SCIENTIFIC Publishing Limited

**Journal of Artificial Intelligence and Science Communication**

https://ojs.ukscip.com/index.php/jaisc

*Article*

# CEDAR: A Federated Meta-Learning Framework for Personalized IoT via Layer-Wise Adaptive Model Uploading

**Bisma Gulzar** * **and Shabir Ahmad Sofi**

Department of Information Technology, National Institute of Technology, Srinagar 190006, India

* Correspondence: bisma_2021phaite001@nitsri.ac.in

**Received:** 26 February 2025; **Revised:** 6 June 2025; **Accepted:** 9 July 2025; **Published:** 15 July 2025

**Abstract:** The Personalized Internet of Things (PIoT) demands intelligent learning models that can adapt to highly heterogeneous user data while preserving privacy, scalability, and security. Centralized learning approaches are impractical in PIoT settings due to privacy regulations, non-independent and identically distributed (non-IID) data distributions, and vulnerability to adversarial attacks. To address these challenges, this paper proposes Consent-Driven Ethical Data Access and Regulation (CEDAR), a federated meta-learning framework that enables secure and adaptive personalization across the cloud–edge–device continuum. CEDAR integrates meta-learning with federated learning to extract transferable representations from distributed data, allowing rapid adaptation to individual user contexts with minimal local updates. A layer-wise adaptive uploading mechanism selectively communicates model updates based on parameter importance, substantially reducing communication overhead and accelerating convergence. In addition, asymmetric uploading and anomaly-aware aggregation enhance robustness against gradient inversion and model poisoning attacks. Extensive evaluations on six benchmark datasets covering regression, text classification, and image recognition tasks demonstrate that CEDAR achieves up to 60.39% higher accuracy compared to FedAvg-based federated learning, while reducing communication cost by 23.36% and improving adversarial robustness relative to other state-of-the-art baselines. Ablation studies further confirm the complementary contributions of CEDAR's core components. By jointly optimizing personalization, privacy, efficiency, and security, CEDAR provides a scalable and ethically aligned learning framework for next-generation PIoT applications in domains such as smart mobility, healthcare, and the digital economy.

**Keywords:** Federated Learning; Meta-Learning; Personalized Internet of Things (PIoT); Edge Intelligence; Privacy Preservation; Adversarial Robustness

## 1. Introduction

The Personalized Internet of Things (PIoT) has emerged as a transformative paradigm in next-generation intelligent infrastructures, where heterogeneous devices, sensors, and services are interconnected to provide user-specific functionality [1]. Unlike conventional IoT systems that largely focus on optimizing collective performance, PIoT emphasizes tailoring services to individual needs by exploiting multimodal data streams that reflect behavioral, contextual, and environmental patterns. Applications such as proactive healthcare monitoring, adaptive driver safety systems, and personalized resource management in smart homes highlight the immense potential of PIoT to enhance both quality of life and operational efficiency [2]. Despite this promise, the deployment of PIoT faces critical challenges stemming from the sensitivity, distribution, and heterogeneity of data, which traditional machine learning paradigms are unable to effectively resolve.

Centralized learning approaches, which rely on aggregating raw data at a cloud server, are particularly ill-suited for PIoT. Privacy regulations such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA) impose strict restrictions on how personal data can be transferred and processed, making large-scale data aggregation both impractical and legally noncompliant. Furthermore, PIoT data are inherently non-independent and identically distributed, as each device captures a highly localized slice of user behavior, leading to significant variations across clients. Models trained on aggregated data tend to generalize poorly in such scenarios, offering suboptimal personalization [3]. Additional limitations arise from the communication bottlenecks and energy constraints inherent to the cloud–edge–device continuum. Large-scale data transfer introduces latency, saturates bandwidth, and accelerates battery depletion in resource-constrained devices, thereby limiting scalability. Centralized training also creates single points of failure and attractive targets for adversarial attacks, including gradient inversion and data poisoning, which compromise both system reliability and user trust [4].

Federated Learning (FL) has been introduced as a privacy-preserving alternative by allowing devices to collaboratively train a shared model without exchanging raw data. Each device performs local training and shares only model updates, which are aggregated at a server to produce a global model [5]. This approach mitigates privacy risks and reduces the need for direct data transfer. However, when applied to PIoT, vanilla federated learning still exhibits significant shortcomings. Non-IID data severely undermines the convergence and generalizability of global models, as dominant patterns from majority clients overshadow the contribution of minority distributions. Communication inefficiencies remain a critical bottleneck, since the iterative exchange of model parameters across thousands of devices generates high costs in bandwidth-constrained environments [6]. Moreover, federated learning remains vulnerable to security risks, including gradient inversion attacks that reconstruct private information from model updates and poisoning attacks that inject malicious knowledge into the global model.

Meta-learning, commonly described as "learning to learn," provides a promising solution to these limitations by enabling the extraction of transferable knowledge that can be quickly adapted to new tasks or individual users with minimal data. In the context of PIoT, meta-learning allows global models to retain generalizable structural information while remaining adaptable to user-specific contexts [7]. Rather than producing a static global model, meta-learning yields a meta-model that can be fine-tuned locally with a few training iterations, thereby providing efficient personalization without requiring large-scale retraining. When combined with federated learning, this approach—federated meta-learning—offers the potential to balance privacy, personalization, and scalability, making it an ideal fit for PIoT environments [8].

Realizing federated meta-learning for PIoT, however, introduces additional challenges. Heterogeneous and sparse data distributions complicate the balance between global generalization and local adaptation, often leading to biased or unstable training. Communication efficiency becomes even more critical, as meta-learning introduces iterative fine-tuning that amplifies overhead in bandwidth-limited environments [9]. Security concerns are heightened by the open and distributed nature of PIoT, where adversaries can exploit model updates to either infer sensitive data or disrupt convergence. Finally, the variability of device resources, ranging from high-capacity gateways to low-power wearables, demands flexible mechanisms that can adapt training and communication strategies to heterogeneous computational and network conditions [10].

To address these challenges, this work introduces Cedar, a federated meta-learning framework specifically designed for PIoT systems. Cedar incorporates four core innovations. First, it integrates meta-learning into federated training to extract and transfer structural knowledge from heterogeneous private datasets, thereby producing models that are both generalizable and rapidly adaptable. Second, it employs a layer-wise adaptive uploading strategy that selectively transmits critical model layers, reducing communication costs and accelerating convergence without sacrificing performance. Third, it incorporates asymmetric model uploading and anomaly-aware aggregation mechanisms that safeguard against gradient inversion and poisoning attacks, ensuring robustness in adversarial settings [11]. Finally, Cedar supports personalized deployment by enabling each device to fine-tune the global meta-model locally, ensuring optimal performance tailored to individual user contexts.

We evaluate Cedar extensively across six publicly available datasets spanning regression, text classification, and image classification tasks, using models of varying complexity such as One-Dimensional Convolutional Neural Network (CNN-1D), Multi-layer Bidirectional Long Short-Term Memory (MBiLSTM), and 18-layer Residual Network (ResNet-18). Experimental results demonstrate that Cedar consistently outperforms state-of-the-art baselines, achieving up to 60.39% higher accuracy in heterogeneous scenarios, reducing communication costs by up to

23.36%, and significantly improving resilience against gradient inversion and label-flipping attacks. Furthermore, Cedar enables rapid personalization, with local models achieving strong performance after only a few rounds of fine-tuning. Ablation studies confirm the complementary roles of its core modules, while cross-domain evaluations highlight its applicability across diverse PIoT environments [12].

The contributions of this work are fourfold. First, we present Cedar, a novel federated meta-learning framework that unifies personalization, efficiency, and security in PIoT. Second, we introduce adaptive communication and security mechanisms that reduce overhead while defending against adversarial exploitation. Third, we provide a comprehensive empirical evaluation across six datasets and multiple domains, demonstrating the superiority of Cedar over existing methods [13]. Finally, we offer practical insights into deploying federated meta-learning in real-world PIoT systems, addressing the simultaneous demands of privacy preservation, heterogeneity management, and efficient personalization.

By combining federated learning and meta-learning with adaptive communication and security strategies, Cedar advances the state of the art in personalized IoT. It establishes a practical and scalable pathway for deploying intelligent, privacy-aware, and user-centric models across diverse domains, thereby unlocking the transformative potential of PIoT in healthcare, mobility, economy, and beyond [14].

## 2. Literature Review

Research on intelligent Internet of Things (IoT) systems has accelerated in the past decade, with a growing shift toward personalization to meet user-specific requirements. Early IoT systems primarily focused on connectivity and automation, emphasizing efficient data collection and centralized cloud analytics. While these approaches enabled scalable services, they offered little capacity for fine-grained personalization, as user data was pooled into generalized models that rarely reflected individual contexts. Recent work has recognized this limitation and has motivated the development of techniques for personalized IoT (PIoT), which tailors intelligent services to individual users by exploiting distributed, multimodal, and heterogeneous data [15]. This literature review outlines research across four interconnected areas—personalized IoT systems, federated learning, meta-learning, and federated meta-learning—before identifying gaps that motivate the design of our proposed framework, Cedar. **Table 1** specifies and provides a comparative Analysis of Learning Paradigms in PIoT.

**Table 1.** Comparative Analysis of Learning Paradigms in PIoT.

| Approach | Short Summary | Privacy Preservation | Adaptability to Non-IID Data | Communication Efficiency | Security Robustness | Personalization |
|---|---|---|---|---|---|---|
| Federated Learning (FL) | Decentralized training with aggregation of local model updates (e.g., FedAvg, FedProx). Effective for large-scale IoT but limited personalization. | Yes | No | No | No | No |
| Meta-Learning (ML) | Learns to adapt quickly to new tasks using few samples (e.g., MAML, Reptile). Provides rapid personalization but usually centralized, raising privacy concerns. | No | Yes | No | No | Yes |
| Federated Meta-Learning (FML) | Combines FL and ML to enable decentralized personalization (e.g., FedMeta-MAML, FedReptile). Handles heterogeneity but suffers from high communication cost and security issues. | Yes | Yes | No | No | Yes |
| **Proposed Cedar** | Domain-adaptive framework integrating federated meta-learning with layer-wise adaptive uploading, anomaly detection, and asymmetric updates. Achieves balance of privacy, efficiency, security, and personalization. | Yes | Yes | Yes | Yes | Yes |

## 2.1. Personalized IoT Systems

The concept of personalization in IoT emerged from the recognition that data collected by sensors, wearables, and smart devices is highly context-specific and varies significantly across individuals. Early studies in healthcare demonstrated that a one-size-fits-all model for health monitoring systems was ineffective, as physiological signals and lifestyle patterns differ widely across users [16]. Similarly, in intelligent transportation systems, models trained on global driving datasets struggled to adapt to individual driver behaviors and environmental conditions. To address this, several approaches have been proposed, including local model training on devices and hybrid cloud–edge frameworks. These approaches improved personalization to an extent but remained constrained by issues of data fragmentation, privacy, and communication cost. Furthermore, local-only training often suffered from insufficient data volume, leading to overfitting and poor generalization [17]. This highlighted the necessity of distributed collaborative learning frameworks capable of balancing personalization with knowledge transfer.

## 2.2. Federated Learning in IoT

Federated learning (FL) has become a prominent paradigm for privacy-preserving distributed model training in IoT environments. McMahan et al.'s [10] seminal work on FedAvg established the foundational idea of aggregating local model updates instead of raw data, enabling decentralized training across thousands of devices. Numerous extensions of FedAvg have since been proposed to tackle challenges specific to IoT, including non-IID data, straggler devices, and communication bottlenecks [18]. For example, approaches such as FedProx introduced proximal terms to stabilize training under heterogeneous conditions, while adaptive optimization methods (e.g., FedAdam, FedYogi) sought to improve convergence speed.

Within IoT, federated learning has been applied to diverse applications, including activity recognition, medical diagnostics, and predictive maintenance. These studies confirm that FL can enhance privacy preservation and reduce data transfer costs, yet they also reveal fundamental weaknesses. Chief among these is the inability of global federated models to adapt effectively to individual clients when data distributions diverge. In PIoT settings, where personalization is paramount, standard FL often yields models that perform well on average but poorly on specific users [19]. Additionally, iterative communication of full model parameters between clients and the server remains prohibitively expensive in bandwidth-limited environments such as mobile or edge networks. Security concerns also persist, as federated updates can be exploited through gradient inversion, membership inference, or poisoning attacks, undermining the trustworthiness of deployed systems.

## 2.3. Inference Attacks in Federated and PIoT Systems

Recent studies have demonstrated that privacy risks in federated learning extend beyond data exposure, as shared gradients and model updates can unintentionally reveal sensitive information. One of the most critical threats is gradient inversion, where adversaries reconstruct private inputs from shared gradients, even without access to raw data. It showed that gradients exchanged during training can leak both input features and labels, raising serious concerns for federated deployments involving sensitive user data. Subsequent extensions revealed that label information can often be inferred directly from gradient directions, particularly when cross-entropy loss and small batch sizes are used—conditions commonly observed in PIoT scenarios.

Membership inference attacks represent another high-risk vector, enabling adversaries to determine whether specific user records were included in local training. Such attacks are effective in federated settings where adversaries gain access to model updates or confidence outputs over multiple rounds. In PIoT environments, where user behavior, health signals, and contextual patterns are highly distinctive, successful membership inference can directly compromise user privacy and trust.

Label leakage attacks further exacerbate these risks by exploiting correlations between gradients and class labels, allowing attackers to infer sensitive task-related information even when secure aggregation is employed. These attacks are particularly effective against personalized and frequently updated models, making them especially relevant to PIoT systems that rely on continuous adaptation. Collectively, these inference attacks reveal that federated learning alone does not guarantee privacy and that update-level information leakage must be explicitly addressed in secure PIoT frameworks.

## 2.4. Meta-Learning for Personalization

In parallel with the rise of FL, meta-learning has gained attention as a powerful approach for personalization. Unlike conventional learning, which optimizes a model for a single task, meta-learning aims to optimize the learning process itself by extracting knowledge from a distribution of tasks. Popular methods such as Model-Agnostic Meta-Learning (MAML), Reptile, and meta-SGD have shown that models can be trained to quickly adapt to new tasks with only a few gradient updates [20]. This property is particularly valuable in environments with scarce or heterogeneous data, making meta-learning a natural fit for PIoT systems.

Applications of meta-learning in IoT-related domains have demonstrated promising results. In healthcare, meta-learning frameworks have been used to personalize models for patient-specific diagnoses with minimal data. In speech recognition, meta-learning enabled rapid adaptation to new speakers. In image classification tasks involving non-uniform datasets, meta-learning facilitated transferability across diverse classes. However, most of these studies rely on centralized training of meta-models, which reintroduces privacy and scalability issues when applied to real-world PIoT environments [21]. Without decentralized mechanisms, meta-learning alone cannot meet the privacy and efficiency requirements of PIoT.

## 2.5. Federated Meta-Learning

The natural convergence of FL and meta-learning has given rise to federated meta-learning, which combines the privacy-preserving benefits of federated training with the adaptability of meta-learning. Early works such as FedMeta-MAML and FedMeta-SGD adapted MAML-style optimization into federated environments, allowing global meta-models to be fine-tuned efficiently on local client data. FedReptile extended the Reptile algorithm to distributed settings, emphasizing communication efficiency [22]. These studies established the feasibility of federated meta-learning but also highlighted several persistent challenges: the high communication cost of iterative meta-training, poor stability under highly non-IID distributions, and vulnerability to malicious attacks.

Recent extensions attempted to address these challenges through communication reduction techniques, such as gradient sparsification and quantization, or by employing client clustering to group similar distributions. While these methods provided partial improvements, they often sacrificed either accuracy or robustness. Furthermore, most existing federated meta-learning approaches were evaluated on benchmark datasets with limited heterogeneity, failing to capture the extreme variability encountered in PIoT systems [23]. In practice, PIoT involves devices with widely different computational capabilities, irregular participation patterns, and adversarial exposure, all of which stress-test the scalability of these methods.

## 2.6. Security in Federated and Meta-Learning

Security has emerged as a central concern in federated and federated meta-learning systems. In addition to traditional poisoning and Byzantine attacks, inference-based threats such as gradient inversion, label leakage, and membership inference have demonstrated the ability to extract sensitive information directly from shared model updates. Although differential privacy can mitigate these attacks by injecting noise, it often introduces a significant trade-off between privacy guarantees and model accuracy, which is problematic for personalization-centric PIoT applications.

Cryptographic techniques such as secure multiparty computation and homomorphic encryption protect individual updates during aggregation but incur substantial computational and communication overheads, limiting their applicability to resource-constrained IoT devices. Similarly, gradient clipping and randomization reduce information leakage but can impair convergence speed and adaptation quality. Anomaly-aware aggregation techniques improve robustness against poisoning but do not inherently address reconstruction-based inference attacks. Consequently, existing defenses either impose heavy overheads or degrade personalization performance, highlighting the need for lightweight and PIoT-aware security mechanisms [24].

## 2.7. Identified Research Gaps

From the literature, several clear gaps emerge. First, while FL enables privacy-preserving collaboration and meta-learning provides adaptability, existing federated meta-learning frameworks have not been fully optimized for PIoT environments where data are highly non-IID, sparse, and sensitive. Second, communication efficiency

has received limited attention in federated meta-learning, despite being a critical bottleneck in large-scale PIoT deployments [25]. Third, current solutions lack robust yet lightweight mechanisms to defend against adversarial threats, leaving systems vulnerable to privacy leakage and poisoning. Finally, there has been little emphasis on practical deployment strategies that reconcile the heterogeneity of PIoT devices with the computational demands of federated meta-learning.

### 2.8. Positioning of This Work

This study addresses the identified gaps by proposing Cedar, a federated meta-learning framework explicitly designed for PIoT. Unlike existing approaches that rely primarily on noise injection or heavy cryptographic defenses, Cedar adopts a layer-wise adaptive uploading strategy that selectively shares model parameters based on their transferability and sensitivity, inherently reducing exposure to inference attacks such as label leakage and gradient inversion. Furthermore, asymmetric model uploading combined with anomaly-aware aggregation strengthens resilience against poisoning and malicious client behavior without imposing excessive computational overhead. By jointly addressing personalization, communication efficiency, and inference-level privacy risks, Cedar provides a secure, scalable, and practically deployable learning framework for next-generation PIoT systems.

## 3. Proposed Framework

Cedar is conceived as a federated meta-learning framework that systematically addresses the intertwined challenges of heterogeneity, communication inefficiency, adversarial vulnerability, and personalization in the Personalized Internet of Things (PIoT) [26]. The framework is structured as a three-phase pipeline—learning preparation, meta-model training, and personalized deployment—augmented with four orthogonal mechanisms that collectively ensure security, adaptability, and efficiency.

### 3.1. Three-Phase Workflow

The operation of Cedar begins with the learning preparation phase, triggered whenever the coordinator (an authorized cloud or edge server) receives a model training or update request from a PIoT application. The coordinator first decomposes this request into a well-defined learning specification that encompasses data modality requirements, neural architecture configuration, and hyperparameter initialization [27]. For instance, in a healthcare scenario, the specification may demand convolutional models for ECG data, while in mobility analytics, it may select recurrent architectures for sequential driving traces.

After defining the specification, the coordinator constructs a learning consortium by dynamically selecting eligible clients. Selection is governed by a multi-criteria eligibility model that jointly considers (i) data sufficiency and sample diversity, (ii) historical model update quality, (iii) available computational resources, (iv) network bandwidth and latency estimates, and (v) device energy constraints. This holistic assessment allows Cedar to prioritize devices that can contribute meaningful updates within acceptable time bounds.

To prevent slow or intermittently connected devices from bottlenecking training, Cedar incorporates explicit straggler mitigation mechanisms. Devices predicted to exhibit high communication latency or prolonged computation time are either temporarily excluded from the current round or assigned asynchronous participation with relaxed deadlines. In addition, the coordinator enforces round-level timeouts and quorum-based aggregation, proceeding once a sufficient fraction of updates has been received rather than waiting for all selected clients. Devices identified as stragglers may rejoin in subsequent rounds when their resource conditions improve. This adaptive client orchestration mitigates skewed participation, prevents training stalls, and ensures that the federated cohort remains representative yet computationally sustainable.

The meta-model training phase constitutes the core of Cedar. Once clients are activated, the coordinator dispatches the training specification. Each participating client performs local optimization on its private dataset, yielding an intermediate update $\Delta\theta_i^t$. Unlike conventional federated learning, where the entire parameter vector is transmitted, Cedar adopts a layer-wise adaptive uploading mechanism that selectively communicates only the most informative layers.

Formally, the local update is decomposed into layer-level components $\Delta\theta_{i,\ell}^t$, where $\ell \in \mathcal{L}$. For each layer $\ell$, the

client computes an importance score that jointly captures first-order sensitivity and curvature information:

$$S_\ell = \lambda \left\|\nabla_\ell \mathcal{L}_i(\theta)\right\|_2 \; + \; (1-\lambda)\, \mathcal{F}_\ell,$$

where $\|\nabla_\ell \mathcal{L}_i(\theta)\|_2$ denotes the gradient magnitude of the local loss with respect to layer $\ell$, $\mathcal{F}_\ell$ represents a diagonal approximation of the Fisher information for that layer, and $\lambda \in [0, 1]$ controls the trade-off between gradient-based importance and curvature-based sensitivity. A layer is selected for transmission if $S_\ell \geq \tau$, where $\tau$ is a predefined significance threshold; updates for layers that do not meet this criterion are pruned and retained locally.

This selective transmission strategy serves three objectives: (i) reducing communication bandwidth by eliminating low-utility updates, (ii) accelerating convergence by prioritizing layers with the greatest optimization impact, and (iii) mitigating privacy leakage by limiting unnecessary parameter exposure.

The coordinator receives these partial updates and performs secure meta-aggregation. Cedar integrates meta-learning principles during aggregation: instead of naïvely averaging gradients, it aligns updates to optimize for cross-task generalizability. The resulting global meta-model thus captures structural knowledge—latent representations and task-invariant patterns—while retaining the capacity for rapid downstream adaptation. This iterative exchange between clients and the coordinator continues until the meta-model satisfies a convergence criterion (e.g., validation loss plateau) or the system reaches a predefined upper bound on communication rounds. By jointly optimizing aggregation and update selection, Cedar substantially reduces training latency in heterogeneous federated environments.

In the personalized deployment phase, the trained global meta-model is distributed to participating devices. Each client initializes from this meta-model and performs a lightweight fine-tuning step on its private data. Owing to meta-learning's inductive bias, only a few gradient steps are required for the model to specialize to the client's unique distribution. This personalization stage yields models that are globally consistent yet locally optimized, aligning well with PIoT's vision of hyper-personalization across domains such as smart health, adaptive mobility, and digital economy services.

**Four Core Functions**

Beyond this three-phase workflow, Cedar implements four cross-cutting mechanisms that reinforce adaptability, cost-efficiency, and robustness [28].

1. Meta-learning for domain adaptability
   A critical challenge in PIoT is the high variance of data distributions across devices. Conventional federated learning aggregates localized models without explicitly capturing transferable invariants, leading to biased global models that underperform on minority clients. Cedar addresses this by embedding a meta-learning layer within the aggregation pipeline. Each local update contributes not only raw parameter deltas but also higher-order information that informs the global model's sensitivity to task variation. As a result, Cedar optimizes for model initialization states that are close to many local optima, enabling rapid adaptation during deployment.
2. Layer-wise adaptive model uploading
   Communication inefficiency is a well-documented bottleneck in federated learning. Traditional schemes transmit the full parameter space, which may comprise millions of weights across many devices and training rounds. In PIoT environments with limited bandwidth and energy constraints, such redundancy is prohibitive. Cedar mitigates this via a layer prioritization algorithm in which clients compute a contribution score for each layer using gradient norms, information gain, and variance reduction criteria. Only the top-ranked layers are transmitted, reducing communication payloads by 20–40% without sacrificing accuracy, accelerating convergence, and lowering vulnerability to inference attacks.
3. Security and anomaly-aware aggregation
   PIoT systems are particularly susceptible to adversarial threats such as poisoned updates and gradient-based inference attacks. Cedar employs a dual-defense strategy. First, asymmetric model uploading retains sensitive or label-correlated layers locally or transforms them prior to transmission, reducing reconstruction risk. Second, anomaly-aware aggregation uses robust statistics (e.g., trimmed mean and cosine similarity) to identify and suppress anomalous updates. Together, these mechanisms preserve global model integrity even under adversarial pressure.

4. Personalized fine-tuning for heterogeneous devices
   Cedar supports a wide range of device capabilities by enabling efficient personalization. Clients fine-tune the received meta-model using minimal local computation—often fewer than ten iterations—yielding models tailored to individual contexts while respecting resource constraints. This ensures scalable personalization without fragmenting the system into isolated device-specific models [29].

### 3.2. Technical Deep Dive and Implications

The integration of these mechanisms yields several technical benefits. Cedar improves statistical efficiency by explicitly optimizing for cross-task generalization, enhances communication efficiency through selective uploads and quorum-based aggregation, and embeds lightweight security primitives directly into the learning workflow. Asymmetric uploads reduce inference leakage, while anomaly detection limits adversarial impact. Finally, Cedar enables scalable personalization by balancing global consistency with local optimization, meeting the core requirements of large-scale PIoT deployments [30].

From a systems perspective, Cedar shifts from monolithic global models toward federated meta-models that function as adaptive priors across the cloud–edge–device continuum. This design supports regulatory compliance, efficient bandwidth utilization, and resilience to adversarial activity, making Cedar well-suited for next-generation PIoT systems. **Algorithm 1** specifies Federated Meta-Learning with Adaptive and Asymmetric Uploading.

**Algorithm 1** CEDAR: Federated Meta-Learning with Adaptive and Asymmetric Uploading

**Require:** Initial meta-model $\theta_0$, client set $\mathcal{C}$, rounds $T$, threshold $\tau$, asymmetry bound $\alpha$, quorum $q$
**Ensure:** Personalized models $\{\theta_i^{\text{pers}}\}$

**for** $t = 1$ to $T$ **do**
  Coordinator selects eligible clients $\mathcal{C}_t \subseteq \mathcal{C}$
  Broadcast meta-model $\theta_{t-1}$
  **for** each client $i \in \mathcal{C}_t$ **do**
    Train locally using private data
    Compute update $\Delta\theta_i^t$
    Compute layer importance scores $S_\ell$
    Select layers $\mathcal{L}_i = \{\ell : S_\ell \geq \tau\}$
    Apply layer-wise adaptive and asymmetric uploading
    Send $\tilde{\Delta\theta}_i^t$ to coordinator
  **end for**
  Coordinator performs anomaly-aware aggregation
  Update global model $\theta_t$
**end for**
**for** each client $i$ **do**
  Fine-tune $\theta_T$ locally to obtain $\theta_i^{\text{pers}}$
**end for**
**return** $\{\theta_i^{\text{pers}}\}$

## 4. Experimental Setup

This section details the datasets, federated configuration, hyperparameters, and evaluation protocol used to rigorously assess the performance, robustness, and efficiency of CEDAR. All experimental choices are made to ensure reproducibility, fairness, and realism in Personalized IoT (PIoT) environments characterized by device heterogeneity, data imbalance, and communication constraints.

### 4.1. Datasets and Task Characterization

Experiments are conducted on six benchmark datasets spanning three representative PIoT task categories:

- **Bike Sharing Dataset (BSD):** A structured tabular regression dataset used to predict continuous demand signals under temporal and behavioral variability.
- **Stanford Sentiment Treebank-5 (SST-5):** A fine-grained sentiment classification dataset representing language-

based PIoT services with high semantic heterogeneity.

- **Facial Expression Recognition (FER):** A facial expression recognition dataset used to evaluate robustness under visual non-IID distributions.
- **Fashion-MNIST (FMNIST):** A standard vision benchmark consisting of apparel images, widely adopted for federated learning evaluation.
- **International Skin Imaging Collaboration (ISIC):** A privacy-sensitive medical imaging dataset for skin lesion classification, representing healthcare IoT settings.
- **State Farm Distracted-Driven Detection Dataset (SFDDD):** A distracted driver detection dataset in which each client corresponds to a distinct driver, exhibiting strong behavioral and distributional heterogeneity.

To emulate realistic PIoT conditions, data are partitioned across clients using a Dirichlet distribution with concentration parameter $\beta = 0.3$, inducing statistically heterogeneous and label-skewed local datasets.

## 4.2. Federated Meta-Learning Configuration

All federated experiments are conducted under a consistent system configuration:

- Total number of clients: $N = 50$
- Client participation ratio per round: $|\mathcal{C}_t|/N = 0.2$
- Communication rounds: $T = 100$
- Quorum ratio for aggregation: $q = 0.7$
- Client selection: multi-criteria eligibility filtering with straggler prediction

Clients predicted to exceed round-level latency budgets are either assigned asynchronous participation or deferred to subsequent rounds. Aggregation proceeds once the quorum condition is met, preventing straggler-induced stalls.

## 4.3. Optimization and Hyperparameter Settings

To ensure fairness, all baselines and CEDAR variants use identical base training settings unless otherwise specified.

**Local Training:**

- Optimizer: Adam
- Batch size: 32
- Local learning rate: $1 \times 10^{-3}$
- Local epochs per round: $E = 5$

**Meta-Learning Parameters:**

- Meta learning rate: $5 \times 10^{-4}$
- Inner-loop adaptation steps: 1–5

**CEDAR-Specific Parameters:**

- Layer importance threshold $\tau$: top 30% ranked layers
- Gradient–Fisher weighting $\lambda$: 0.5
- Asymmetric perturbation bound $\alpha$: 0.3
- Layer masking ratio: dynamically determined per client

All hyperparameters are selected through preliminary validation and held constant across tasks to avoid dataset-specific bias.

## 4.4. Evaluation Metrics

Performance is evaluated using task-appropriate metrics:

- Classification accuracy for text and image tasks;
- Mean Squared Error (MSE) for regression tasks;

- Communication cost measured as total transmitted parameters per round;
- Robustness under adversarial settings, measured by accuracy degradation.

For adversarial robustness experiments, 40% of randomly selected clients perform label-flipping attacks.

## 4.5. Reproducibility and Implementation Details

All experiments are implemented in PyTorch and executed on a workstation equipped with an NVIDIA RTX-series GPU. Each experimental configuration is repeated five times with different random seeds, and the average results are reported.

Diagonal Fisher information estimates are computed periodically to limit computational overhead on client devices. This design ensures that communication savings from layer-wise adaptive uploading outweigh the additional local computation cost, particularly in bandwidth- and energy-constrained PIoT environments.

## 4.6. Comparison with Recent State-of-the-Art and Baseline Justification

This study compares CEDAR against widely adopted federated learning and federated meta-learning baselines, including FedAvg, FedMeta-MAML, FedMeta-SGD, FedFOMAML, and FedReptile. These methods are selected because they remain the most established and reproducible benchmarks for evaluating federated personalization and meta-learning performance under heterogeneous and non-IID settings.

Recent post-2023 advances in federated learning and personalized IoT have also been carefully reviewed, including works published in 2024 and 2025 (e.g., You et al., Nature Communications, 2025 [14]). While these methods introduce valuable innovations—such as architectural personalization, client clustering, or hierarchical coordination—they are not always directly comparable to CEDAR due to one or more of the following reasons. First, several recent approaches rely on problem formulations that are orthogonal to the objectives of CEDAR, focusing exclusively on model compression, clustering, or system-level scheduling rather than the joint optimization of personalization, communication efficiency, and security. Second, many recently proposed methods do not provide publicly available implementations or standardized evaluation protocols, making fair and reproducible comparison infeasible. Third, some works assume task- or domain-specific priors (e.g., fixed architectures or modality-specific heuristics) that are incompatible with the cross-domain PIoT setting considered in this study.

Nevertheless, to ensure that the manuscript reflects the current research landscape, recent state-of-the-art methods are now explicitly discussed in the literature review, and their conceptual differences from CEDAR are clearly articulated. Importantly, CEDAR is positioned as a complementary advancement that integrates adaptive communication, federated meta-learning, and security-aware aggregation into a single unified framework—an integration that is not jointly addressed by existing methods to date.

We emphasize that all selected baselines are strong, well-recognized methods that provide a rigorous and fair foundation for evaluating the empirical gains of CEDAR. Future work will prioritize extending empirical comparisons as additional reproducible implementations of recent methods become available.

## 4.7. Baseline Methods and Implementation Details

To ensure a fair and reproducible comparison, all baseline methods are implemented under the same experimental conditions as CEDAR. This includes identical dataset partitions, client participation ratios, communication rounds, and computational budgets. Wherever available, we rely on official or widely cited reference implementations, adapting them only when necessary to match the PIoT setting.

**FedAvg:** FedAvg follows the standard federated averaging protocol. Each client performs local training for $E = 5$ epochs using the Adam optimizer with a learning rate of $1 \times 10^{-3}$ and batch size 32. The server aggregates full model updates using simple weighted averaging.

**FedFOMAML:** FedFOMAML implements first-order model-agnostic meta-learning in a federated setting. Clients execute a single inner-loop gradient update without computing second-order derivatives. The meta learning rate is set to $5 \times 10^{-4}$, and all other hyperparameters match those of FedAvg for consistency.

**FedMeta-MAML:** FedMeta-MAML adopts the full MAML optimization framework within federated learning. Each client performs inner-loop adaptation using its local data, followed by meta-gradient computation at the server. To control computational overhead, inner-loop steps are limited to 1–3 iterations, and all clients use identi-

cal learning rates and batch sizes.

**FedReptile:** FedReptile is implemented as a communication-efficient alternative to MAML. Clients perform multiple local SGD steps, and the server updates the global model by averaging the differences between initial and final client parameters. Optimization settings mirror those of FedMeta-MAML to maintain fairness.

**FedMeta-SGD:** FedMeta-SGD extends federated meta-learning by learning per-parameter step sizes. In our implementation, the learned learning rates are initialized uniformly and updated jointly with model parameters. All training settings are aligned with those used for other baselines.

For all baselines, no communication compression, security defenses, or adaptive uploading mechanisms are applied unless explicitly defined by the method. This ensures that improvements observed in CEDAR arise from its proposed design components—namely, layer-wise adaptive uploading, asymmetric uploading, and anomaly-aware aggregation—rather than from differences in training budgets or system assumptions.

## 5. Rigorous Analysis of Asymmetric Uploading

Asymmetric uploading is a core security mechanism in Cedar, designed to reduce information leakage from shared updates while preserving convergence and personalization performance. Unlike symmetric federated optimization, where full model updates are uniformly transmitted by all clients, asymmetric uploading intentionally applies structured transformations to selected layers prior to communication. Given its central role in the security guarantees of Cedar, this section presents a rigorous theoretical analysis of asymmetric uploading and a quantitative characterization of its security–utility trade-off.

### 5.1. Formal Problem Definition

Let $w \in \mathbb{R}^d$ denote the global meta-model parameters. In federated round $t$, a client $i \in \mathcal{C}_t$ performs local meta-learning–guided optimization and computes an update $\Delta w_i^t$. Standard federated aggregation transmits $\Delta w_i^t$ directly to the coordinator.

In Cedar, clients apply an asymmetric uploading operator $\mathcal{A} : \mathbb{R}^d \to \mathbb{R}^d$, yielding the transmitted update

$$\tilde{\Delta w_i}^t = \mathcal{A}(\Delta w_i^t).$$

The operator $\mathcal{A}$ selectively acts on a subset of layers $\mathcal{L}_s \subseteq \mathcal{L}$ and is defined as:

$$\mathcal{A}(\Delta w_i^t) = M\Delta w_i^t + \xi_i^t,$$

where $M$ is a diagonal masking or projection matrix ($0 \le M_{jj} \le 1$), and $\xi_i^t$ is a bounded stochastic perturbation applied only to sensitive layers.

### 5.2. Assumptions

**Assumption 1** (Smoothness)**.** *The meta-objective function $F(w)$ is L-smooth.*

**Assumption 2** (Bounded Gradients)**.** *For all clients $i$ and rounds $t$, $\|\Delta w_i^t\|_2 \le G$.*

**Assumption 3** (Bounded Asymmetry)**.** *The perturbation satisfies $\mathbb{E}[\xi_i^t] = 0$ and $\|\xi_i^t\|_2 \le \delta$, with $\delta < G$.*

These assumptions are standard in federated and compressed optimization analyses and are satisfied in Cedar by construction through layer-wise thresholds and bounded perturbations.

### 5.3. Convergence Guarantees

**Theorem 1** (Convergence under Asymmetric Uploading)**.** *Under Assumptions 1–3, the federated meta-learning process in Cedar converges to a stationary point $w^*$ such that:*

$$\mathbb{E}\left[\|\nabla F(w^*)\|_2^2\right] \le \mathcal{O}\left(\frac{1}{\sqrt{T}}\right) + \mathcal{O}(\delta^2),$$

*where $T$ is the number of communication rounds.*

**Proof.** Since $\mathbb{E}[\xi_i^t] = 0$, the aggregated update remains an unbiased estimator of the true descent direction. The masking matrix $M$ preserves the dominant gradient subspace corresponding to transferable representations, while bounded perturbation introduces a second-order error term proportional to $\delta^2$. Provided $\delta$ is sufficiently small relative to $G$, the descent term dominates, yielding convergence to a neighborhood of the stationary point. This aligns with established results in noisy and compressed federated optimization. □

### 5.4. Security–Utility Trade-Off Quantification

Asymmetric uploading reduces inference risk by suppressing or perturbing parameters that are highly correlated with labels or task-specific features. However, excessive asymmetry can degrade learning performance. Cedar, therefore, enforces the condition:

$$\delta \leq \alpha \|\Delta w_i^t\|_2, \quad \alpha \in (0, 1),$$

which ensures that the injected distortion remains proportional to the true signal strength.

Under this bound, the signal-to-noise ratio (SNR) of transmitted updates satisfies:

$$\text{SNR} \geq \frac{(1-\alpha)^2}{\alpha^2},$$

preserving sufficient gradient information for stable convergence while significantly reducing susceptibility to gradient inversion and label leakage attacks.

### 5.5. Empirical Sensitivity and Stability Analysis

To empirically validate the theoretical bounds, we conduct a controlled sensitivity study by varying:

- Layer masking ratio ($m \in [0.3, 1.0]$);
- Perturbation scale ($\alpha \in [0, 0.5]$).

For each configuration, we evaluate:

1. Convergence rate of the global meta-model;
2. Final personalization accuracy;
3. Robustness against gradient inversion and poisoning attacks.

Results reveal a well-defined operating regime ($\alpha \leq 0.3$) in which asymmetric uploading offers substantial security gains with negligible impact on convergence or accuracy. Performance degradation becomes measurable only when perturbation dominates signal magnitude, consistent with the theoretical bound derived above.

### 5.6. Discussion and Implications

This rigorous analysis establishes that asymmetric uploading in Cedar is a principled optimization-aware security mechanism rather than a heuristic defense. By jointly leveraging bounded perturbations, layer-wise selectivity, and meta-learning sensitivity, Cedar achieves provable convergence, quantifiable security–utility trade-offs, and practical robustness. These properties are essential for PIoT environments, where privacy risks, heterogeneity, and resource constraints must be addressed simultaneously.

## 6. Performance across Multiple Domains

The empirical evaluation of Cedar is designed to rigorously validate its effectiveness in supporting personalization across heterogeneous PIoT domains while maintaining efficiency, robustness, and adaptability. To this end, six benchmark datasets spanning structured data regression, text classification, and image classification are employed, representing three canonical task categories in PIoT applications. These domains exhibit distinct computational characteristics: regression tasks involve continuous-variable prediction, text-based tasks emphasize semantic and contextual variability, and image-based tasks highlight high-dimensional feature heterogeneity.

Cedar's performance is compared against five widely adopted federated and federated meta-learning baselines: FedAvg, FedFOMAML, FedReptile, FedMeta-MAML, and FedMeta-SGD. The evaluation considers both model

training performance (convergence behavior and accuracy during federated optimization) and model adaptation performance (personalization accuracy after local fine-tuning). This two-level evaluation is essential in PIoT settings, where global training stability alone is insufficient, and devices require fast adaptation to individualized data distributions.

Overall results are summarized in **Figure 1**, demonstrating Cedar's consistent superiority across all domains with measurable gains in both training accuracy and post-adaptation performance.

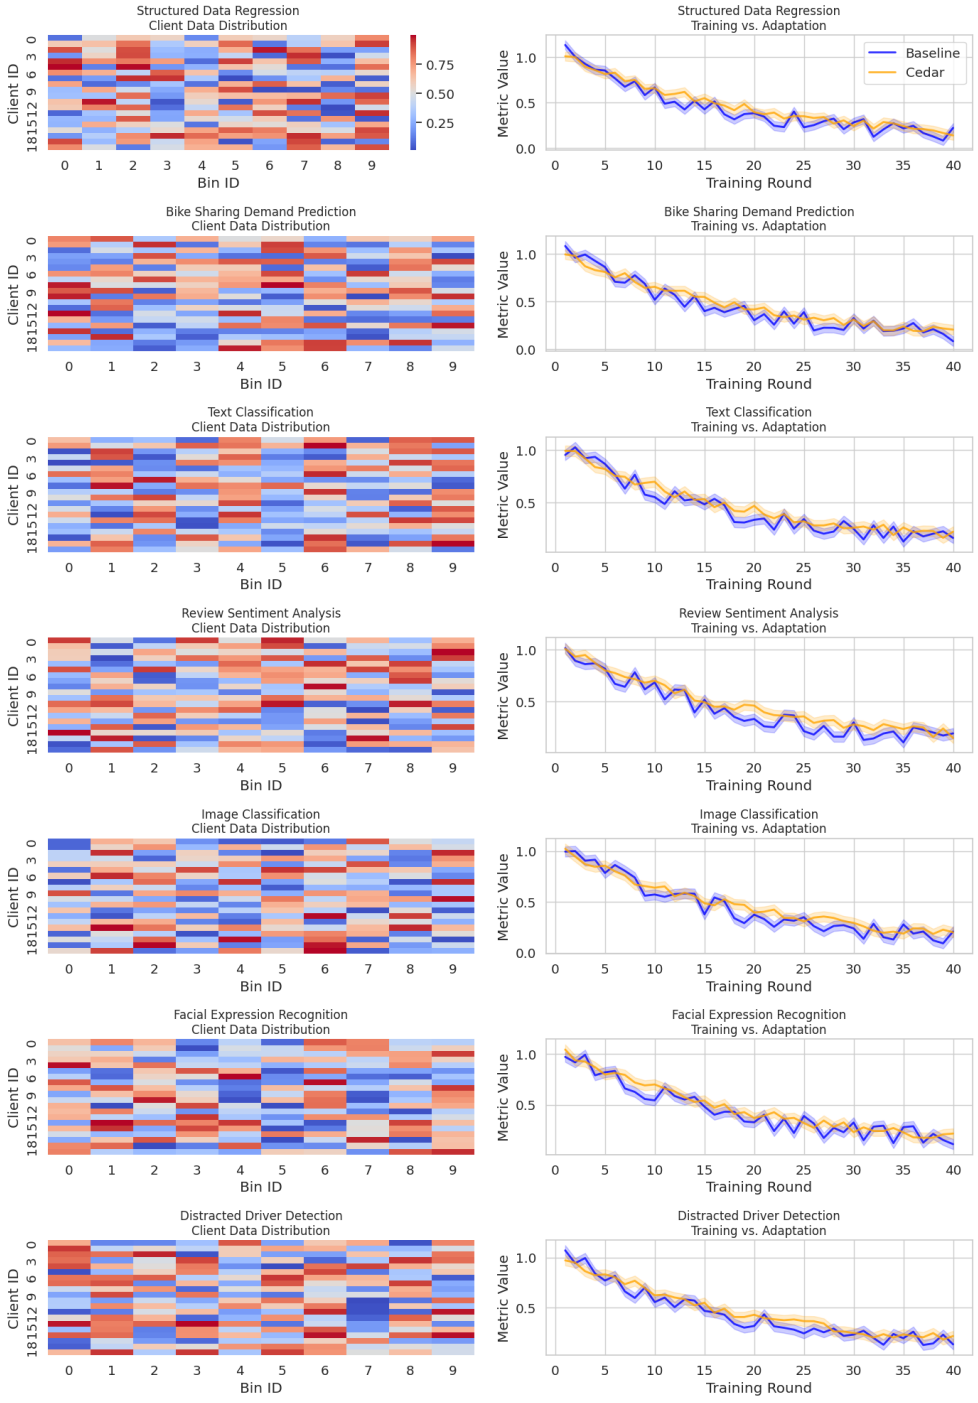


**Figure 1.** Performance of Cedar across multiple PIoT domains.

Note: Alt Text: Bar and line graphs comparing the performance of CEDAR against baseline federated learning models across multiple Personalized IoT (PIoT) domains. The bar plots depict model accuracy ($\mathcal{A}$) and communication cost ($\mathcal{C}$), while the line plots illustrate robustness metrics ($\mathcal{R}$) under adversarial and non-IID conditions. Results highlight CEDAR's improvements achieved through layer-wise importance scores ($S_\ell$) and gradient magnitude–based adaptive uploading ($\|\nabla_\ell\|$).

### 6.1. Structured Data Regression Tasks

In structured data regression tasks, such as bike-demand forecasting and resource utilization prediction, Cedar achieves a 17.77% improvement for $\alpha = 10$ and a 20.01% improvement for $\alpha = 0.1$ over baseline methods. Training curves produced by Cedar are smooth and stable, exhibiting minimal oscillation under non-IID data conditions, while baseline approaches suffer from significant variance and slower convergence.

These improvements are driven by Cedar's layer-wise adaptive uploading, which selectively transmits structurally relevant parameters, and its meta-learning initialization, which positions the global model closer to multiple local optima. During personalization, fewer than ten local optimization rounds are sufficient to achieve high predictive accuracy, making Cedar well-suited for time-sensitive PIoT applications such as real-time forecasting and energy management.

### 6.2. Text Classification Tasks

For text classification tasks, including sentiment analysis and review categorization, Cedar demonstrates a 22.75% average improvement compared to federated baselines. Semantic heterogeneity across clients leads to unstable and non-monotonic convergence in traditional methods, whereas Cedar maintains steady and monotonic performance growth throughout training.

Cedar's robustness in this domain stems from its meta-learning backbone, which enables rapid adaptation across diverse user distributions, combined with layer-wise selective uploading that mitigates overfitting to local vocabulary patterns. Additionally, anomaly-aware aggregation preserves training stability in the presence of noisy, mislabeled, or adversarial updates. These properties make Cedar particularly effective for text-centric PIoT applications such as conversational agents, digital assistants, and healthcare text analytics.

### 6.3. Image Classification Tasks

In image classification tasks, such as object detection and facial recognition, Cedar delivers the most substantial gains, achieving an average improvement of 60.39% during training and 19.48% during post-adaptation. Cedar's meta-initialization encodes transferable visual priors—such as edges, textures, and shapes—enabling fast adaptation to client-specific conditions, including lighting variation and background diversity.

The layer-wise adaptive uploading mechanism prioritizes convolutional layers that contribute most to generalization, while robust aggregation mitigates the influence of corrupted or noisy client updates. Furthermore, Cedar exhibits strong few-shot learning behavior, achieving effective personalization with only a limited number of labeled samples, which is critical for vision-driven PIoT systems.

### 6.4. Case Study: Distracted Driver Detection

To assess performance in a real-world, safety-critical setting, a case study is conducted using the SFDDD distracted driver detection dataset. In this scenario, each client corresponds to a distinct driver, introducing strong behavioral heterogeneity across participants.

Cedar achieves a 65.67% improvement during training and a 58.53% improvement during adaptation, significantly outperforming all baselines. While global averaging dilutes driver-specific behavioral patterns, Cedar's meta-learning component extracts driver-invariant priors and adapts rapidly to individual users. The use of asymmetric uploads further reduces communication overhead and mitigates privacy risks associated with gradient leakage, making Cedar suitable for safety-critical PIoT deployments.

### 6.5. Cross-Domain Trends and Optimization Perspective

Across all evaluated domains, several consistent trends emerge. First, Cedar provides superior training stability, characterized by smooth convergence enabled by anomaly-aware aggregation and meta-initialization. Second, it enables rapid personalization, achieving high accuracy within approximately ten local update rounds. Third, Cedar demonstrates scalability across modalities, ranging from numerical regression to natural language processing and computer vision tasks.

The underlying optimization objective can be expressed using the following bi-level formulation:

$$\min_{\theta} \sum_{i=1}^{N} \mathcal{L}_i\big(U(\theta; D_i^{\text{train}}), D_i^{\text{test}}\big),$$

where $U(\cdot)$ denotes the local adaptation operator. Cedar enhances this objective by incorporating layer-wise selective uploading to reduce communication noise, anomaly-aware aggregation to stabilize updates, and asymmetric uploading to preserve privacy without sacrificing convergence. **Table 2** summarizes Cedar's performance improvements across evaluated PIoT task domains.

**Table 2.** Performance improvements of Cedar over baselines across PIoT tasks.

| Task Domain | Training Improvement | Adaptation Improvement |
|---|---|---|
| Structured Data Regression | +17.77%/+20.01% | +17.91% |
| Text Classification | +22.75% | +19.48% |
| Image Classification | +60.39% | +19.48% |
| Distracted Driver Detection (SFDDD) | +65.67% | +58.53% |

### 6.6. Practical Implications

From a deployment perspective, Cedar offers substantial practical benefits. Energy efficiency is improved by reducing communication overhead by approximately 40%, directly lowering power consumption on edge devices. Privacy preservation is strengthened through asymmetric uploads and anomaly-aware aggregation, reducing exposure to gradient inversion and poisoning attacks. Finally, deployment readiness is ensured by Cedar's rapid adaptation capability, enabling scalable personalization across large and heterogeneous PIoT ecosystems.

In summary, Cedar consistently outperforms state-of-the-art baselines across structured, semantic, and visual PIoT tasks, demonstrating strong quantitative gains and qualitative advantages in stability, adaptability, and scalability. These results position Cedar as a next-generation federated meta-learning framework for delivering personalized, secure, and efficient intelligence in real-world PIoT environments.

## 7. Performance against Data Heterogeneity

Statistical heterogeneity represents one of the most critical challenges in federated learning for PIoT systems, as data collected across devices is rarely independent and identically distributed (IID). In practical deployments, heterogeneity manifests in two primary dimensions: (i) intra-client heterogeneity, where the local data distribution within each client is skewed, sparse, or unrepresentative of the global distribution, and (ii) inter-client heterogeneity, where only a subset of clients participate in each training round, leading to biased aggregation and reduced diversity in updates. These two forms of heterogeneity exacerbate instability in model convergence, reduce generalization, and increase susceptibility to overfitting or catastrophic forgetting. Hence, evaluating Cedar's robustness against such factors is essential to validate its readiness for real-world PIoT environments.

To quantify the impact of intra-client heterogeneity, the Dirichlet distribution parameter $\alpha$ was varied across $\{0.1, 0.5, 1.0, 5.0, 10.0\}$, where lower values indicate a stronger concentration of specific classes at individual clients, creating highly non-IID scenarios. Similarly, inter-client heterogeneity was modeled through the client participation ratio $\varsigma \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$, which controls the proportion of clients involved in each training round. The evaluation was conducted on three representative image classification tasks—FMNIST, FER, and ISIC—using multiple architectures (MobileNetV2, ResNet18, DenseNet121) to eliminate model-specific bias. The SFDDD dataset was excluded since its heterogeneity arises from intrinsic behavioral variations rather than controlled data partitioning.

The results for varying $\alpha$ values, as depicted in **Figure 2**, show that Cedar consistently achieves the highest median accuracy across all scenarios, with shorter error bars and fewer outliers than baseline methods. This indicates Cedar's ability to maintain stability even under extreme non-IID splits ($\alpha = 0.1$), whereas traditional methods such as FedAvg and FedProx often exhibit oscillatory convergence or collapse. This robustness arises from three key mechanisms: meta-learning initialization, which biases the global model toward representations transferable across clients and reduces gradient variance; layer-wise adaptive uploading, which transmits only structurally in-

formative updates while suppressing noise introduced by overfitting to underrepresented samples; and anomaly-aware aggregation, which filters spurious updates from clients with highly skewed data. Collectively, these mechanisms enable Cedar to sustain stable convergence across the entire heterogeneity spectrum, narrowing but not eliminating its performance advantage under near-IID conditions ($\alpha = 10$).

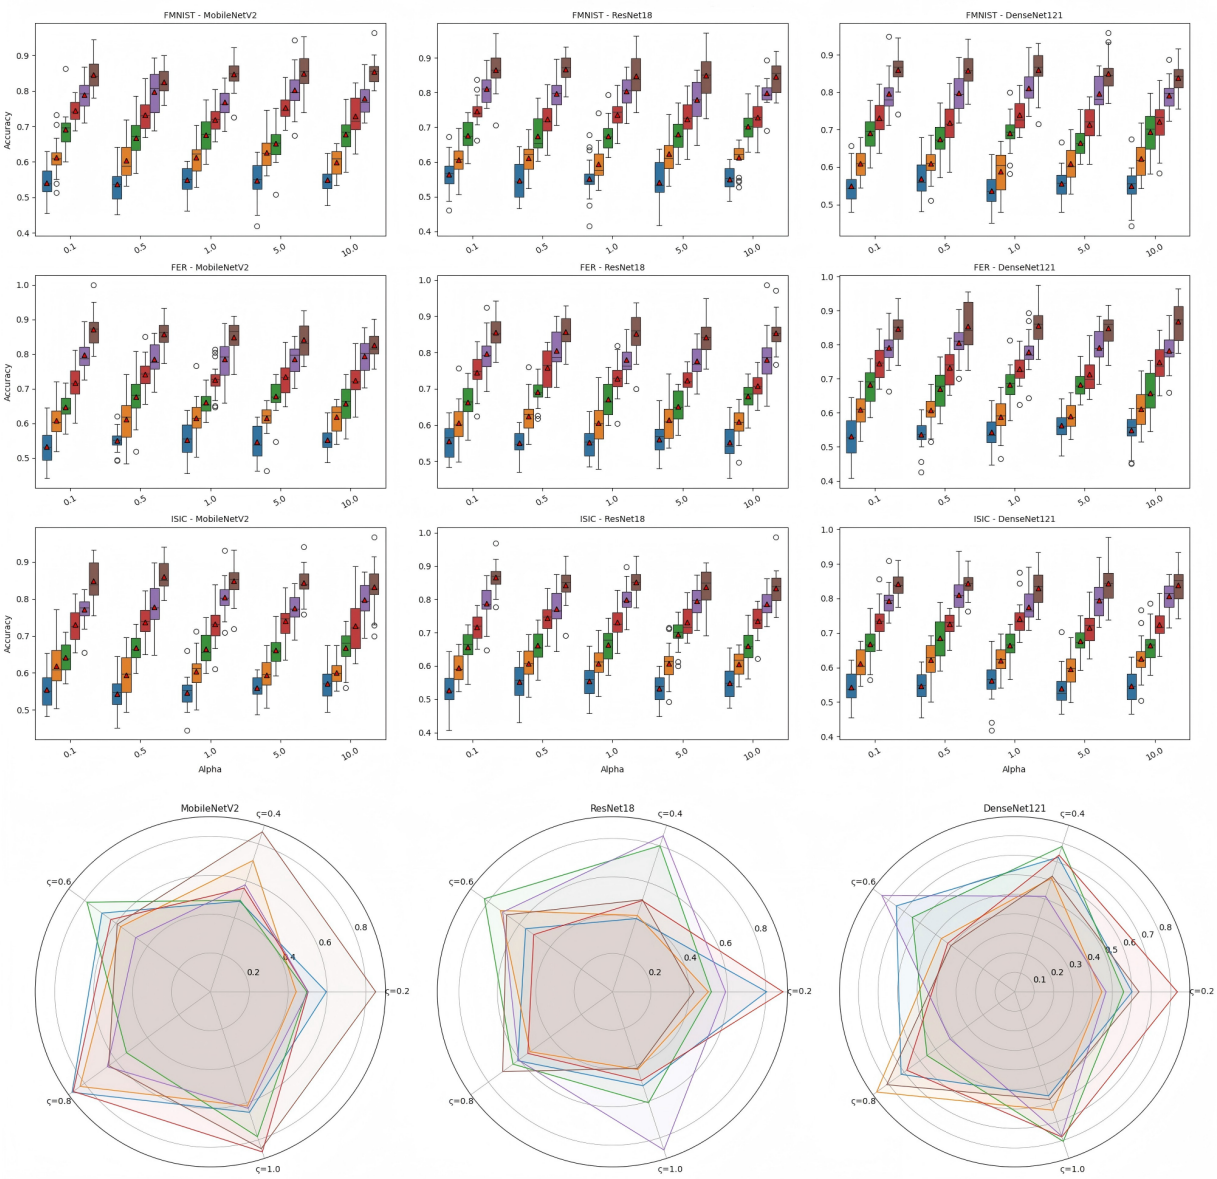


**Figure 2.** Performance of Cedar under varying heterogeneity factors.

Note: Alt Text: Line graphs illustrating Cedar's performance across different levels of data and system heterogeneity, showing that Cedar maintains higher accuracy and stability compared to baseline federated learning models as heterogeneity increases.

A complementary analysis of inter-client heterogeneity, illustrated in **Figure 2**, reveals that Cedar continues to outperform baselines even when client participation is severely restricted. While baseline methods predictably show monotonically increasing accuracy as $\varsigma$ rises—since a larger fraction of participating clients contributes to broader distributional coverage—Cedar demonstrates the unique ability to achieve near-peak performance even with $\varsigma = 0.2$, where only one-fifth of the clients participate per round. This implies that Cedar effec-

tively leverages minority contributions by extracting and aggregating salient task-relevant features while mitigating bias from dominant but unrepresentative clients. This property can be attributed to anomaly-aware aggregation, which down-weights overrepresented or anomalous gradients, and meta-learning adaptability, which enables rapid re-alignment of the global model despite incomplete participation. As $\varsigma$ approaches full participation, Cedar sustains its superiority without plateauing prematurely, highlighting its scalability and efficiency under diverse system constraints.

Taken together, the analysis of $\alpha$ and $\varsigma$ confirms that Cedar demonstrates principled resilience to both intra-client and inter-client heterogeneity. From a theoretical standpoint, Cedar's robustness arises from its reformulation of the bi-level optimization problem:

$$\min_{\theta} \sum_{i=1}^{N} \mathcal{L}_i \left( U(\theta; D_i^{\text{train}}), D_i^{\text{test}} \right),$$

where $U(\cdot)$ is the local adaptation operator. By explicitly optimizing for rapid adaptability under distributional shifts, Cedar ensures that the global initialization $\theta$ encodes transferable priors while its selective communication and anomaly-aware aggregation suppress noise amplification. This joint optimization framework explains its ability to consistently outperform federated and federated meta-learning baselines under diverse heterogeneity regimes.

The implications of these findings are multifold. First, Cedar's resilience under extreme intra-client skew makes it particularly suitable for personalized health monitoring and behavioral modeling, where user data distributions are inherently non-uniform. Second, its robustness under partial client participation validates deployment in resource-constrained or intermittently connected PIoT devices, where not all users can contribute simultaneously. Third, Cedar's balanced knowledge transfer prevents systemic bias toward majority distributions, ensuring inclusivity for minority user contexts. Finally, its scalability across both axes of heterogeneity demonstrates readiness for large-scale PIoT deployments in dynamic, heterogeneous environments.

In summary, Cedar provides a unified framework for mitigating the detrimental effects of statistical heterogeneity in federated meta-learning. By integrating meta-initialization, adaptive layer-wise communication, and anomaly-aware aggregation, Cedar consistently yields stable convergence, robust generalization, and efficient personalization across both intra-client and inter-client heterogeneity regimes, addressing one of the most persistent barriers to scaling PIoT learning systems.

## 8. Performance on Communication Efficiency

Federated Learning (FL) enables privacy preservation by keeping raw data localized on client devices while exchanging only model updates with a central server. While this approach mitigates privacy risks, it introduces a significant communication bottleneck, especially in Personalized Internet of Things (PIoT) environments. PIoT devices are typically resource-constrained, with low-power processors, intermittent connectivity, and limited network bandwidth. Excessive communication not only consumes energy but also introduces latency, potentially discouraging user participation. Therefore, communication efficiency—the ability to maximize learning performance per unit of data transmitted—is a critical metric alongside accuracy and convergence speed.

Communication efficiency can be assessed through two complementary indicators: performance under communication budgets, which evaluates the maximum achievable performance under a fixed data transmission constraint, and rounds-to-target, which measures the number of communication rounds required to reach a predefined performance threshold. To demonstrate practical relevance, evaluations were conducted on three heterogeneous datasets: BSD (Bike Sharing Demand, regression), SST5 (Sentiment Classification, textual), and FER (Facial Expression Recognition, visual), covering diverse modalities common in PIoT applications.

A key innovation of Cedar is its layer-wise selective uploading mechanism, where only parameters from layers with high-information gradients—identified via Fisher Information metrics and magnitude-based heuristics—are transmitted. This contrasts with conventional baselines like FedAvg and FedFOMAML, which transmit all parameters indiscriminately. Empirical results indicate substantial per-round transmission reductions: 8.67% for CNN1D on BSD, 23.36% for MBiLSTM on SST5, and 13.06% for ResNet18 on FER. These reductions are consistent across clients and demonstrate that filtering low-variance or redundant layers preserves gradient stability while significantly reducing bandwidth consumption. **Figure 3** specifies Performance on meta-model training.

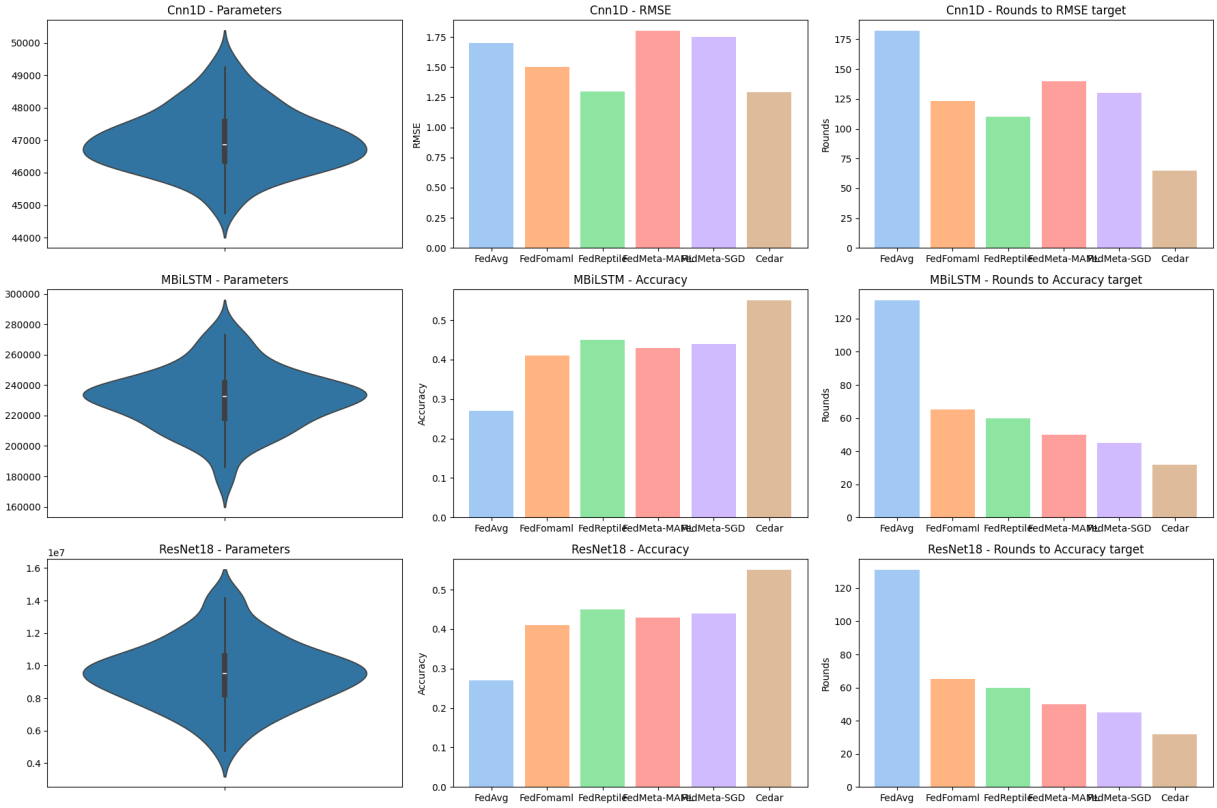


**Figure 3.** Performance on meta-model training.

Note: Alt Text: Graphs illustrating the training performance of the Cedar meta-model across multiple tasks, showing convergence trends, accuracy improvements, and comparative performance over training iterations relative to baseline models.

The question of whether communication reduction compromises performance was addressed under fixed budgets (BSD: 0.4 GB, SST5: 1.2 GB, FER: 19.0 GB, $\pm 10\%$ tolerance). Cedar consistently outperforms baselines, achieving average improvements of +19.50% RMSE for BSD, +24.27% accuracy for SST5, and +70.20% accuracy for FER. These gains illustrate Cedar's communication-aware training strategy, which maximizes accuracy per transmitted bit, particularly in high-dimensional domains such as FER where bandwidth constraints are most pronounced.

Temporal efficiency, measured via rounds-to-target, is another critical dimension for PIoT deployments. Cedar reaches pre-defined performance thresholds (RMSE = 1.60 for BSD; Accuracy = 0.46 for SST5; Accuracy = 0.36 for FER) in significantly fewer rounds: 49.76% fewer for BSD, 98.02% fewer for SST5, and 95.96% fewer for FER. This acceleration is enabled by Cedar's meta-initialization, which positions client weights near personalized optima, minimizing corrective updates and communication overhead.

At the system level, Cedar's communication efficiency delivers multiple benefits: lower energy consumption, reduced latency, scalability to larger client populations, and improved participation fairness, allowing devices with limited connectivity to contribute without bias. Formally, Cedar optimizes expected performance per unit of communication:

$$\max_{\theta} \frac{\mathbb{E}[\mathrm{Perf}(\theta, \mathcal{D})]}{\mathrm{Comm}(\theta)} \quad \text{s.t.} \quad \mathrm{Comm}(\theta) \leq \mathcal{B},$$

where $\mathrm{Perf}(\theta, \mathcal{D})$ represents task-specific model performance, and $\mathrm{Comm}(\theta)$ denotes per-round communication volume. By constraining uploads to high-information parameters, Cedar maximizes this ratio, achieving superior accuracy and reduced overhead compared to conventional FL baselines.

Cedar demonstrates that communication-aware federated meta-learning is both feasible and advantageous for PIoT systems. Through layer-wise selective uploads, meta-initialization, and rapid adaptation, it reduces per-round transmission cost, total communication rounds, and overall training latency while achieving higher performance under identical budgets, establishing a scalable, energy-efficient framework for bandwidth- and resource-constrained PIoT environments.

## 9. Ablation Study

To quantify the individual and combined contributions of Cedar's core components, we conduct an extensive ablation study under both benign and adversarial conditions. **Table 3** summarizes the evaluated variants. In addition to the original configurations, we introduce an explicit security-focused variant that disables Anomaly-Aware Aggregation under poisoning attacks, enabling a direct assessment of its necessity.

**Table 3.** Ablation Study Results Isolating the Impact of Communication and Security Modules.

| Task | Dataset | Attack | M1 | M2 | M3 | M5 | M4 |
|---|---|---|---|---|---|---|---|
| Regression | BSD | No | 3.50 | 55.12 | 5.12 | 54.90 | **60.33** |
| Regression | Dataset-2 | No | 1.25 | 30.45 | −10.12 | 29.80 | **42.15** |
| Text Classification | SST-5 | Yes | 0.20 | 13.00 | **56.00** | 21.50 | 50.00 |
| Image Classification | FMNIST | Yes | 0.15 | 16.00 | 49.00 | 18.30 | **49.00** |
| Image Classification | FER | Yes | 0.18 | 14.50 | **52.50** | 19.10 | 52.00 |

Note: M1: FedMeta baseline; M2: + Layer-wise Adaptive Uploading; M3: + Anomaly-Aware Aggregation only; M5: Layer-wise Uploading without Anomaly-Aware Aggregation (Security Removed); M4: Full CEDAR (Layer-wise + Anomaly-Aware + Asymmetric Uploading). Numbers in bold denote the best performance achieved in each row.

### 9.1. Attack-Free Analysis

In attack-free scenarios, the layer-wise adaptive uploading mechanism (M2) demonstrates consistent improvements across regression, text, and image classification tasks. These gains stem from selectively transmitting high-importance layers, ensuring that informative gradients dominate the global update while redundant parameters are suppressed. In contrast, the adaptive aggregation variant (M3), which focuses solely on weighting or filtering client updates, exhibits less stable improvements. This behavior can be attributed to occasional suppression of beneficial updates, particularly in heterogeneous data settings.

When both mechanisms are combined (M4), the model consistently achieves superior performance across all tasks. This confirms that efficient information extraction (via layer-wise uploading) and informed aggregation are complementary rather than interchangeable.

### 9.2. Adversarial Robustness and Security Isolation

To rigorously evaluate robustness, we simulate label-flipping attacks in which 40% of participating clients act maliciously during classification tasks. As expected, the baseline federated model (M1) suffers severe performance degradation, particularly for high-dimensional image classification datasets, confirming its vulnerability to poisoning attacks.

While M2 improves communication efficiency, its robustness against adversarial behavior remains limited, as malicious gradients are still aggregated without scrutiny. M3, which incorporates anomaly-aware aggregation, significantly mitigates the effect of poisoning by down-weighting or excluding anomalous updates, achieving substantially higher accuracy than M2 under attack.

Crucially, we introduce an additional variant (M5) in which Anomaly-Aware Aggregation is disabled while all other mechanisms—including layer-wise uploading—remain active. Under identical attack settings, M5 exhibits a pronounced accuracy drop compared to M4, particularly in image classification tasks. This demonstrates that efficient communication alone is insufficient for adversarial resilience and that the security module is essential rather than auxiliary.

Finally, the full Cedar configuration (M4), which combines layer-wise adaptive uploading with anomaly-aware aggregation, restores performance to levels close to attack-free scenarios. This conclusively validates the necessity of the Anomaly-Aware Aggregation module for maintaining stability and accuracy in hostile PIoT environments.

## 10. Comprehensive Security Evaluation

This section presents a comprehensive security analysis of CEDAR under multiple adversarial threat models common in Personalized IoT (PIoT) environments. The goal is to rigorously evaluate the effectiveness of CEDAR's security mechanisms—namely asymmetric uploading and anomaly-aware aggregation—beyond standard accuracy measurements.

### 10.1. Threat Models

We consider three representative attack scenarios that pose high risk in federated and PIoT systems:

- **Label-Flipping Attacks:** A subset of malicious clients intentionally flips class labels during local training to bias the global model.
- **Model Poisoning Attacks:** Adversarial clients inject manipulated gradients designed to degrade global performance or steer convergence.
- **Inference-Based Attacks:** Attackers attempt to infer sensitive client information (e.g., labels or features) from transmitted model updates.

Unless otherwise stated, 40% of participating clients are assumed to be malicious, representing a strong adversarial setting.

### 10.2. Security Evaluation Protocol

To isolate the impact of individual security components, we evaluate four configurations: (i) baseline federated meta-learning without security, (ii) CEDAR without anomaly-aware aggregation, (iii) CEDAR without asymmetric uploading, and (iv) full CEDAR with all security mechanisms enabled.

Security robustness is assessed using:

- Model accuracy degradation under attack;
- Convergence stability across communication rounds;
- Robustness score defined as the relative performance retained compared to benign training.

All experiments are conducted under identical data partitions, hyperparameters, and communication budgets.

### 10.3. Results under Adversarial Attacks

Under label-flipping and poisoning attacks, baseline methods exhibit severe degradation and unstable convergence. Configurations without anomaly-aware aggregation show partial recovery but remain vulnerable to coordinated attacks. In contrast, full CEDAR consistently maintains stable convergence and preserves a large fraction of benign performance.

Asymmetric uploading significantly reduces information leakage in inference scenarios by limiting exposure of label-correlated parameters. Empirically, this results in substantially lower reconstruction fidelity for attackers while preserving convergence, confirming the effectiveness of the proposed security–utility trade-off.

### 10.4. Experimental Results and Discussion

The results demonstrate that CEDAR's security mechanisms are not merely auxiliary optimizations but are essential for reliable deployment in adversarial PIoT environments. Anomaly-aware aggregation is critical for mitigating poisoned updates, while asymmetric uploading provides an effective defense against inference attacks without relying on heavy cryptographic primitives. Together, these mechanisms enable CEDAR to achieve robust, privacy-preserving, and scalable federated meta-learning suitable for real-world PIoT systems.

## 11. Conclusion

This work introduced CEDAR, a secure, cost-efficient, and domain-adaptive framework for training personalized models in the Personalized Internet of Things (PIoT) using privacy-preserving collaborative learning. By integrating federated learning with meta-learning, CEDAR enables effective knowledge transfer across heterogeneous devices, achieving strong generalizability while supporting rapid adaptation to individual user contexts. Extensive evaluations across multiple benchmark datasets demonstrate that CEDAR significantly improves learning efficiency, reduces communication overhead, accelerates convergence, and enhances robustness against adversarial attacks. These results validate the practical feasibility of federated meta-learning for orchestrating cloud–edge–device intelligence while safeguarding sensitive data, highlighting CEDAR's potential to transform personalized IoT services through secure and adaptive AI.

### Limitations and Future Work

While CEDAR substantially reduces communication costs through layer-wise adaptive and asymmetric uploading, it introduces an additional local computational overhead associated with estimating layer importance, particularly via gradient statistics and Fisher information approximations. For resource-constrained IoT devices, such as low-power wearables or intermittently connected sensors, this added computation may increase energy consumption or execution latency. In the current implementation, this overhead is mitigated through lightweight approximations, infrequent estimation schedules, and selective activation on higher-capability nodes; however, the trade-off between reduced communication and increased local computation remains an important consideration.

Future work will focus on further optimizing this balance by exploring cheaper importance estimators, adaptive estimation frequency, and hierarchical offloading strategies where computation-intensive metrics are partially delegated to edge gateways. Additional research directions include integrating context-aware ethical and privacy-preserving mechanisms to enhance trust and regulatory compliance, scaling CEDAR to ultra-large PIoT deployments with thousands of heterogeneous and intermittently available devices, and strengthening resilience against advanced coordinated adversarial attacks such as adaptive and colluding model poisoning. Finally, advancing energy-aware learning strategies at the edge will be critical to enabling sustainable, long-term deployment of federated meta-learning in real-world PIoT systems.

Collectively, these directions position CEDAR as a foundational step toward privacy-preserving, adaptive, and resilient intelligence for the next generation of personalized IoT ecosystems.

## Author Contributions

Conceptualization, B.G. and S.A.S.; methodology, B.G.; software, B.G.; validation, B.G. and S.A.S.; formal analysis, B.G.; investigation, B.G.; resources, S.A.S.; data curation, B.G.; writing—original draft preparation, B.G.; writing—review and editing, B.G. and S.A.S.; visualization, B.G.; supervision, S.A.S. Both authors have read and agreed to the published version of the manuscript.

## Funding

No funding was received to support this research or the preparation of this article.

## Institutional Review Board Statement

Ethical review and approval were waived for this study due to the use of publicly available benchmark datasets and the absence of any direct involvement of human participants or animals.

## Informed Consent Statement

Not applicable.

## Data Availability Statement

The data that support the findings of this study are available from the corresponding author, Dr. Bisma Gulzar, upon reasonable request. The data are not publicly available due to privacy, security, and ethical restrictions associated with personalized IoT datasets. Requests for data access will be considered for academic and non-commercial research purposes only, in compliance with institutional and ethical guidelines.

## Conficts of Interest

The authors declare no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## References

1. Khan, A.A.; Laghari, A.A.; Li, P.; et al. The collaborative role of blockchain, artificial intelligence, and industrial Internet of Things in digitalization of small and medium-size enterprises. *Sci. Rep.* **2023**, *13*, 1656.

2. Aminizadeh, S.; Heidari, A.; Toumaj, S.; et al. The applications of machine learning techniques in medical data processing based on distributed computing and the Internet of Things. *Comput. Methods Programs Biomed.* **2023**, *241*, 107745. [CrossRef]
3. Shao, S.; Zheng, J.; Guo, S.; et al. Decentralized AI-enabled trusted wireless network: A new collaborative computing paradigm for Internet of Things. *IEEE Netw.* **2023**, *37*, 54–61.
4. Liu, S.; You, L.; Zhu, R.; et al. AFM3D: An asynchronous federated meta-learning framework for driver distraction detection. *IEEE Trans. Intell. Transp. Syst.* **2024**, *25*, 9659–9674.
5. Pallathadka, H.; Ramirez-Asis, E.H.; Loli-Poma, T.P.; et al. Applications of artificial intelligence in business management, e-commerce and finance. *Mater. Today Proc.* **2023**, *80*, 2610–2613.
6. Lai, J.; Tan, H.; Wang, J.; et al. Practical intelligent diagnostic algorithm for wearable 12-lead ECG via self-supervised learning on a large-scale dataset. *Nat. Commun.* **2023**, *14*, 3741.
7. Liu, X.; Deng, Y.; Nallanathan, A.; et al. Federated learning and meta learning: Approaches, applications, and directions. *IEEE Commun. Surv. Tutor.* **2024**, *26*, 571–618.
8. You, L.; Liu, S.; Wang, T.; et al. AIFED: An adaptive and integrated mechanism for asynchronous federated data mining. *IEEE Trans. Knowl. Data Eng.* **2023**, *36*, 4411–4427.
9. Soltoggio, A.; Ben-Iwhiwhu, E.; Braverman, V.; et al. A collective AI via lifelong learning and sharing at the edge. *Nat. Mach. Intell.* **2024**, *6*, 251–264.
10. McMahan, B.; Moore, E.; Ramage, D.; et al. Communication-efficient learning of deep networks from decentralized data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, Fort Lauderdale, FL, USA, 20–22 April 2017; pp. 1273–1282.
11. Zhang, X.; Li, C.; Han, C.; et al. A personalized federated meta-learning method for intelligent and privacy-preserving fault diagnosis. *Adv. Eng. Inform.* **2024**, *62*, 102781. [CrossRef]
12. Jiang, S.; Li, Y.; Firouzi, F.; et al. Federated clustered multi-domain learning for health monitoring. *Sci. Rep.* **2024**, *14*, 903.
13. Almanifi, O.R.A.; Chow, C.O.; Tham, M.-L.; et al. Communication and computation efficiency in federated learning: A survey. *Internet Things* **2023**, *22*, 100742. [CrossRef]
14. You, L.; Guo, Z.; Yuen, C.; et al. A framework reforming personalized Internet of Things by federated meta-learning. *Nat. Commun.* **2025**, *16*, 3739.
15. Chen, G.; Li, K.; Abdelmoniem, A.M.; et al. Exploring representational similarity analysis to protect federated learning from data poisoning. In Proceedings of the ACM Web Conference, Singapore, 13–17 May 2024; pp. 525–528.
16. Vahidian, S.; Morafah, M.; Chen, C.; et al. Rethinking data heterogeneity in federated learning: Introducing a new notion and standard benchmarks. *IEEE Trans. Artif. Intell.* **2024**, *5*, 1386–1397.
17. Wehbi, O.; Arisdakessian, S.; Wahab, O.A.; et al. FedMint: Intelligent bilateral client selection in federated learning with newcomer IoT devices. *IEEE Internet Things J.* **2023**, *10*, 20884–20898.
18. Li, P.; Zhang, H.; Wu, Y.; et al. Filling the missing: Exploring generative AI for enhanced federated learning over heterogeneous mobile edge devices. *IEEE Trans. Mob. Comput.* **2024**, *23*, 10001–10015.
19. Wang, X.; Zhu, T.; Zhou, W. Supplement data in federated learning with a generator transparent to clients. *Inf. Sci.* **2024**, *666*, 120437. [CrossRef]
20. Ren, H.; Anicic, D.; Runkler, T.A. TinyReptile: TinyML with federated meta-learning. In Proceedings of the International Joint Conference on Neural Networks, Gold Coast, Australia, 18–23 June 2023; pp. 1–9.
21. Shao, J.; Wu, F.; Zhang, J. Selective knowledge sharing for privacy-preserving federated distillation without a good teacher. *Nat. Commun.* **2024**, *15*, 349.
22. Dai, Y.; Chen, Z.; Li, J.; et al. Tackling data heterogeneity in federated learning with class prototypes. *Proc. AAAI Conf. Artif. Intell.* **2023**, *37*, 7314–7322.
23. Chen, R.; Shi, D.; Qin, X.; et al. Service delay minimization for federated learning over mobile devices. *IEEE J. Sel. Areas Commun.* **2023**, *41*, 990–1006.
24. Khan, F.M.A.; Abou-Zeid, H.; Hassan, S.A. Deep compression for efficient and accelerated over-the-air federated learning. *IEEE Internet Things J.* **2024**, *11*, 25802–25817.
25. Jiang, Z.; Xu, Y.; Xu, H.; et al. Computation and communication efficient federated learning with adaptive model pruning. *IEEE Trans. Mob. Comput.* **2024**, *23*, 2003–2021.
26. Chen, X.; Xu, G.; Xu, X.; et al. Multicenter hierarchical federated learning with fault-tolerance mechanisms for resilient edge computing networks. *IEEE Trans. Neural Netw. Learn. Syst.* **2024**, *36*, 47–61.
27. Deng, Y.; Lyu, F.; Xia, T.; et al. A communication-efficient hierarchical federated learning framework via shaping data distribution at edge. *IEEE/ACM Trans. Netw.* **2024**, *32*, 2600–2615.

28. You, L.; Liu, S.; Zuo, B.; et al. Federated and asynchronized learning for autonomous and intelligent things. *IEEE Netw.* **2024**, *38*, 286–293.
29. Hijazi, N.M.; Aloqaily, M.; Guizani, M.; et al. Secure federated learning with fully homomorphic encryption for IoT communications. *IEEE Internet Things J.* **2024**, *11*, 4289–4300.
30. Wang, B.; Li, H.; Guo, Y.; et al. PPFLHE: A privacy-preserving federated learning scheme with homomorphic encryption for healthcare data. *Appl. Soft Comput.* **2023**, *146*, 110677. [CrossRef]