al., 2024), though it remains vulnerable to probabilistic, ranking-based inference.

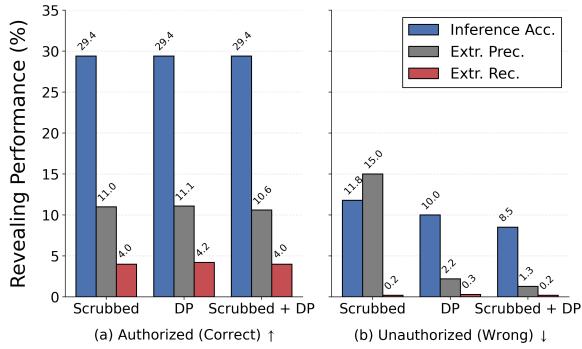


Figure 2: Gating module performance in routing queries to the correct adapter under inference and extraction attacks across various defenses on Llama-1B.

When *tokens are incorrect or absent*, all models demonstrate strong resistance to extraction, with recall dropping near zero. For instance, Llama 3.2-3B recall decreases to at most 1.12% on Yelp, representing up to a 17.07 \times reduction relative to authorized access. Similar trends hold for Qwen 3-1.7B, proving that the gating module reliably suppresses sensitive disclosures without authorization.

Overall, these results show that SECUREGATE enforces access-aware privacy: sensitive information is revealed only with valid authorization, while incorrect or missing tokens reliably trigger a secure state. Candidate pool scaling ($c = 50, 100, 500$) is analyzed in Appendix B.5.

Evaluating the Effectiveness of SECUREGATE’s Gating Model Across Defenses. We evaluate SECUREGATE per client under multiple defense strategies (data scrubbing, DP, and scrubbing+DP). As shown in Fig. 2a, with a *valid authorization token*, the gating module achieves 100% routing accuracy to the revealing adapter, yielding a consistent 29.41% revealing performance across all attacks, since defenses are not activated.

For *unauthorized queries*, the gating module reliably routes requests to the secure adapter. As illustrated in Fig. 2b, the hybrid SCRUBBED+DP defense offers the strongest protection, reducing extraction precision to 1.30% (11.54 \times lower than scrubbing alone at 15.00%) and inference accuracy to 8.51% (3.46 \times lower than the authorized baseline). The perfect routing accuracy confirms that any residual leakage stems from the defense mechanisms themselves rather than failures in token-based gating. Overall, SECUREGATE effectively integrates layered defenses to achieve strong privacy guarantees while preserving authorized utility.

Multi-Adapter Evaluation: Authorization, Isolation, and Privacy Floors To assess how well

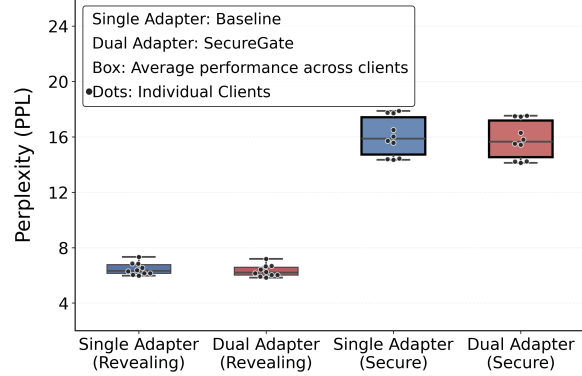


Figure 3: PPL across ten clients for Llama-1B, showing that SECUREGATE matches the utility of standalone secure and revealing adapter baselines.

SECUREGATE manages dataset-specific and role-based permissions, we tested a configuration featuring several “revealing” adapters alongside a single “secure” adapter. In this Llama-1B setup, each client hosts four specialized heads: one adapter fine-tuned on ECHR, one on Yelp, a joint adapter trained on the union of ECHR and Yelp (“E + Y”), and a secure adapter trained on a scrubbed, combined corpus where PII has been masked. Each adapter is bound to a distinct keyed token, and the router is trained to map every key deterministically to its corresponding adapter, thereby implementing per-key access control over personalization paths.

Table 3 summarizes these outcomes, showing that the gating module achieves 100% routing reliability by successfully directing all requests to the authorized adapter path without collision. Crucially, the results reveal a significant *Authorization Gap*; for the ECHR dataset, an attacker with valid credentials achieves an inference accuracy of 20.00%, which drops to a *Privacy Floor* of 2.44% when no valid Key is provided. This represents an 8.20 \times reduction in PII leakage for unauthorized users. There is also clear *cross-dataset isolation*: using the Yelp key on ECHR candidates reaches just 4.94% accuracy, and using incorrect or missing keys on Yelp tops out at 6.76%, indicating that miskeyed requests expose only a small fraction of the sensitive information that is available under the correct key. Finally, the joint “E + Y” adapter attains 13.41% on ECHR—lower than the 20.00% of the dedicated ECHR adapter—which suggests that mixing heterogeneous sources during training can dilute dataset-specific memorization and slightly dampen leakage.

Perplexity Analysis and Utility Preservation. Fig. 3 compares SECUREGATE’s dual-adapter de-

Table 3: Inference attack accuracy on a single client using Llama-1B. The diagonal (light blue) demonstrates perfect adapter routing under valid authorization keys.

Dataset	Auth Key Used	Inference Accuracy \uparrow	Adapter Selected (Routing Result)			
			E + Y	ECHR	Yelp	Masked
ECHR	ECHR & Yelp Key	13.41	100.00	0.00	0.00	0.00
	ECHR Key	20.00	0.00	100.00	0.00	0.00
	Yelp Key	4.94	0.00	0.00	100.00	0.00
	Wrong/No Key	2.44	0.00	0.00	0.00	100.00
Yelp	ECHR & Yelp Key	12.20	100.00	0.00	0.00	0.00
	ECHR Key	5.19	0.00	100.00	0.00	0.00
	Yelp Key	13.92	0.00	0.00	100.00	0.00
	Wrong/No Key	6.76	0.00	0.00	0.00	100.00

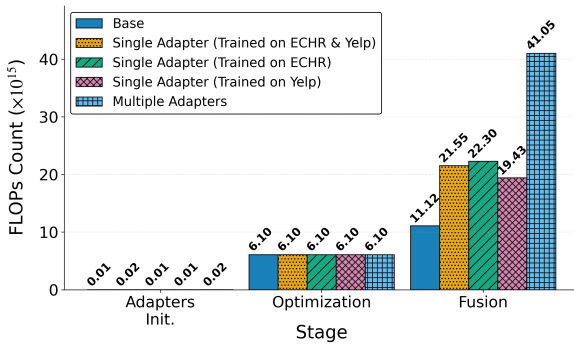


Figure 4: Client-side computational cost across the three Federated LLM fine-tuning stages (initialization, optimization, fusion), showing that multi-adapter configurations increase overhead only during the fusion stage.

sign with single-adapter baselines (secure or revealing) on Llama-1B using the ECHR dataset. With a valid token, SECUREGATE achieves an average Perplexity (PPL) of 6.32, matching the revealing-adapter baseline (6.46). Under unauthorized access (missing or incorrect token), it reliably switches to the secure adapter, yielding a PPL of 15.89, close to the secure-adapter baseline of 16.03. Detailed results in Table 9 (Appendix C.2) confirm that the gating mechanism preserves *personalized model utility* while enforcing privacy constraints.

Computational Cost Analysis. Fig. 4 reports the phase-wise FLOPs across four configurations. The optimization phase remains constant at 6.10×10^{15} FLOPs, while the fusion phase scales with the number of revealing adapters, increasing from 11.12×10^{15} to 41.05×10^{15} FLOPs (3.69 \times). This shows that computational overhead is dominated by local adapter fusion rather than server-side training or communication. A detailed configuration and communication cost is provided in Appendix B.3.

Gating Overhead Analysis We evaluate the computational overhead introduced by the gating mechanism in SECUREGATE compared to a single-

Table 4: Gating overhead analysis: Llama 1B gated adapters vs single-adapter baseline (ECHR dataset)

	Baseline		SECUREGATE (Ours)	
	Secure	Revealing	Secure	Revealing
Time (ms) (\downarrow)				
Gating Pass	–	–	197.14	197.14
Adapter Pass	1501.71	1463.55	1501.71	1463.55
Gating Overhead (\times)	1 \times	1 \times	1.13 \times	1.13 \times

adapter baseline. As shown in Table 4, the baseline performs only the adapter pass, whereas SECUREGATE introduces an additional gating pass prior to adapter execution.

The gating pass incurs a fixed cost of 197.14 ms, while the adapter pass remains unchanged across both baseline and SECUREGATE configurations. This results in a modest overhead of approximately 1.13 \times per inference, consistent across both secure and revealing settings.

Notably, the gating computation is lightweight compared to the adapter pass, contributing only a small fraction of the total latency. These results demonstrate that SECUREGATE achieves controlled and predictable overhead, making it practical for deployment in real-world systems where both privacy guarantees and efficiency are critical.

6 Conclusion

We propose SECUREGATE, a privacy-aware federated fine-tuning framework for LLMs that ensures strong privacy, high utility, and efficiency. It combines a dual-adapter architecture with token-controlled gating to enable fine-grained personalization while protecting sensitive PII from inference and extraction attacks. Extensive evaluations show that SECUREGATE achieves strong privacy-utility trade-offs across heterogeneous settings with minimal overhead, offering a scalable and practical approach to secure federated LLMs.

7 Limitations

Despite the effectiveness of SECUREGATE, several practical limitations remain. The additional “gating pass” executed prior to adapter inference increases the token prefill latency, which may affect deployment in latency-sensitive, real-time applications. Maintaining multiple active LoRA adapters (secure, revealing, and global) also incurs additional memory overhead compared to single-adapter personalization, potentially limiting adoption on resource-constrained edge devices. Although SECUREGATE achieves 100% routing reliability in our experiments, its security ultimately depends on the secrecy and uniqueness of key-token representations, motivating future work on stronger cryptographic embedding mechanisms to further harden access control at scale. Finally, our evaluation focuses on decoder-only architectures (e.g., Llama, Qwen, Gemma), and extending the proposed dual-adapter design to substantially larger models or alternative architectures remains an open research direction.

References

- Alan Akbik, Tanja Bergmann, Duncan Blythe, Kashif Rasul, Stefan Schweter, and Roland Vollgraf. 2019. Flair: An easy-to-use framework for state-of-the-art nlp. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics (demonstrations)*, pages 54–59.
- Atilla Akkus, Masoud Poorghaffar Aghdam, Mingjie Li, Junjie Chu, Michael Backes, Yuyang Zhang, and Sinem Sav. 2025. Generated data with fake privacy: Hidden dangers of fine-tuning large language models on generated data. In *34th USENIX Security Symposium (USENIX Security 25)*, pages 8075–8093.
- Youssef Allouah, Rachid Guerraoui, and John Stephan. 2025. Towards trustworthy federated learning with untrusted participants. In *Forty-second International Conference on Machine Learning*.
- Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, and 1 others. 2021. Extracting training data from large language models. In *30th USENIX security symposium (USENIX Security 21)*, pages 2633–2650.
- Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. 2019. Neural legal judgment prediction in english. *arXiv preprint arXiv:1906.02059*.
- Shuai Cheng and 1 others. 2025. [Understanding pii leakage in large language models: A systematic survey](#). In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pages 10409–10417. International Joint Conferences on Artificial Intelligence Organization. Survey Track.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv e-prints*, pages arXiv–2407.
- Muneeb Ul Hassan, Mubashir Husain Rehmani, and Jinjun Chen. 2019. Differential privacy techniques for cyber physical systems: A survey. *IEEE Communications Surveys & Tutorials*, 22(1):746–789.
- Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. 2019. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Chengsong Huang, Qian Liu, Bill Yuchen Lin, Tianyu Pang, Chao Du, and Min Lin. 2023. Lorahub: Efficient cross-task generalization via dynamic lora composition. *arXiv preprint arXiv:2307.13269*.
- Meirui Jiang and 1 others. 2024. Heterogeneous personalized federated learning by local-global updates mixing via convergence rate. In *The Twelfth International Conference on Learning Representations*.
- Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, and 1 others. 2021. Advances and open problems in federated learning. *Foundations and trends® in machine learning*, 14(1–2):1–210.
- Bill Yuchen Lin and 1 others. 2022. [FedNLP: Benchmarking federated learning methods for natural language processing tasks](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 157–175, Seattle, United States. Association for Computational Linguistics.
- Chen Ling, Xujiang Zhao, Jiaying Lu, Chengyuan Deng, Can Zheng, Junxiang Wang, Tanmoy Chowdhury, Yun Li, Hejie Cui, Xuchao Zhang, and 1 others. 2023. Domain specialization as the key to make large language models disruptive: A comprehensive survey. *ACM Computing Surveys*.
- Bingyan Liu, Nuoyan Lv, Yuanchun Guo, and Yawen Li. 2024. Recent advances on federated learning: A systematic survey. *Neurocomputing*, 597:128019.
- Guodong Long, Tao Shen, Jing Jiang, Michael Blumenstein, and 1 others. 2024. Dual-personalizing adapter for federated foundation models. *Advances in Neural Information Processing Systems*, 37:39409–39433.

- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Wei Lu, Rachel K Luu, and Markus J Buehler. 2025. Fine-tuning large language models for domain adaptation: Exploration of training strategies, scaling, model merging and synergistic capabilities. *npj Computational Materials*, 11(1):84.
- Nils Lukas, Ahmed Salem, Robert Sim, Shruti Tople, Lukas Wutschitz, and Santiago Zanella-Béguelin. 2023. Analyzing leakage of personally identifiable information in language models. In *2023 IEEE Symposium on Security and Privacy (SP)*, pages 346–363. IEEE.
- Mohamad Mansouri, Melek Önen, Wafa Ben Jaballah, and Mauro Conti. 2023. Sok: Secure aggregation based on cryptographic schemes for federated learning. *Proceedings on Privacy Enhancing Technologies*.
- Brendan McMahan and 1 others. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR.
- David Nadeau and Satoshi Sekine. 2007. A survey of named entity recognition and classification. *Linguistic Investigations*, 30(1):3–26.
- Venkatesh Balavadhani Parthasarathy, Ahtsham Zafar, Aafaq Khan, and Arsalan Shahid. 2024. The ultimate guide to fine-tuning llms from basics to breakthroughs: An exhaustive review of technologies, research, best practices, applied research challenges and opportunities. *arXiv preprint arXiv:2408.13296*.
- Yuefeng Peng, Junda Wang, Hong Yu, and Amir Houmansadr. 2024. Data extraction attacks in retrieval-augmented generation via backdoors. *arXiv preprint arXiv:2411.01705*.
- Ildikó Pilán, Pierre Lison, Lilja Øvrelid, Anthi Papadopoulou, David Sánchez, and Montserrat Batet. 2022. The text anonymization benchmark (tab): A dedicated corpus and evaluation framework for text anonymization. *Computational Linguistics*, 48(4):1053–1101.
- Samar Pratap and 1 others. 2025. [The fine art of fine-tuning: A structured review of advanced llm fine-tuning techniques](#). *Natural Language Processing Journal*, 11:100144.
- Lang Pu, Jingjing Gu, Chao Lin, and Xinyi Huang. 2025. Janus: Dual-server multi-round secure aggregation with verifiability for federated learning. In *Forty-second International Conference on Machine Learning*.
- Jiaxing Qi, Zhongzhi Luan, Shaohan Huang, Carol Fung, Hailong Yang, and Depei Qian. 2024. Fd-lora: Personalized federated learning of large language model via dual lora tuning. *arXiv preprint arXiv:2406.07925*.
- Sashank Reddi and 1 others. 2020. Adaptive federated optimization. *arXiv preprint arXiv:2003.00295*.
- Lei Shen, Zhenheng Tang, Lijun Wu, Yonggang Zhang, Xiaowen Chu, Tao Qin, and Bo Han. Hot-pluggable federated learning: Bridging general and personalized fl via dynamic selection. In *The Thirteenth International Conference on Learning Representations*.
- Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. 2013. On the importance of initialization and momentum in deep learning. In *International conference on machine learning*, pages 1139–1147. pmlr.
- Gemma Team. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Xinghao Wu, Xuefeng Liu, Jianwei Niu, Haolin Wang, Shaojie Tang, and Guogang Zhu. 2024. Fedlora: When personalized federated learning meets low-rank adaptation.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, and et al. 2025. Qwen3: Large language models with hybrid thinking modes and extensive multilingual support. *arXiv preprint arXiv:2505.09388*.
- Rui Ye, Wenhao Wang, Jingyi Chai, Dihan Li, Zexi Li, Yinda Xu, Yaxin Du, Yanfeng Wang, and Siheng Chen. 2024. Openfedllm: Training large language models on decentralized private data via federated learning. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 6137–6147.
- Liping Yi and 1 others. 2023. pfedlora: Model-heterogeneous personalized federated learning with lora tuning. *arXiv preprint arXiv:2310.13283*.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28.
- Zeyuan Zuo, Ningxin Su, Baochun Li, and Teng Zhang. 2024. Pack: Towards communication-efficient homomorphic encryption in federated learning. In *Proceedings of the 2024 ACM Symposium on Cloud Computing*, pages 470–486.

Appendices

A Detailed Experimental Settings

A.1 Reproduction Hyperparameters

To ensure reproducibility, all LoRA adapters were configured with a scaling factor of $4 \times r$ and a dropout rate of 0.1. Weight decay was set to $\gamma = 0.001$, and global aggregation applied a momentum of $m = 0.5$. Personalized adapter fusion used a regularization factor $\psi = 0.01$. All other experimental parameters follow (Lukas et al., 2023).

A.2 Evaluation Metrics

We describe the evaluation metrics used to assess privacy, utility, and computational efficiency as:

- **Precision:** Fraction of generated PII that matches the training data, indicating actual leakage.
- **Recall:** Fraction of training PII successfully reproduced, reflecting potential leakage risk.
- **FLOPs:** Total floating-point operations, independent of hardware; lower values indicate higher computational efficiency.
- **Perplexity (PPL):** Measures model uncertainty in predicting text; lower values indicate better predictive performance:

$$\text{PPL} = \exp \left(-\frac{1}{H} \sum_{i=1}^H \log P(\kappa_i | \kappa_1, \dots, \kappa_{i-1}) \right),$$

where H is the total number of tokens and $P(\kappa_i | \kappa_1, \dots, \kappa_{i-1})$ is the predicted probability of token κ_i .

A.3 Gating Module Architecture

The design of our SECUREGATE’s gating module comprises:

- A linear layer with GLU activation, followed by layer normalization and dropout.
- A second linear layer with GELU activation, layer normalization, and dropout for feature refinement.
- A final linear layer producing logits over available adapters.

GLU dynamically controls feature flow for routing, while GELU introduces smooth non-linear transformations to stabilize training. Together, they enable flexible and efficient adapter selection.

Table 5: Inference attack accuracy (%) for Llama-1B across non-IID clients (Clients 1–5: ECHR, Clients 6–10: Yelp) under correct-token and wrong/no-token settings.

Client #	Training Dataset	Correct Token	Wrong/No Token
1	ECHR	36.00	5.33
2	ECHR	33.33	4.48
3	ECHR	31.33	6.10
4	ECHR	22.37	10.67
5	ECHR	31.65	4.94
6	Yelp	17.91	1.59
7	Yelp	32.00	6.94
8	Yelp	24.68	7.79
9	Yelp	18.31	3.23
10	Yelp	12.86	6.06

A.4 Computing Environment

All experiments ran on Ubuntu 24.04.2 LTS with an Intel Core i9-14900K CPU @ 3.2GHz and an NVIDIA RTX PRO 6000 Blackwell GPU (96GB VRAM). This configuration supports concurrent loading of multiple LoRA adapters in SECUREGATE.

A.5 Implementation Details

The framework is implemented using OpenFedLLM (Ye et al., 2024). Named Entity Recognition (NER) employs the Flair library (Akbiik et al., 2019), using hyperparameters consistent with (Lukas et al., 2023) to ensure comparability and reproducibility.

B In-Depth Evaluation of SECUREGATE

B.1 Inference Performance under Non-IID Client Distributions

Table 5 reports inference attack accuracy (%) for Llama-1B across *ten non-IID* clients (organizations), where clients 1–5 are trained on ECHR and clients 6–10 on Yelp, explicitly modeling a heterogeneous federated LLMs setting.

Under *authorized access* (correct token), ECHR clients exhibit consistently higher revealing percentage (inference success), reaching up to 36.00% (Client 1) and remaining above 30% for several clients (Clients 2, 3, and 5). In contrast, most Yelp clients achieve lower accuracy, with some below 25% (e.g., 17.91% for Client 6 and 12.86% for Client 10). This suggests that models trained on structured legal text are more vulnerable to inference attacks than those trained on informal review data. Nevertheless, notable attack success on some Yelp clients (e.g., 32.00% for Client 7 and 24.68%

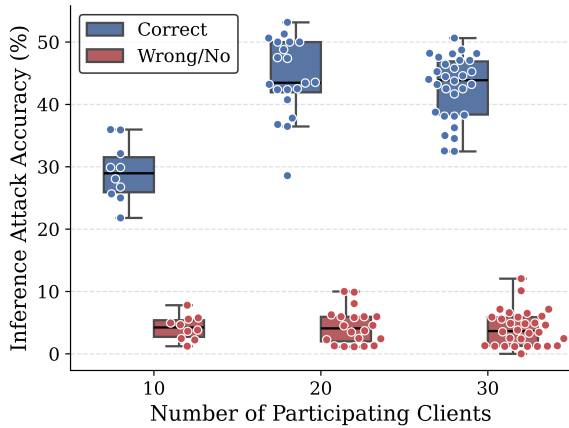


Figure 5: Scalability analysis of inference attack accuracy for Llama 3.2-1B across 10, 20, and 30 clients. The gating mechanism successfully decouples authorized utility from unauthorized leakage, maintaining a stable privacy boundary ($\approx 4.2\%$) as the network scales.


for Client 8) confirms that data heterogeneity does not eliminate the threat but instead induces client-specific vulnerability profiles.

Under *unauthorized access* (wrong or missing token), inference accuracy drops sharply across all clients, typically to the 1–11% range (e.g., 5.33% for Client 1, 10.67% for Client 4, and 1.59% for Client 6), demonstrating the effectiveness of token-controlled gating in suppressing attacks by routing queries to the secure adapter. A small number of outliers (e.g., Client 4 at 10.67% and Client 8 at 7.79%) still exhibit residual leakage, indicating that the privacy strength ultimately depends on the underlying defense used in the secure adapter. These results underscore the need for stronger or client-adaptive defenses to ensure uniform privacy guarantees under highly non-IID data distributions.

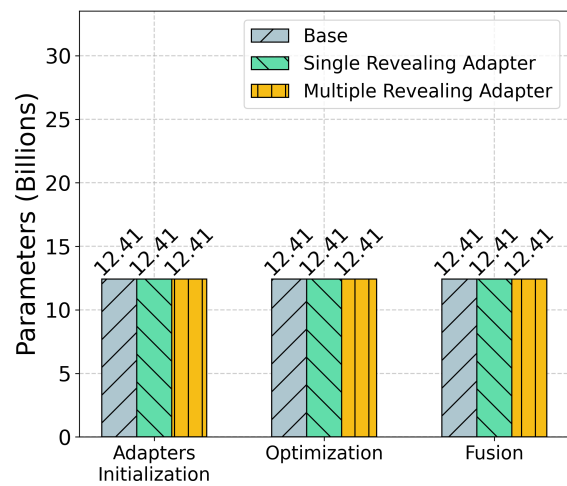
B.2 Evaluating Scalability Across Client Distributions

To assess the scalability and robustness of our SECUREGATE, Fig. 5 reports inference attack accuracy for the Llama 3.2–1B model under increasing numbers of federated clients. Under authorized access (correct token), the mean attack success rate rises from 29.09% ($\sigma = 4.62$) with 10 clients to 44.32% ($\sigma = 6.23$) with 20 clients, reflecting improved model utility as additional client data contributes to fine-tuning.

In contrast, under unauthorized access, leakage remains consistently low and largely invariant to network scale. Mean attack success rates are 4.20% ($\sigma = 1.95$) for 10 clients, 4.36% ($\sigma = 2.83$) for 20



(a) Trainable Parameters



(b) Non-Trainable Parameters

Figure 6: Comparative analysis of trainable vs. non-trainable parameter counts (billions) across framework stages. The analysis demonstrates the minimal footprint of authorization-specific adapters relative to the frozen base model. Values are rounded to two decimals.

clients, and 4.11% ($\sigma = 2.86$) for 30 clients. This stability demonstrates that SECUREGATE maintains strong privacy guarantees as the federated network scales, ensuring that increased client participation does not introduce additional privacy risk.

In summary, these results showcase two core strengths of SECUREGATE. Authorized users with correct tokens retain full access to personalized knowledge, while any unauthorized or miskeyed requests are automatically isolated within a secure adapter that minimizes PII exposure. This per-request gating provides a robust, auditable access control framework that protects sensitive data without degrading legitimate model utility.

Table 6: Performance comparison across PII classes for correct and wrong/no token cases, along with adapter selection ratios, on inference attack accuracy (%) for Llama-3B on the Yelp dataset.

PII Class	Correct Token (\uparrow)		Wrong/No Token (\downarrow)	
	Acc. (%)	Reveal Adapter	Acc. (%)	Secure Adapter
PERSON	51.72	100.0%	5.36	100.0%
LOC	77.14	100.0%	24.24	100.0%
ORG	84.91	100.0%	10.71	100.0%
PRODUCT	85.71	100.0%	6.90	100.0%
Mean	74.87	100.0%	11.80	100.0%

B.3 Evaluating SECUREGATE Communication Overhead

To further evaluate the overhead of the proposed authorization mechanism, we analyze the parameter distribution across the operational phases of SECUREGATE using the LLaMA 3.2-1B model. This analysis builds directly upon the previously described authorization experiment (see Section 5.2) conducted on the ECHR and Yelp datasets, comparing base performance against configurations with single or multiple revealing adapters. As shown in Figure 6, the frozen parameters remain constant at 12.41 billion regardless of setup, while trainable parameters rise from 0.06 billion to 0.23 billion during the initialization and fusion phases when scaling to multiple revealing adapters, a $3.83\times$ increase in local capacity. Crucially, the optimization phase remains unaffected by these configurations, maintaining a minimal trainable footprint of 0.06 billion. By confining the training of multiple personalized adapters to the client side, the mechanism minimizes communication overhead and prevents interference with global training.

B.4 Inference Attack Performance by PII Class

Table 6 reports inference attack accuracy of Llama-3B on the Yelp dataset across different PII classes under Correct and Wrong/No Tokens. When provided the correct token, attack accuracy is high for all PII categories (PERSON, LOC, ORG, PRODUCT), with a mean accuracy of 74.87%, indicating that the model reliably reveals targeted PII when properly authorized. Under Wrong/No Tokens, accuracy drops to the 5.36%–24.24% range with a mean of 11.80%, showing that the defended adapter substantially suppresses unauthorized PII extraction, though location entities (LOC) remain

Table 7: Inference attack accuracy (%) for Llama 3.2-3B on ECHR (Date class) across different candidate sizes (c) for a single client.

Token Condition	Candidate Pool Size (c)		
	$c = 50$	$c = 100$	$c = 500$
Correct Token (\uparrow)	15.85	10.10	7.45
Wrong/No Token (\downarrow)	10.84	8.33	4.49

relatively more vulnerable than other PII types. Notably, our routing mechanism achieves a 100.0% selection rate for the intended adapter (Reveal or Secure) across all PII classes. This perfect adapter selection confirms that the residual inference accuracy observed in categories such as LOC (24.24%) is not due to routing failures, but rather the inherent difficulty of masking location-based context even within the secure state.

B.5 Effect of Candidate Pool Size on Inference Accuracy

Table 7 illustrates the relationship between the number of candidates (c) and inference attack accuracy for the ECHR Date class. As expected, increasing the candidate pool size from $c = 50$ to $c = 500$ results in a consistent reduction in accuracy for both token conditions. While correct authorization provides higher accuracy across all scales, the unauthorized accuracy (Wrong/No Token) drops significantly as the search space expands, reaching its lowest point (4.49%) at $c = 500$. This trend highlights that the defense becomes increasingly robust against brute-force or high-entropy inference attacks as the candidate pool grows.

C Ablation Study of SECUREGATE Components

C.1 Evaluating LoRA Rank Impact

Table 8 evaluates the impact of LoRA ranks ($r = 4, 8, 12, 16$) on the robustness of SECUREGATE against inference attacks using the Llama-1B model. When the correct token is utilized, inference accuracy generally increases with higher LoRA ranks, peaking at $r = 16$ for most clients, such as Client 10 reaching 35.90%. Notably, Client 1 achieves the highest overall authorized disclosure of 40.48% at $r = 12$. In unauthorized scenarios using wrong/no tokens, leakage remains consistently low, typically staying below 10% across all ranks.

Table 8: Inference attack accuracy (%) for Llama-1B on the ECHR dataset under different LoRA ranks (r) and numbers of clients.

Client #	Correct Token				Wrong/No Token			
	$r = 4$	$r = 8$	$r = 12$	$r = 16$	$r = 4$	$r = 8$	$r = 12$	$r = 16$
1	25.64	27.16	40.48	35.96	3.66	4.71	1.52	3.61
2	18.42	24.64	20.25	25.64	1.64	0.00	4.84	5.56
3	15.28	20.97	22.50	26.74	1.35	2.47	4.41	4.65
4	17.33	22.03	26.32	21.79	11.39	6.49	6.35	4.94
5	27.27	25.76	25.00	29.87	5.00	2.50	6.94	2.22
6	17.11	18.31	17.86	25.00	2.94	9.09	6.56	3.80
7	17.11	20.83	21.43	28.05	3.53	3.66	0.00	5.75
8	20.29	22.73	27.16	29.87	1.59	2.70	1.72	7.79
9	22.22	36.23	30.38	32.10	1.37	0.00	3.03	1.22
10	25.97	33.33	22.35	35.90	2.50	2.47	1.43	2.44

Table 9: Client-wise Perplexity (PPL) comparison between SECUREGATE and standalone revealing or secure adapter baselines. Values are rounded to two decimals.

Client	Standalone Baselines		SECUREGATE (Ours)	
	Revealing	Secure	Correct Token	Wrong/No Token
1	6.86	16.02	6.69	15.81
2	6.03	17.71	5.92	17.49
3	5.97	15.57	5.83	15.43
4	6.55	15.71	6.42	15.50
5	6.29	14.34	6.16	14.12
6	6.85	17.87	6.66	17.53
7	6.16	14.44	6.03	14.22
8	7.33	17.74	7.20	17.47
9	6.14	14.38	6.02	14.23
10	6.37	16.50	6.24	16.29
Average	6.46	16.03	6.32	15.89

The defense remains robust even as model capacity increases; for instance, at $r = 16$, Client 8’s accuracy drops from 29.87% (Correct) to 7.79% (Wrong), a reduction of 22.08 percentage points and approximately $3.83\times$ lower success. While some fluctuations occur—such as Client 4 reaching 11.39% at $r = 4$ —the leakage at higher ranks like $r = 16$ often remains more stable, as seen with Client 10’s 2.44% leakage. This disparity highlights that for some users, higher ranks actually reduce the risk of unintended exposure compared to $r = 4$, where Client 4 sees a $2.31\times$ higher leakage risk. This indicates that while higher ranks can improve authorized utility, the gating module effectively maintains a secure state regardless of the LoRA rank; however, it is important to highlight that effective LoRA privacy tuning necessitates careful rank selection, tailored to each client’s specific requirements.

Gating Module Convergence

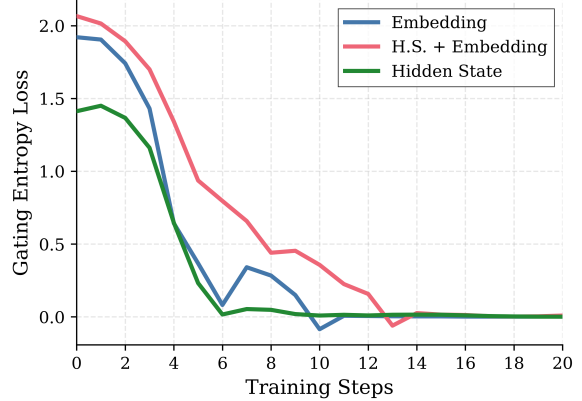


Figure 7: Training loss convergence of the SECUREGATE gating module across different input features. The rapid convergence within 20 steps highlights the module’s lightweight computational requirements.

C.2 Language Modeling Performance (Perplexity)

Table 9 presents the full perplexity (PPL) metrics for all 10 clients using the Llama-1B model on the ECHR dataset. These results compare our SECUREGATE framework against standalone revealing and secure adapter baselines. The data confirms that SECUREGATE maintains linguistic utility similar to the standalone adapters while successfully enforcing token-based privacy boundaries.

C.3 Gating Feature Analysis

We evaluate three candidate feature sources for the gating module: (i) raw token embeddings, (ii) final-layer hidden states (h_{key}), and (iii) their concatenation. Although all configurations converge within 20 training steps, using only the final-layer

hidden states consistently yields the most reliable routing signal. This choice results in lower gating entropy, more accurate adapter selection, and improved training stability, while avoiding the increased dimensionality and overfitting risks introduced by feature concatenation.

Based on these observations, we adopt the final-layer hidden state as the default gating input in the SECUREGATE architecture. Fig. 7 illustrates the corresponding gating entropy loss curves, demonstrating both rapid convergence and a steady reduction in routing uncertainty during training.

D Security Analysis of Keyed Tokenization and Routing

We consider an adversary with full access to *public* model artifacts (base weights, public tokenizer, public logs, and prompts) but without privileged access to organization-specific keys, local tokenizers, or key-management systems. The adversary may (i) inspect or modify public prompts, (ii) attempt to infer or brute-force keys from public artifacts, and (iii) observe adapter routing via black-box queries. **Security Properties.** Under this threat model, SECUREGATE provides the following guarantees:

1. **Key confidentiality.** Keys exist only in local tokenizers and are absent from public vocabularies, preventing exposure or enumeration by external tools.
2. **Atomic key representation.** Each key is a single vocabulary entry and never split into subtokens, preventing partial leakage via subtoken analysis.
3. **Deterministic, non-revealing embeddings.** Key embeddings $e_{\text{key}} \in \mathbb{R}^d$ are generated from a keyed hash and pseudorandom generator, making inversion computationally infeasible.
4. **Low-cost, deterministic routing.** A compact gating MLP maps key embeddings to adapter indices deterministically, supporting reproducible, auditable routing.
5. **Efficient key rotation and revocation.** Key updates only require replacing the keyed embedding and updating the gating module, without retraining the base model or adapters.
6. **Resistance to tokenization-based attacks.** Public tokenizers never contain key tokens, mit-

igating prompt-scraping, token-guessing, and vocabulary-based attacks.

7. **Operational robustness.** Keys, local tokenizers, and gating checkpoints are managed under least-privilege control, with versioning and secure backups to enable safe rotation and roll-back.

Discussion. Items (1)–(2) ensure keys remain atomic and secret within the organization, preventing leakage via public tooling. Item (3) provides reproducible embeddings that are opaque to outsiders. Items (4)–(5) guarantee deterministic, auditable routing while supporting rapid key rotation. Finally, items (6)–(7) reduce the attack surface against token-based, prompt-manipulation, and operational attacks, ensuring secure adapter selection without exposing secrets.

D.1 Deterministic, Non-Revealing Embeddings

Special tokens representing authorization keys are added to the tokenizer vocabulary as $T = \{t_1, \dots, t_N\}$, where each $t_k \in \Sigma^L$ is a unique string of length L over alphanumeric Σ . In SECUREGATE, we use two tokenizers: a *public tokenizer* that never contains key tokens and can be shared across organizations, and a *local private tokenizer* maintained by each organization that stores the protected keys after hashing. The local tokens are first generated as plaintext strings and then hashed before being added into the private tokenizer. The key space entropy is

$$H(K) = L \log_2 |\Sigma|, \quad (10)$$

which results in a key space of $|\Sigma|^L$ and provides exponential growth in entropy.

The base seed is derived from the full set of key tokens:

$$S_{\text{base}} = \left\lfloor \frac{1}{N} \sum_{k=1}^N \text{hash}(t_k) \right\rfloor, \quad (11)$$

where $\text{hash}(t_k)$ extracts a fixed-length value from the MD5 hexadecimal digest of string t_k and converts it into an integer. This S_{base} functions as a collective key derived from all tokens.

For each token $t_k \in T$, we compute

$$S_{t_k} = \text{hash}(t_k). \quad (12)$$

$$S_{\text{comb}} = S_{\text{base}} + S_{t_k}. \quad (13)$$

The key embedding $\mathbf{e}_{\text{key}}(t_k) \in \mathbb{R}^d$ is generated as

$$\mathbf{e}_{\text{key}}(t_k) \sim \mathcal{N}(0, \sigma^2 I_d) \mid S_{\text{comb}}. \quad (14)$$

In our deployment, we use $L = 20$ alphanumeric characters with $|\Sigma| = 62$ (uppercase A–Z, lowercase a–z, and digits 0–9), yielding $H(K) \approx 119$ bits. The brute-force success probability with Q queries is

$$P_{\text{success}} = \min\left(1, \frac{Q}{|\Sigma|^L}\right). \quad (15)$$

For $Q = 10^9$, this yields approximately 2^{-89} , rendering brute-force attacks infeasible under realistic query budgets.

This transformation ensures deterministic key embeddings while maintaining practical non-invertibility. Specifically, recovering authorization strings from \mathbf{e}_{key} is infeasible due to information loss introduced by hashing and the stochastic embedding process.

E Synthetic Data Generation and Augmentations for Gating Training

The routing dataset $\mathcal{D}_{\text{gate}}$ is constructed to systematically cover nominal, malformed, and adversarial input conditions. It includes: (i) valid keyed prompts; (ii) malformed keys; (iii) empty-key positions; and (iv) no-key baselines. For the purpose of these examples, identifiers such as ALPHA, BETA, and GAMMA serve as human-readable placeholders for the high-entropy, secret keys actually utilized by the system. When the gating module encounters an invalid or unauthorized key, the system is designed to deterministically fall back to a designated secure adapter. The following examples characterize the diverse input types within $\mathcal{D}_{\text{gate}}$ used to train the gating module for robust routing and fallback detection:

Illustrative Data Examples: The following samples demonstrate the diverse input types within $\mathcal{D}_{\text{gate}}$ used to harden the routing mechanism:

Valid Key Tokens: Representations of authorized requests using well-formed keys.

- [SPECIAL_TOKEN=ALPHA] Example input for module processing
- [SPECIAL_TOKEN=BETA] A prompt designed to assess token routing
- [SPECIAL_TOKEN=GAMMA] Request requiring special token handling

Algorithm 1: SECUREGATE Training and Personalized Fusion

Input: Raw client datasets $\{\mathcal{D}_n^{\text{raw}}\}_{n=1}^N$;
Frozen base model M ;

Local secure tokenizer \mathcal{T}_{sec}

Output: Personalized adapters
 $\{\Delta\mathbf{w}_{p,n}^{(\text{sec})}, \Delta\mathbf{w}_{p,n}^{(\text{rev})}\}$; Trained
gating MLP G

1 (1) Local Initialization (Sec. 4.1):

1. Generate masked datasets $\mathcal{D}_n^{\text{mask}}$ via NER and scrubbing.
2. Initialize revealing and secure adapters $\Delta\mathbf{w}_n^{(\text{rev})}, \Delta\mathbf{w}_n^{(\text{sec})}$ following (2).

(2) Federated Fine-Tuning (Sec. 4.2):

1. **Inner Loop:** Minimize local objective on $\mathcal{D}_n^{\text{mask}}$ to produce $\Delta\mathbf{w}_n^{t+1}$.
2. **Outer Loop:** Aggregate secure updates $\Delta\bar{\mathbf{w}}^{t+1}$ using (3).
3. Update momentum $\Delta\mathbf{v}^{t+1}$ via (4).
4. Compute global model $\Delta\mathbf{w}^{t+1}$ using (5).

(3) Personalized Fusion (Sec. 4.3):

1. Fuse secure personalized adapter $\Delta\mathbf{w}_{p,n}^{(\text{sec})}$ using (6).
2. Fuse revealing personalized adapter $\Delta\mathbf{w}_{p,n}^{(\text{rev})}$ using (8).
3. Optimize fusion coefficients (α, β) via (7).

(4) Gating Module Training (Sec. 4.4):

1. Train G on synthetic key patterns using the cross-entropy loss $\mathcal{L}_{\text{gating}} = -\sum_i y_i \log p_i$.
-

Malformed/Incorrect Keys: Samples designed to test the router’s resistance to corruption.

- [SPECIAL_TOKEN=ALPH] Example input for module processing
- [SPECIAL_TOKEN=BETA_CORRUPT] Scenario of ambiguous adapter routing

Empty/No Key Baselines: Cases that must trigger the default fallback logic.

- [SPECIAL_TOKEN=] Empty string placeholder
- [SPECIAL_TOKEN=] An instruction that triggers fallback adapter
- A prompt designed to assess token routing

By training across these distinct key states, the gating network learns to enforce a deterministic mapping for valid authorization tokens while identifying corrupted or missing keys as candidates for the secure fallback adapter.

Algorithm 2: SECUREGATE Two-Pass Inference Gating

Input: User request X_{raw} ; Personalized adapters $\{\Delta\mathbf{w}_{p,n}^{(sec)}, \Delta\mathbf{w}_{p,n}^{(rev)}\}$;

Gating MLP G ; Confidence threshold τ

Output: Authorized or sanitized model response Y

1 Pass 1: Gating & Authorization

1. Forward request through M and extract hidden state h_{key} via (9).
2. Compute selection probabilities p_i using the softmax $p_i = \frac{\exp(z_i)}{\sum_{j \in \mathcal{I}} \exp(z_j)}$.
3. Select adapter index $a^* = \arg \max p_i$.
4. **Enforce Policy:** If $p_{a^*} \leq \tau$, default to secure adapter $\Delta\mathbf{w}_{p,n}^{(sec)}$.

Pass 2: Clean Generation

1. Strip key token from prompt to prevent output contamination.
2. Generate final response Y using selected adapter $\Delta\mathbf{w}_{p,n}^{(a^*)}$.

return Y

F Algorithms

Algorithm 1 provides the complete training workflow for the SECUREGATE framework, which is executed in four primary stages: (1) *local initialization* involving data masking through NER and scrubbing, (2) *federated fine-tuning* utilizing inner-loop local updates and outer-loop momentum-based aggregation, (3) *the personalized fusion* of secure and revealing adapters, and (4) *the final optimization* of the gating module.

Algorithm 2 details the two-pass inference process. The first pass focuses on gating and authorization by extracting hidden states and enforcing security policies through a confidence threshold τ , while the second pass performs clean generation by stripping key tokens to prevent output contamination.