



Delta: A Cloud-assisted Data Enrichment Framework for On-Device Continual Learning

Chen Gong
Shanghai Jiao Tong University

Zhenzhe Zheng
Shanghai Jiao Tong University

Fan Wu
Shanghai Jiao Tong University

Xiaofeng Jia
Beijing Big Data Centre

Guihai Chen
Shanghai Jiao Tong University

Abstract

In modern mobile applications, users frequently encounter various new contexts, necessitating on-device continual learning (CL) to ensure consistent model performance. While existing research predominantly focused on developing lightweight CL frameworks, we identify that data scarcity is a critical bottleneck for on-device CL. In this work, we explore the potential of leveraging abundant cloud-side data to enrich scarce on-device data, and propose a private, efficient and effective data enrichment framework *Delta*. Specifically, *Delta* first introduces a directory dataset to decompose the data enrichment problem into device-side and cloud-side sub-problems without sharing sensitive data. Next, *Delta* proposes a soft data matching strategy to effectively solve the device-side sub-problem with sparse user data, and an optimal data sampling scheme for cloud server to retrieve the most suitable dataset for enrichment with low computational complexity. Further, *Delta* refines the data sampling scheme by jointly considering the impact of enriched data on both new and past contexts, mitigating the catastrophic forgetting issue from a new aspect. Comprehensive experiments across four typical mobile computing tasks with varied data modalities demonstrate that *Delta* could enhance the overall model accuracy by an average of 15.1%, 12.4%, 1.1% and 5.6% for visual, IMU, audio and textual tasks compared with few-shot CL, and consistently reduce the communication costs by over 90% compared to federated CL.

*Zhenzhe Zheng is the corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. *ACM MobiCom '24, November 18–22, 2024, Washington D.C., DC, USA*
© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0489-5/24/11

<https://doi.org/10.1145/3636534.3690701>

CCS Concepts

• **Human-centered computing** → **Mobile computing**; • **Computing methodologies** → **Machine learning**.

Keywords

Continual Learning, On-Device Training, Data Enrichment

ACM Reference Format:

Chen Gong, Zhenzhe Zheng, Fan Wu, Xiaofeng Jia, and Guihai Chen. 2024. Delta: A Cloud-assisted Data Enrichment Framework for On-Device Continual Learning. In *The 30th Annual International Conference on Mobile Computing and Networking (ACM MobiCom '24)*, November 18–22, 2024, Washington D.C., DC, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3636534.3690701>

1 Introduction

Machine learning (ML) models have become the indispensable components in modern mobile applications and services, such as image tagging in Google Smart Lens [3], speech recognition in Siri [30], text summarization and rewriting in Apple Intelligence [1] and etc. In a wide range of mobile applications, users encounter dynamic contexts in their daily lives and exhibit varying behaviors, leading to a non-stationary data distribution observed and collected by mobile devices. Consequently, on-device ML models are expected to evolve incrementally as new contextual data becomes available. This evolution, known as *continual learning (CL)* [66, 70], enables on-device ML models to gradually learn individual user preferences in different contexts and behaviors, and thus becoming more personalized and intelligent over time.

Unlike conventional ML built on the premise of learning static data distributions, CL involves learning from dynamic data distributions. A significant challenge in CL is balancing the model's learning plasticity (*i.e.* ability to assimilate new knowledge from emerging context) and memory stability (*i.e.* ability to preserve past knowledge from historical contexts). For cloud servers with abundant hardware and data resources, many CL approaches have been proposed to address this challenge, such as regularizing model parameter updates [34, 82], replaying historical data [15, 42, 44] and designing context-adaptive model architectures [38, 48, 60]. For resource-constrained devices, previous research focused

on optimizing the usage of limited hardware resources to facilitate the efficient on-device deployment of cloud-side CL solutions [26, 36], such as saving storage through data quantization [29, 56], accelerating data loading via hierarchical memory management [39, 46], and speeding up computation by optimizing the allocation of hardware resources [35, 40].

Data Bottleneck on Mobile Devices. However, we identify that the scarce data resource on mobile devices is the key bottleneck for on-device CL. *First, data scarcity is a pervasive issue across various mobile applications.* For example, for image analysis applications, an average European citizen takes only 4.9 photos daily [10]. For virtual assistant applications, a mere 16% of iPhone users reports using Siri several times a day [84]. *Second, the utilization of data resources fundamentally determines the performance ceiling for on-device CL,* whereas the optimization of hardware resources only influences the efficiency with which this ceiling can be reached. On one hand, limited data resources for a single context often results in the well-known issue of model overfitting [25, 79]. On the other hand, the inadequate data resources for both past and new contexts exacerbate the mutual interference between their learning processes, which impedes knowledge transfer for new context and deteriorates the model performance on past contexts, a phenomenon commonly referred to as catastrophic forgetting [20, 34, 50].

Limitation of Existing Work. To tackle the challenge of data scarcity for CL, *few-shot CL* and *federated CL* are two representative approaches to mitigate the issues of overfitting and catastrophic forgetting from the aspects of model initialization and training algorithms (elaborated in §2.2). (1) *Few-shot CL* [49, 61, 69] involves pre-training ML models on common contexts with extensive data to capture general knowledge, which can be transferred to new contexts through model initialization and transfer learning techniques. However, this approach is ineffective for on-device settings due to the unpredictability and diversity of upcoming user contexts. (2) *Federated CL* [19, 80] suggests leveraging a cloud server to periodically aggregate the local models trained on distributed devices, which mitigates the overfitting problem on a single device and enables knowledge transfer across multiple devices. However, the model performance and convergence rate of federated CL are sensitive to device participation rate and data heterogeneity across devices [23, 41, 63], leading to high communication overhead and unstable training process for real-world applications.

Our Motivation. The data bottleneck of mobile devices coupled with the limitations of existing approaches motivate us to consider leveraging the abundant cloud-side data resources to enrich the sparse device-side data, fundamentally addressing the data scarcity problem. As we will elaborate in §2.2, simply increasing the training data size from 10 to 50 can yield a 10% improvement in model accuracy compared to

the best few-shot CL approach, while incurring less than 5% communication costs compared with federated CL. The feasibility of such a cloud-assisted data enrichment framework is underpinned by two key observations: (1) *Abundant cloud-side data resource.* Cloud servers typically possess extensive datasets sourced from various channels, such as public datasets released by organizations (e.g. ImageNet [59]), open-source data crawled from the Internet webs (e.g. Common Crawl [2]), crowdsourced data contributed by authorized mobile users (e.g. DonateClient service of Huawei [54] and learn from this app in Apple [6]). (2) *Similarities among user contexts and behaviors.* Previous investigations have demonstrated that the preferences and behaviors of different mobile users in various contexts share similar patterns rather than being entirely unique [7, 24, 45]. This indicates the existence of a cloud-side data-subset that exhibits a similar distribution with the device-side data, offering an opportunity to enhance on-device CL performance.

Challenges. A feasible data enrichment framework for practical on-device CL needs to be *private, effective* and *efficient*, which are challenging to be achieved simultaneously.

- *Privacy vs. Efficiency.* In contemporary mobile applications, user data stored on devices is subject to stringent privacy regulations like GDPR [53]. However, to enrich device-side data with an optimal data-subset from cloud, one must either upload raw user data to the cloud for precise similarity comparison [76], or download numerous data-subsets from cloud and conduct trial-and-error processes to identify the appropriate data-subset [14]. Therefore, achieving efficient data enrichment without violating user privacy is challenging.

- *Effectiveness vs. Efficiency for New Context.* Given the diverse sources of cloud-side data, a randomly selected data-subset is likely to deviate significantly from the device-side data distribution, thereby degrading the CL performance over personal contexts. However, to identify the data-subset with the highest data enrichment performance for on-device CL, the cloud server needs to evaluate an exponential number of candidate data-subsets from the vast cloud-side dataset, which introduces prohibitively high time complexity and computational burden. Consequently, simultaneously reaching high effectiveness and efficiency poses another challenge.

- *Effectiveness for Both Past and New Contexts.* As the data distributions of new contexts encountered by mobile users are dynamic, independently conducting data enrichment for each emerging context would compromise the on-device model's memory stability over past contexts, as the mutual interference among different contexts' learning processes can be escalated. Additionally, there is a lack of theoretical analysis or insight into the correlation between the enriched data of new context and model performance over past contexts, which further complicates the data enrichment problem for

CL. Therefore, designing a data enrichment strategy that is effective for both new and past contexts is challenging.

Our Design. We propose *DeI ta*, a cloud-assisted data enrichment framework designed for on-device CL with high privacy protection, efficiency and effectiveness. First, we provide a generic formalization of the data enrichment problem for on-device CL, and analyze its practical challenges concerning user privacy and computation efficiency. Second, to mitigate privacy concerns, we propose the construction of a compact “directory” dataset for cloud-side data. This approach helps to decompose the original data enrichment problem into two sub-problems, which can be independently solved by mobile device and cloud server without necessitating the exchange of sensitive raw data. Third, to achieve both efficient and effective data enrichment for each new context, we develop a soft data matching strategy to accurately solve the device-side sub-problem with sparse on-device data, and a theoretically optimal data sampling scheme for cloud-side data selection, which can be computed with a constant time complexity. Fourth, to maintain high effectiveness across both new and past contexts, we theoretically analyze the impact of new context’s enriched data on model performance over all contexts, and re-optimize cloud-side data sampling strategy from a holistic perspective.

Contributions of this work are summarized as follows:

- We identify the data bottleneck in on-device CL for dynamic user contexts, and explore the potential of utilizing cloud-side abundant data to enrich device-side data.
- We formalize the data enrichment problem for on-device CL and propose the first practical cloud-assisted data enrichment framework that simultaneously achieves privacy protection, effectiveness and efficiency.
- We evaluate *DeI ta* across four typical mobile computing tasks with diverse data modalities and models, demonstrating its broad applicability and superior performance over baselines in overall accuracy and communication efficiency.

2 Background and Motivation

2.1 On-Device Continual Learning

In mobile applications, users often encounter dynamic contexts and exhibit varying behaviors, leading to a non-stationary distribution of data collected by devices. For example, mobile users can encounter unseen objects, weather conditions and digital corruptions in image analytics applications [9, 33], experience new activities, physical conditions and device placements in human activity recognition (HAR) applications [37], or come across articles on various topics and in different languages in text analysis applications [32]. These applications necessitate timely and accurate responses from on-device ML models to ensure high service quality, driving




Figure 1: On-device continual learning pipeline.

the need for on-device CL. Figure 1 depicts the four stages a new context undergoes in on-device CL.

- *Context Detection*: When a new context is experienced by the user, it can be detected by mobile device through existing human-involved or automatic approaches [9, 33, 78]. For example, in HAR application, the former approach would suggest users to confirm a new activity, whereas the latter would detect a shift in sensor data distribution [78].
- *Data Collection*: For each new context, data samples following a new distribution are collected by mobile device as training data for the subsequent on-device CL process. In mobile applications, the data collected from an individual user’s daily life is sparse, personalized and private, such as photos taken by user or interactions with a virtual assistant.
- *Enhancement*: Prior to conducting on-device CL for a new context, various enhancement techniques need to be applied to mitigate the severe impact of data scarcity, such as few-shot CL based on model initialization and federated CL approaches based on training algorithms. *Our work focuses on the design of this stage from the data perspective.*
- *Continual Learning*: The training data of both new and past contexts are mixed to update the on-device model, which has been recognized as one of the most effective methods to assimilate knowledge from new contexts without forgetting the knowledge of past contexts¹ [11, 39, 44, 46, 55].


2.2 Limitation of Existing Approaches

In this section, we elaborate the limitations of existing few-shot CL and federated CL approaches in mitigating device-side data scarcity problem through preliminary experiments².

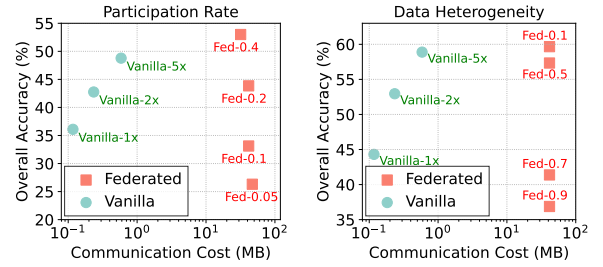
Few-shot CL [49, 61, 83] proposes pretraining ML models on base contexts with massive public data to capture general knowledge, which is then transferred to new contexts through transfer learning techniques. Representative methods include: 1) knowledge distillation [83] (FS-KD), which distills past contexts’ knowledge to the new context’s model by keeping the model outputs of historical data samples unchanged, 2) robust optimization [61] (FS-RO), which constrains model parameters within the common flat minima of

¹It is noteworthy that our data enrichment framework can also benefit other classic CL approaches, such as parameter regularization [34, 82] and context-adaptive model architectures [48, 60], as illustrated in §7.

²The detailed experimental settings are introduced in §6.1.



(a) Performance of few-shot CL approaches without (■) and with (□) prior information on user contexts, and performance of vanilla CL with different amount of available training data (Vanilla- $n\times$).



(b) Communication cost and accuracy of federated CL with varying device participation rates and data heterogeneity degree (Fed- p denotes that $p\times 100\%$ devices hold data from different contexts).

Figure 2: Preliminary experiments on image classification task to illustrate the limitations of existing solutions.

all contexts' training objective functions, and 3) parameter freezing [49] (FS-PF), which freezes the important parameters with high value of the previously trained model.

However, most of these few-shot CL approaches depend on a powerful model pre-training process, which pretrain either a large model on data from diverse contexts to fully capture the general knowledge, or a tiny model on a customized dataset to learn personalized knowledge. Unfortunately, both of them are impractical for on-device scenarios due to limited hardware resources and unpredictable contexts. On one hand, the limited memory and computational capabilities of mobile devices restrict the size and capacity of deployed models, impeding effective model pretraining over diverse data. On the other hand, the uncertainty of future user contexts prevents the pre-selection of a tailored data-subset for pretraining before model deployment. Our preliminary experiments shown in Figure 2(a) reveal that the performance of few-shot CL declines significantly without prior information on user contexts, with model accuracy reduction ranging from 8.6–15.3% for FS-PF, 1.9–7.9% for FS-RO and 3.9–7.2% for FS-KD. In contrast, simply increasing the training data size to 50 can outperform all few-shot CL approaches, underscoring the potential of data enrichment.

Federated CL [19, 80] utilizes a cloud server to periodically aggregate the parameters of models trained on distributed devices, thereby mitigating the overfitting issue on individual devices and facilitating knowledge transfer across multiple devices. However, the substantial communication overheads and unstable model training process render federated CL impractical for mobile devices. First, the frequent exchange of model parameters between mobile devices and the cloud server incurs significant communication costs and prolongs the wall-clock training time for on-device models. Second, the model performance of federated CL is relatively sensitive to the device participation rate (or amount) and the data heterogeneity across devices [41, 63]. Experimental results shown in Figure 2(b) indicate that: 1) Federated CL

achieves superior performance only when $\geq 20\%$ devices participate in each round of model aggregation or when more than $\geq 30\%$ mobile users experience similar contexts, which can be unrealistic in real-world settings; 2) In comparison to federated CL, transmitting data with a suitable distribution from cloud to each device could reach the same target accuracy with communication costs reduced to less than 1%.

3 Problem Definition

In this section, we present a generic formalization of the cloud-assisted data enrichment problem for on-device CL. We consider a scenario where a mobile user sequentially encounters T new contexts. Each context $t=1, \dots, T$ has an underlying data distribution \mathcal{D}_{de}^t and the device collects an empirical dataset $\widehat{\mathcal{D}}_{de}^t$ for training on-device model. Due to the scarcity of user data, a similar data-subset \mathcal{S}^t is expected to be retrieved from the cloud-side dataset \mathcal{D}_{cl} to enrich the on-device empirical dataset $\widehat{\mathcal{D}}_{de}^t$ and thereby enhance the CL performance.

To assess the effectiveness of data enrichment, we first define a metric to evaluate the similarity between two datasets in terms of their impacts on the model training process. In on-device CL, model parameters are typically fine-tuned by on-device data via gradient descent methods. Therefore, the similarity between two datasets \mathcal{D}_1 and \mathcal{D}_2 with respect to the training process of model θ can be quantified by the maximal difference between the average gradients of \mathcal{D}_1 and \mathcal{D}_2 within a nearby parameter space $\{\theta' \mid \|\theta' - \theta\| \leq \epsilon\}$:

$$Sim(\mathcal{D}_1, \mathcal{D}_2 | \theta) \triangleq - \max_{\|\theta' - \theta\| \leq \epsilon} \|\nabla L(\mathcal{D}_1, \theta') - \nabla L(\mathcal{D}_2, \theta')\|, \quad (1)$$

where $L(\mathcal{D}, \theta) = \mathbb{E}_{(x,y) \in \mathcal{D}} [l(x, y, \theta)]$ denotes the expected loss of model θ over dataset \mathcal{D} . A high similarity between two datasets \mathcal{D}_1 and \mathcal{D}_2 implies their comparable performance in updating model parameters for multiple steps, resulting in similar impacts on on-device model training.

Problem Formulation. For each new context t , the cloud server aims to select the most similar data-subset $\mathcal{S}^{t,*} \subseteq \mathcal{D}_{cl}$

to update the current on-device model θ^{t-1} in a similar way with the device-side underlying data distribution \mathcal{D}_{de}^t , which means that $\mathcal{S}^{t,*}$ and \mathcal{D}_{de}^t should exhibit high similarity as measured by the metric in Equation (1). Consequently, the data enrichment problem can be formally expressed as:

$$\begin{aligned} \mathcal{S}^{t,*} &= \arg \max_{\mathcal{S}^t \subseteq \mathcal{D}_{cl}, |\mathcal{S}^t| \leq B} \text{Sim}(\mathcal{S}^t, \mathcal{D}_{de}^t | \theta^{t-1}) \\ &\approx \arg \max_{\mathcal{S}^t \subseteq \mathcal{D}_{cl}, |\mathcal{S}^t| \leq B} \text{Sim}(\mathcal{S}^t, \widehat{\mathcal{D}}_{de}^t | \theta^{t-1}), \end{aligned} \quad (2)$$

where B represents the maximum allowable size of the selected data-subset and is constrained by the communication cost budget of each device. This formulation enables the device to enhance model training performance by expanding the training data from the collected dataset $\widehat{\mathcal{D}}_{de}^t$ to the enriched larger-scale dataset \mathcal{S}^t , while ensuring that the enriched data follows a similar distribution.

Practical Challenges. Directly solving the data enrichment problem in Equation (2) brings severe privacy concerns for mobile users and high computational burden for cloud server. First, the mobile device needs to upload both the current model θ^{t-1} and raw user data $\widehat{\mathcal{D}}_{de}^t$ to the cloud server, which poses a severe breach of user privacy. Second, the cloud server has to compute the similarity score $\text{Sim}(\mathcal{S}^t, \widehat{\mathcal{D}}_{de}^t | \theta^{t-1})$ for every possible data-subset $\mathcal{S}^t \subseteq \mathcal{D}_{cl}$, $|\mathcal{S}^t| \leq B$, resulting in exponential computation complexity.

4 Framework Design

Delta incorporates three key components to render data enrichment systematically practical: the construction of a directory dataset to address privacy concerns (§4.1), device-side soft data matching strategy coupled with a cloud-side data sampling scheme to efficiently and effectively enrich data for new contexts (§4.2), and a re-optimization of the cloud-side data sampling to further enhance its effectiveness across both past and new contexts (§4.3). *Each component is inspired and supported by theoretical analysis presented in §5 and the overall design rationale is illustrated in Figure 3.*

4.1 Directory Dataset Construction

To address privacy concerns, Delta introduces the concept of “directory” dataset, which facilitates decomposing the data enrichment problem (2) into two sub-problems, and allows the device and cloud to collaboratively solve the sub-problems without the need to share raw user data.

Design Rationale. Inspired by the directory structures in storage systems [16], Delta constructs a compact directory dataset consisting of a few data samples to represent the extensive cloud-side dataset, denoted as $\mathcal{D}_{cl}^{dir} = \{(\bar{x}_c, \bar{y}_c)\}_{c=1}^{|\mathcal{D}_{cl}^{dir}|}$. This directory dataset can be pre-downloaded by mobile devices along with the model deployment. As illustrated in

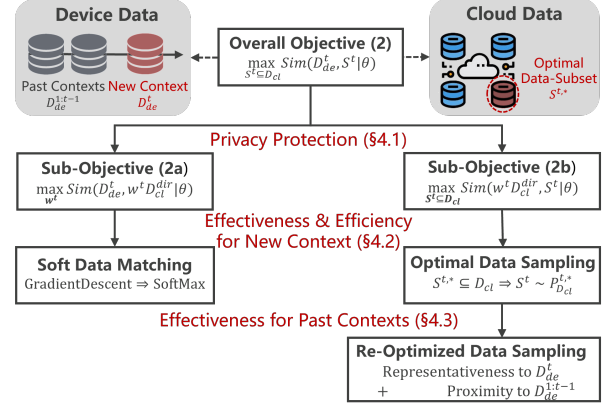


Figure 3: Design Rationale of Delta.

Figure 3 and supported in Theorem 1, the objective function (2) of the data enrichment problem can be decomposed into the sum of two sub-objective functions:

- *Sub-objective (2a)*: similarity between the device-side dataset \mathcal{D}_{de}^t and the weighted directory dataset $w^t \mathcal{D}_{cl}^{dir}$, where each data sample $(\bar{x}_c, \bar{y}_c) \in \mathcal{D}_{cl}^{dir}$ is assigned a weight w_c^t . The weight vector w^t is a variable to be optimized.
- *Sub-objective (2b)*: similarity between the weighted directory dataset $w^t \mathcal{D}_{cl}^{dir}$ and the cloud-side data-subset \mathcal{S}^t , where \mathcal{S}^t is the variable to be optimized.

These two two sub-objective functions can be optimized sequentially and independently by the mobile device and cloud server through the exchange of non-sensitive information:

- 1) *Mobile device* optimizes sub-objective (2a) by computing the optimal weight $w^{t,*}$ for the directory dataset \mathcal{D}_{cl}^{dir} to represent the device-side data distribution \mathcal{D}_{de}^t .
- 2) *Cloud server* optimizes sub-objective (2b) by searching for the optimal cloud-side data-subset $\mathcal{S}^{t,*} \subseteq \mathcal{D}_{cl}$ to align with the weighted directory dataset $w^{t,*} \mathcal{D}_{cl}^{dir}$, with $w^{t,*}$ being uploaded by the mobile device after device-side optimization.
- 3) *Device-cloud communication* involves the cloud-side directory dataset \mathcal{D}_{cl}^{dir} and the device-side optimized weight $w^{t,*}$, which do not involve any raw user data and thus protect user privacy akin to classic federated learning [51]. Detailed discussion and comparison are presented in §8.

Practical Implementation. The practical effectiveness of the above decomposition process relies on an appropriate directory dataset that accurately represents the cloud-side public dataset. While classical data clustering methods can be used to select cluster centroids as the directory dataset elements, we observe that directly clustering raw data samples may not fully capture the influence of data on model training, due to the diverse sources, wide-ranging distributions and varying dimensions of cloud-side data. To address this issue, we take advantage of the typical paradigm of on-device model training [12, 13, 17], where the feature extractor ϕ is pre-trained on extensive cloud-side data for generalization

ability and the classifier ψ is trained on device-side data for personalization performance. We propose clustering data samples $(x, y) \in \mathcal{D}_{cl}$ based on the feature extractor outputs $\phi(x)$ rather than raw input x , and selecting the cluster centroids as elements of the directory dataset, which offers two advantages: 1) features as model’s intermediate outputs have a consistent dimension and are more relevant to model training than raw inputs, 2) the features of most cloud-side data samples are already available from the pre-training process of feature extractor, incurring minimal additional costs.

4.2 Data Enrichment for New Context

While the directory dataset safeguards user privacy by decomposing the data enrichment problem into device-side and cloud-side sub-problems, it is non-trivial to solve them in an effective and efficient manner due to the scarcity of on-device data and diversity of cloud-side data.

- *Device-side ineffectiveness*: Solving sub-problem (2a) requires determining the optimal weight $w^{t,*}$ to align the weighted directory dataset $w^t \mathcal{D}_{cl}^{dir}$ with the device-side data distribution \mathcal{D}_{de}^t . However, the underlying data distribution is typically approximated by the sparse empirical dataset $\widehat{\mathcal{D}}_{de}^t$ stored by mobile device, which can cause conventional gradient descent algorithms to converge to local optima. Consequently, the derived weight becomes overfitted to the limited empirical dataset and ineffective in representing the device-side data distribution.
- *Cloud-side inefficiency*: Exactly solving sub-problem (2b) involves evaluating the similarity score for each potential cloud-side data subset, which requires exploring a vast feasible region of candidate data-subsets $\mathcal{S}^t \subseteq \mathcal{D}_{cl}$, $|\mathcal{S}^t| \leq B$ and results in exponential computation and time complexity for cloud server, leading to low efficiency.

To achieve an efficient and effective data enrichment process for each coming context, we propose a soft data matching strategy for mobile device to derive a representative directory weight by fully leveraging the limited on-device data, and a data sampling scheme for cloud server to sample an optimal data-subset with constant time complexity.

Device-Side: Soft Data Matching. To prevent the directory weight w^t from overfitting to scarce on-device data, we propose to assign physical meanings to w^t by interpreting each element w_c^t as the fraction of on-device data that exhibits high similarity with the cloud-side cluster centroid $(\bar{x}_c, \bar{y}_c) \in \mathcal{D}_{cl}^{dir}$. Thus, for each data sample $(x, y) \in \widehat{\mathcal{D}}_{de}^t$ collected by mobile device, its similarities with all the cluster centroids are computed, and the weight of the most similar one is incremented by one step:

$$c^* = \arg \max_c \text{Sim}((x, y), (\bar{x}_c, \bar{y}_c) \mid \theta^{t-1}),$$

$$w_{c^*}^t \leftarrow w_{c^*}^t + 1. \quad (\text{Hard Matching})$$

However, in our experiments, we observe that each on-device data sample can exhibit high similarity with more than one cloud-side cluster centroids, which is influenced by the granularity of cloud-side data clustering (*i.e.* the number of data clusters) during the directory construction process. However, the “hard” matching function *argmax* is incapable of capturing the correlation between one device-side sample and multiple cloud-side clusters. Thus, we propose to employ a “soft” matching function *softmax*, allowing each data sample to contribute to the weights of more than one clusters:

$$\forall c, w_c^t \leftarrow w_c^t + \text{Softmax} \left(\frac{\text{Sim}((x, y), (\bar{x}_c, \bar{y}_c) \mid \theta^{t-1})}{\tau} \right), \quad (3)$$

where τ is a temperature hyperparameter to control the weight increments of clusters with different degrees of similarity. As $\tau \rightarrow 0$, *softmax* gradually degrades to *argmax*.

Cloud-Side: Optimal Data Sampling. To enhance efficiency and reduce the computational overhead on the cloud server, we propose transforming the “hard” data selection process into a “soft” data sampling process. The key difference is that the former seeks to find an exact data-subset $\mathcal{S}^{t,*}$ to optimize sub-problem (2b), whereas the latter aims to compute a data sampling policy $P_{\mathcal{D}_{cl}}^{t,*}$ such that the sampled data-subset is optimal for sub-problem (2b) in expectation:

$$\max_{\mathcal{S}^t \subseteq \mathcal{D}_{cl}} \text{Sim}(\mathcal{S}^t, w^t \mathcal{D}_{cl}^{dir} \mid \theta^{t-1}) \quad (\text{Hard Selection})$$

$$\Rightarrow \max_{P_{\mathcal{D}_{cl}}^{t,*}} \mathbb{E}_{\mathcal{S}^t \sim P_{\mathcal{D}_{cl}}^{t,*}} [\text{Sim}(\mathcal{S}^t, w^t \mathcal{D}_{cl}^{dir} \mid \theta^{t-1})]. \quad (\text{Soft Sampling})$$

This transformation allows the cloud server to directly identify an appropriate data-subset through data sampling policy, which can be computed with constant time complexity.

We outline the specific operations of cloud-side data sampling scheme, with theoretical foundation provided in §5.2. The scheme involves *inter-cluster size allocation* and *intra-cluster data sampling*, which determine *how many* and *which* data samples to select from each cloud-side data cluster:

- *Inter-cluster size allocation.* Given that the size of the selected data-subset is limited by the communication cost budget, the cloud server needs to allocate distinct sampling sizes to different data clusters to maximize the overall similarity between the sampled data-subset and the weighted directory dataset, *i.e.* sub-objective (2b). As demonstrated in Lemma 1, the optimal sampling size $|\mathcal{S}_c^{t,*}|$ for each cluster $\mathcal{D}_{cl,c}$ depends on its directory weight w_c^t and the dispersion degree of intra-cluster feature distribution $\mathbb{E}_x \|\phi(x) - \phi(\bar{x})\|$:

$$|\mathcal{S}_c^{t,*}| \propto w_c^t \cdot \mathbb{E}_{(x,y) \in \mathcal{D}_{cl,c}} \|\phi(x) - \phi(\bar{x}_c)\|. \quad (4)$$

For each cluster, a higher weight suggests a higher similarity with the device-side data for on-device CL, and a wider feature distribution indicates the need for more data samples to comprehensively represent the cluster.

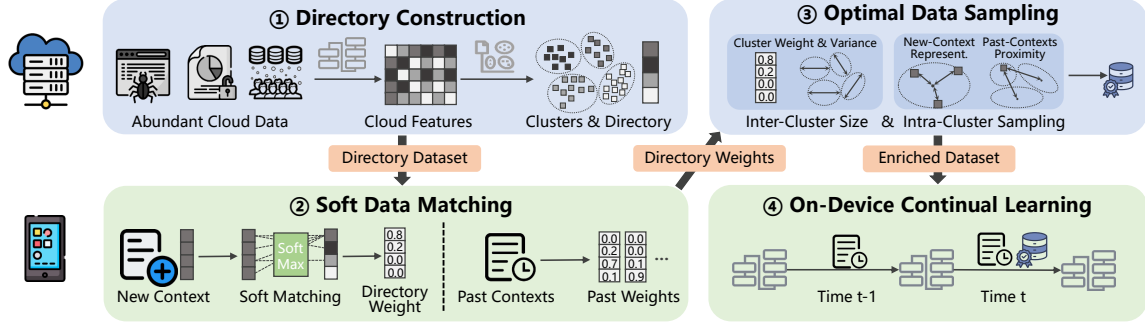


Figure 4: Overall Workflow of Delta Framework. Delta serves as a plug-in for on-device continual learning.

• *Intra-Cluster Data Sampling.* Within each cloud-side data cluster $\mathcal{D}_{cl,c}$, the optimal sampling probability for each data sample (x, y) is proportional to the feature distance between such data sample and the cluster centroid (\bar{x}_c, \bar{y}_c) :

$$P_{\mathcal{D}_{cl,c}}^{t,*}(x, y) = \frac{\|\phi(x) - \phi(\bar{x}_c)\|}{\sum_{(x', y') \in \mathcal{D}_{cl,c}} \|\phi(x') - \phi(\bar{x}_c)\|}. \quad (5)$$

Theoretically, our analysis in Lemma 1 demonstrates that this sampling probability could maximize the expected similarity between each data cluster $\mathcal{D}_{cl,c}$ and the corresponding selected data-subset \mathcal{S}_c^t , thereby optimizing sub-objective (2b) in expectation given fixed directory weights w^t . Intuitively, this sampling strategy favors data samples that are farther from the cluster centroid, which enhances the diversity and informativeness of the selected data-subset while ensuring unbiasedness and representativeness through data re-weighting technique like importance sampling [31].

4.3 Data Enrichment for All Contexts

Although the previous components ensure a private, efficient and effective data enrichment process for each new context, the notorious issue of catastrophic forgetting (*i.e.* inferior memory stability) is also exacerbated. First, as model parameters θ continually adapt to the enriched data $\{\mathcal{S}^i\}_{i=1}^t$, the similarity between each past context i 's enriched data \mathcal{S}^i and the underlying distribution \mathcal{D}_{de}^i gradually diminishes, hindering the use of $\{\mathcal{S}^i\}_{i=1}^t$ for retaining past knowledge. Second, independently enriching data solely for the new context will exacerbate the mutual interference between the model training processes of new and past contexts.

To address these issues, we take the first step to theoretically analyze the correlation between new context's enriched data and the model performance on both new and past contexts. Further, we re-optimize the data sampling scheme for cloud server to identify a data-subset that could contribute to the learning processes of both new and past contexts.

Theoretical Analysis. Theorem 3 reveals that the overall CL performance, quantified by the average loss of model over all contexts, is primarily determined by three terms:

- 1) *New-context representativeness*, which is quantified by the feature distance between the enriched dataset \mathcal{S}^t and the underlying data distribution of new context \mathcal{D}_{de}^t .
 - 2) *Past-contexts proximity*, which is measured by the feature distance between the enriched dataset \mathcal{S}^t and the underlying data distributions of all the past contexts $\{\mathcal{D}_{de}^i\}_{i=1}^{t-1}$.
 - 3) *Cross-Context Heterogeneity*, which is a fixed term and determined by the heterogeneity between the new context and the past contexts encountered by the mobile user.
- Consequently, the original intra-cluster data sampling strategy in Equation (5) can be seen as focusing only on the first term (*i.e.* effectiveness for new context), while overlooking the second term (*i.e.* effectiveness for past contexts).

Practical Implementation. Guided by the theoretical results, we further derive the analytical expression for the re-optimized cloud-side data sampling policy, with the detailed mathematical derivation provided in §5.3. Specifically, for intra-cluster data sampling, the optimal sampling probability for each data sample is proportional to the weighted sum of new-context representativeness and past-contexts proximity:

$$P_{\mathcal{D}_{cl,c}}^{t,*}(x, y) \propto \left\| \phi(x) - \phi(\bar{x}_c) \right\| + \alpha \left\| \phi(x) - \frac{\sum_{i=1}^{t-1} \phi(w^{i,*} \mathcal{D}_{cloud}^{dir})}{t-1} \right\|,$$

where α is a hyperparameter determined by the device to balance the model performance over the new context and past contexts when conducting cloud-assisted data enrichment.

4.4 Overall Framework

We illustrate the overall workflow of Delta framework in Figure 4, which comprises four stages. **1 Directory Construction:** Initially, the cloud server utilizes the pre-trained feature extractor to extract features from diverse datasets and performs data clustering to construct the directory dataset. The directory dataset is then distributed to mobile devices along with the model deployment. **2 Soft Data Matching:** For each coming new context t , the mobile device solves sub-problem (2a) through the soft data matching strategy outlined in Equation (3), and uploads the optimal directory weights for both the new and past contexts to the cloud

server. **⑤ Optimal Data Sampling:** Upon receiving the directory weights, the cloud server computes the analytical expressions for the optimal data sampling scheme, which includes inter-cluster size allocation in Equation (4) and intra-cluster data sampling in Equation (5). The optimal data-subset is then sampled according to the scheme and transmitted back to the mobile device. **④ On-Device Continual Learning:** The mobile device conducts CL process using the enriched datasets of both new and past contexts. *Generally, Delta serves as a plug-in module to enhance on-device CL performance with privacy protection, effectiveness and efficiency.*

5 Theoretical Analysis

In this section, we provide theoretical foundations for the key components of Delta framework. Detailed proofs are provided in supplementary material [22] due to limited space.

5.1 Theory for Directory Construction

To facilitate data enrichment as outlined in Equation (2) without disclosing raw user data, we introduce the directory dataset to decompose the original objective function into two sub-objective functions. The performance of this decomposition is theoretically guaranteed by Theorem 1, which elucidates the relation between the original objective function and two sub-objective functions.

THEOREM 1. *Given directory dataset \mathcal{D}_{cl}^{dir} , the maximal similarity between the device-side dataset \mathcal{D}_{de}^t and the cloud-side data-subset $\mathcal{S}^t \subseteq \mathcal{D}_{cl}$ for model θ^{t-1} can be bounded by*

$$\underbrace{\max_{\mathcal{S}^t \subseteq \mathcal{D}_{cl}} \text{Sim}(\mathcal{D}_{de}^t, \mathcal{S}^t | \theta^{t-1})}_{\text{original objective}} \geq \underbrace{\max_{\mathcal{w}^t} \text{Sim}(\mathcal{D}_{de}^t, \mathcal{w}^t \mathcal{D}_{cl}^{dir} | \theta^{t-1})}_{\text{sub-objective (2a) for optimal weight}} + \underbrace{\max_{\mathcal{S}^t \subseteq \mathcal{D}_{cl}} \text{Sim}(\mathcal{S}^t, \mathcal{w}^{t,*} \mathcal{D}_{cl}^{dir} | \theta^{t-1})}_{\text{sub-objective (2b) for optimal subset}},$$

where $\mathcal{w}^t \mathcal{D}_{cl}^{dir}$ represents the weighted directory dataset.

This theorem shows that the optimal value of the original objective function (2) is bounded from below by the sum of the optimal values of the two sub-objective functions (2a) and (2b). Consequently, Delta essentially optimizes the worst-case performance of data enrichment for diverse contexts. The practical gap between the original and decomposed objective functions is determined by the representativeness of the cloud-side directory dataset. In §6.3, we empirically show that a directory dataset with around 10^2 elements is sufficient to represent a cloud-side dataset consisting of 10^6 data samples across 10^2 contexts.

5.2 Theory for New Context's Enrichment

To provide theoretical guarantees for the optimality of the cloud-side data sampling scheme outlined in Equations (4)

and (5), we first present Theorem 2 to partition the on-device model θ into a feature extractor ϕ and a classifier ψ with a Lipschitz continuity constant L_ψ . The feature extractor is typically pre-trained by cloud server and remains unchanged during the on-device model training process.

THEOREM 2. *The expected similarity between the weighted directory dataset $\mathcal{w}^t \mathcal{D}_{cl}^{dir}$ and the data-subset \mathcal{S}^t selected according to sampling scheme $P_{\mathcal{D}_{cl}}^t$ is bounded by:*

$$\begin{aligned} & \mathbb{E}_{\mathcal{S}^t \sim P_{\mathcal{D}_{cl}}^t} [\text{Sim}(\mathcal{S}^t, \mathcal{w}^t \mathcal{D}_{cl}^{dir} | \theta^{t-1})] \\ & \geq -\mathbb{E}_{\mathcal{S}^t \sim P_{\mathcal{D}_{cl}}^t} L_\psi \left\| \mathbb{E}_{(x,y) \in \mathcal{S}^t} [\phi(x)] - \sum_c \mathcal{w}_c^t \phi(\bar{x}_c) \right\|. \end{aligned}$$

Further, in Lemma 1, we demonstrate that the expected value of sub-objective function (2b) (i.e., the lower bound of the above inequality) is determined by two terms: inter-cluster sampling size $|\mathcal{S}_c^t|$ and intra-cluster sampling probability $P_{\mathcal{D}_{cl,c}}^t(x, y)$ for each cloud-side data cluster c .

LEMMA 1. *The expected similarity between the sampled data-subset \mathcal{S}^t and the weighted directory dataset $\mathcal{w}^t \mathcal{D}_{cl}^{dir}$ is determined by each cluster c 's sampling size $|\mathcal{S}_c^t|$ and intra-cluster data sampling probability $P_{\mathcal{D}_{cl,c}}^t(x, y)$:*

$$\begin{aligned} & \min_{P_{\mathcal{D}_{cl}}^t} \mathbb{E}_{\mathcal{S}^t \sim P_{\mathcal{D}_{cl}}^t} \left\| \mathbb{E}_{(x,y) \in \mathcal{S}^t} [\phi(x)] - \sum_c \mathcal{w}_c^t \phi(\bar{x}_c) \right\| \\ & = \min_{|\mathcal{S}_c^t|, P_{\mathcal{D}_{cl,c}}^t} \sum_c \left(\frac{(\mathcal{w}_c^t)^2}{|\mathcal{S}_c^t|} \cdot \sum_{(x,y) \in \mathcal{D}_{cl,c}} \frac{\|\phi(x) - \phi(\bar{x}_c)\|^2}{|\mathcal{D}_{cl,c}|^2 \cdot P_{\mathcal{D}_{cl,c}}^t(x, y)} \right) \end{aligned}$$

By leveraging Cauchy-Schwarz inequality, we can derive the analytical expressions of the optimal data sampling policy (i.e. $|\mathcal{S}_c^{t,*}|$ and $P_{\mathcal{D}_{cl,c}}^{t,*}$), which can be computed directly using the directory weights uploaded by mobile device:

$$\begin{cases} |\mathcal{S}_c^{t,*}| & \propto \mathcal{w}_c^t \cdot \mathbb{E}_{(x,y) \in \mathcal{D}_{cl,c}} \|\phi(x) - \phi(\bar{x}_c)\| \\ P_{\mathcal{D}_{cl,c}}^{t,*}(x, y) & \propto \|\phi(x) - \phi(\bar{x}_c)\|, \quad \forall (x, y) \in \mathcal{D}_{cl,c}. \end{cases} \quad (6)$$

5.3 Theory for All Contexts' Enrichment

In §4.3, we propose to refine the cloud-side data sampling scheme to ensure that the enriched data for new context can contribute to the learning processes of both new and past contexts. To achieve this, we first analyze the impact of new context's enriched data on the model performance over all contexts in Theorem 3, which consists of three key terms: representativeness to new context, proximity to past contexts and the data heterogeneity across different contexts.

THEOREM 3. *In m -th training round for context t , when the model parameters are updated from $\theta^{t,m}$ to $\theta^{t,m+1}$ using the enriched data \mathcal{S}^t sampled by policy $P_{\mathcal{D}_{cl}}^t$, the expected reduction in model loss (or improvement in model performance) over all contexts' data distribution $\mathcal{D}_{de}^{1:t}$ can be bounded by:*

$$\begin{aligned}
& \mathbb{E}_{S^t \sim P_{\mathcal{D}_{cl}}^t} \left[\underbrace{L(\mathcal{D}_{de}^{1:t}, \theta^{t,m+1}) - L(\mathcal{D}_{de}^{1:t}, \theta^{t,m})}_{\text{loss reduction in } m\text{-th model update}} \right] \\
& \leq \frac{1}{2} (H\eta^2 - \eta) L_{\psi} \underbrace{\mathbb{V}_{S^t \sim P_{\mathcal{D}_{cl}}^t} [\phi(\mathcal{D}_{de}^t) - \phi(S^t)]}_{\text{representativeness to new context } t} + \\
& \quad \underbrace{\frac{\eta L_{\psi}}{2} \mathbb{V}_{S^t \sim P_{\mathcal{D}_{cl}}^t} [\phi(\mathcal{D}_{de}^{1:t-1}) - \phi(S^t)]}_{\text{proximity to past contexts } 1 \sim t-1} + \underbrace{\frac{\eta L_{\psi}}{2} \|\phi(\mathcal{D}_{de}^t) - \phi(\mathcal{D}_{de}^{1:t-1})\|^2}_{\text{heterogeneity across contexts}}
\end{aligned}$$

where $\mathbb{V}_x[f(x)]$ denotes the variance of function $f(x)$.

Building on this analysis, we observe that to improve the overall CL performance and reduce the model loss across all contexts, the cloud-side sampling scheme $P_{\mathcal{D}_{cl}}^t$ should take both the representativeness to new context and the proximity to past contexts into consideration. From a theoretical perspective, we further derive the analytical expression of the re-optimized data sampling scheme $P_{\mathcal{D}_{cl}}^{t,*}$ in Lemma 2.

LEMMA 2. *To optimize the model performance on the overall data distribution of all encountered contexts, the intra-cluster data sampling probability $P_{\mathcal{D}_{cl}}^{t,*}$ needs to be refined as:*

$$P_{\mathcal{D}_{cl}}^{t,*}(x, y) \propto \sqrt{\|\phi(x) - \phi(\bar{x}_c)\|^2 + \alpha \|\phi(x) - \phi(\mathcal{D}_{de}^{1:t-1})\|^2},$$

where $\alpha = \frac{1}{L_{\psi}\eta-1}$ can be regarded as a hyper-parameter to balance the model performance over new and past contexts.

6 Evaluation

6.1 Experimental Setup

Tasks, Datasets and Models. To demonstrate Delta's broad applicability, we evaluate Delta on four typical mobile computing tasks with diverse data modalities, model structures and categories of user contexts (summarized in Table 1).

- **Image Classification (IC).** The Cifar10-C dataset [28] contains around 750,000 images of 10 objects across four context categories: weather, noise, blur and digital corruptions. For each context category, the dataset is processed into 5 subsets with 2 new objects and 1 new context per subset. ResNet-18 [27] is trained for this 10-class image classification task.
- **Human Activity Recognition (HAR).** HHAR [68], UCI [58], MotionSense [47] and Shoaib [64] are four public datasets collected from 73 users performing 6 basic activities (still, walking, upstairs, downstairs, jogging, bike) with 5 device placements (pocket, belt, arm, wrist, waist). For each context category, the dataset is processed into 6 subsets with 1 new activity in a new context. A lightweight CNN-based model DCNN [77] is trained for this 6-class classification task.
- **Audio Recognition (AR).** Google Speech command [73] comprises 100,000 sound files of 20 commands from over 2,000 users with varied tones and environmental conditions. The dataset is processed into 5 subsets for each context category,

each containing 4 new commands in 1 new context. A deep neural network VGG-11 [67] is deployed for this task.

- **Text Classification (TC).** The NC corpus in XGLUE benchmark [43] is a cross-lingual understanding dataset consisting of 50,000 articles on 10 topics and in 5 languages (German, English, Spanish, French, Russian). For each context category, the dataset is processed into 5 subsets with 2 new topics and 1 new context. A transformer-based model BERT [18] is fine-tuned for this 10-class classification task.

Note that we standardize the total number of on-device contexts to approximately 5 to ensure a consistent evaluation of Delta across various tasks, models and modalities, and thus the class number per context may vary for different datasets.

Configurations. For each task, we collect data from 50% users (or randomly select 50% samples for IC and TC tasks) to form the cloud-side public dataset, with the remaining data used to simulate the on-device empirical data across different contexts. For cloud server, data samples from different users and contexts are mixed to reflect the typical scenario where the specific context of each raw data sample is unknown. For mobile device, we use 5 samples per class in each context as empirical data for model fine-tuning, consistent with the statistics that an average European citizen takes over around 4.9 photos daily [10] and uses Siri several times a day [84]. The remaining data samples are used as testing data for each context. For Delta, the temperature τ for device-side soft matching is set to 0.1 and the number of cloud-side data clusters is $20 \times \text{num_class}$ (i.e. 200/120/400/200 for IC/HAR/AR/TC). The hyperparameter α is set to 1.0 to balance the effects of cloud-side data sampling on new and past contexts. The default communication budget is set to 25 samples/class for each new context, and an in-depth analysis of the impacts of such budget and on-device data amount is presented in §6.3.

Baselines. To our best knowledge, Delta is the first data enrichment framework for on-device CL, and we compare it against the model- and algorithm-based baselines (few-shot CL and federated CL) and a random data enrichment baseline. 1) *Few-shot CL* pre-trains model on cloud-side data in advance to capture the general knowledge, which is transferred to device-side new contexts through knowledge distillation [83] (FS-KD), robust optimization [61] (FS-RO) and parameter freezing [49] (FS-PR). 2) *Federated CL* leverages the cloud server to periodically aggregate the models trained on multiple devices per 10 local model updates. In our experiments, the default device number is 50, except for 35 for HAR task. We use Fed- p to denote different settings of device participation rate p . For IC and TC tasks, the data samples from each user are from the same context category to simulate the common data heterogeneity across users (e.g. W/N/B/D for IC and L for TC). The CL performance is evaluated on an independent test dataset, constructed according

Task	Modality	Context Category	Dataset	Model(#params)
IC	Image	Object (O), Weather (W), Noise (N), Blur (B), Digital Corruption (D)	Cifar10-C	ResNet18(11.2M)
HAR	IMU	Activity (A), Physical Condition (P), Device Placement (D)	HHAR, UCI, Motion, Shoaib	DCNN(17.3K)
AR	Audio	User Command (C), Tone (T), Environmental Noise (N)	Google Speech	VGG11(9.75M)
TC	Text	Article Topic (T), Language (L)	XGLUE	BERT(0.178B)

Table 1: Summary of tasks, modalities, contexts, datasets and models.

to the user contexts specified by the experimental setting. 3) *Random* method selects a random cloud-side data-subset to enrich device-side empirical data.

Metrics. We assess the on-device CL performance using four metrics. *Overall performance* measures the inference accuracy of the final model across all the encountered contexts. *Learning plasticity* is the average of each new context’s highest accuracy during its learning process. *Memory stability* is the average ratio between each context’s final accuracy to its maximal accuracy. *System overheads* include the computation latency, communication costs, memory footprint and energy consumption for both the device side and cloud side.

Deployments. We use a cloud server with one NVIDIA 3090Ti GPU and one mobile platform NVIDIA Jetson Nano [52].

6.2 End-to-End Performance

We begin by comparing the end-to-end performance of *Delta* against the baselines across all four tasks.

***Delta* significantly improves the overall performance of on-device CL.** Table 2 summarizes the average accuracy of the final model across all contexts. Compared with the best-performing few-shot CL method, *Delta* achieves a notable improvement, with accuracy increases of 13–16% higher accuracy on IC, 10–14% on HAR, 0.2–2.5% on AR, and 4–7.3% on TC. Note that *Delta*’s improvement on AR task is minimal because its data heterogeneity across contexts is relatively low (i.e. different tones and background noises) and vanilla CL could perform well. When compared to federated CL, *Delta* consistently achieves the highest overall performance across all settings, and reduces total communication costs by 91–99%, demonstrating its superior effectiveness and efficiency in enhancing CL performance. Furthermore, we observe that for most tasks (IC, HAR and TC), all methods tend to perform better on contexts with mixed categories (last line of each task in Table 2). The potential reason is that data samples with different context categories exhibit a greater distribution divergence, making it easier for the on-device model to learn the decision boundary.

***Delta* enhances the learning plasticity of on-device CL with various new contexts.** Figure 5 reports the average value of each new context’s peak accuracy during the learning process, a metric widely adopted to assess the learning plasticity. A key observation is that *Delta* consistently

outperforms the baselines across various tasks, data modalities and context categories, demonstrating high robustness and applicability to diverse new contexts. For example, *Delta* achieves around 90% and 100% accuracy for new context in IC and HAR tasks regardless of context categories and fluctuates less than 3% accuracy on the other two tasks. The high accuracy for new contexts can be attributed to the limited classes within each new context and the enriched data from cloud side. In contrast, the performance of baselines on new contexts is sensitive to the diversity of context categories, such as few-shot CL dropping from 95% to 90% on AR task and federated CL reducing from 93% to 87% on TC task. The rationale behind these is that few-shot CL depends on the high relevance between the on-device context and the base contexts during pre-training to facilitate effective knowledge transfer. Similarly, the performance of federated CL is largely influenced by the data heterogeneity across different users’ ongoing contexts. Conversely, *Delta* can consistently identify an appropriate cloud-side data-subset that contributes to the device-side CL process, making it relatively robust.

***Delta* consistently achieves a low accuracy drop on past contexts and exhibits high memory stability.** Figure 6 plots the average ratio between each context’s final accuracy and its peak accuracy, which indicates that *Delta* can maintain over 90% relative performance for past contexts. The superior memory stability is due to the consideration of the impact of new context’s enriched on all contexts’ overall performance during the cloud-side data sampling process. We also observe that few-shot CL methods can slightly outperform *Delta* in some cases. This is because they achieve significantly lower peak accuracy for new contexts compared to *Delta* (e.g. a 10% accuracy gap in IC task shown in Figure 5), and thus the accuracy drop might be less pronounced.

***Delta* incurs marginal system overheads for both mobile device and cloud server,** as depicted in Figure 7.

- *Device-Side.* The soft matching solution for sub-problem (2a) requires computing the feature of each local data sample and its distance to each element of directory dataset. This results in additional latency of 23.8/1.05/ 4.25/109ms and energy consumption of 0.49/0.30/0.42/2.47J per sample for IC/HAR/AR/TC tasks, respectively. Moreover, Figure (7(c)) shows that soft matching process has a lower memory footprint than CL process for avoiding model backpropagation,

Tasks	Context Category	Vanilla CL	Few-Shot CL			Federated CL			Data Enrichment		Δ Acc.	Δ Comm.
			FS-KD	FS-RO	FS-PF	Fed-0.1	Fed-0.2	Fed-0.4	Random	Delta		
IC	O+W	32.7±1.49	41.7±1.78	39.2±2.13	36.9±2.87	31.8±0.24	46.4±1.65	55.1±0.42	42.5±2.42	57.7±0.54	16.0% ↑	93.7% ↓
	O+N	31.3±1.74	36.2±2.34	35.5±1.65	32.3±1.25	31.1±0.04	40.4±0.51	45.0±0.12	35.8±1.00	50.9±1.66	14.8% ↑	93.5% ↓
	O+B	35.6±0.94	43.7±1.12	40.6±0.24	39.2±0.06	32.6±0.16	39.6±0.24	50.1±0.31	39.9±1.69	57.7±0.98	14.0% ↑	91.1% ↓
	O+D	45.0±2.57	55.1±1.17	51.5±2.66	52.2±3.10	36.9±0.04	49.0±0.51	61.7±0.34	53.7±2.24	72.3±2.27	17.1% ↑	92.2% ↓
	O+W+N+B+D	77.3±0.49	81.2±1.53	80.4±0.81	75.3±0.41	30.0±0.05	39.8±0.71	50.8±0.41	47.8±6.64	94.8±2.74	13.6% ↑	95.3% ↓
HAR	A	52.4±3.67	55.0±3.93	52.9±2.55	48.3±2.69	54.0±0.64	60.0±0.21	61.3±0.55	58.4±0.35	69.3±1.96	14.3% ↑	99.6% ↓
	A+P	51.2±4.53	53.3±3.20	50.1±3.52	49.4±2.95	60.5±1.28	61.1±1.89	63.1±0.85	58.5±0.75	66.6±1.78	13.3% ↑	99.8% ↓
	A+P+D	81.0±4.75	80.3±2.35	78.7±4.37	71.0±4.27	62.2±3.58	66.8±3.97	70.1±4.28	61.1±3.25	90.3±5.09	10.0% ↑	99.7% ↓
AR	C	93.6±0.16	93.5±0.07	92.9±0.65	94.2±0.28	88.1±1.65	88.3±0.83	88.5±1.78	90.4±0.19	94.3±0.17	0.2% ↑	99.9% ↓
	C+T	89.0±0.41	89.4±0.57	89.4±0.38	90.3±0.79	86.5±0.24	88.5±0.62	88.7±0.25	90.3±0.26	91.1±1.17	0.8% ↑	99.9% ↓
	C+T+N	84.7±0.64	84.8±1.52	86.2±0.79	86.9±0.40	87.5±0.54	87.7±0.31	88.0±0.61	88.5±1.45	89.2±1.60	2.3% ↑	99.9% ↓
TC	T	73.2±2.15	73.5±1.35	75.7±4.07	73.3±2.56	79.6±0.37	79.6±0.19	79.8±0.14	73.9±2.69	83.1±2.26	7.3% ↑	99.8% ↓
	T+L	77.7±3.19	82.2±0.29	80.1±3.02	80.0±1.89	84.3±0.14	84.4±0.18	84.7±0.09	79.7±2.21	86.2±2.16	4.0% ↑	99.4% ↓

Table 2: Summary of overall CL performance (average accuracy of final model on all contexts). We also mark Delta’s improvement on accuracy (over few-shot CL) and reduction in communication costs (over federated CL).

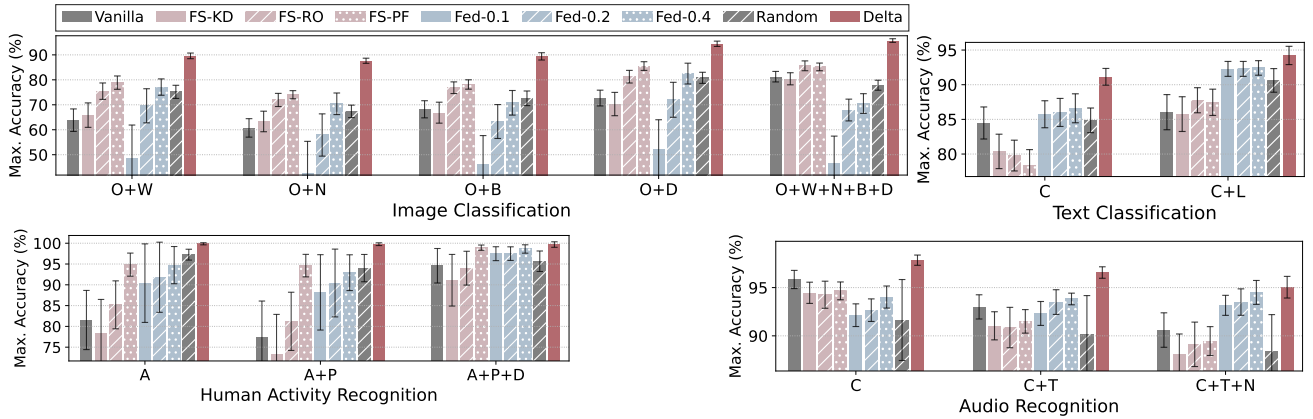


Figure 5: Comparison of learning plasticity (maximum model accuracy for each new context during CL).




Figure 6: Comparison of memory stability in the settings of mixed context categories for each task.

indicating that Delta does not increase peak memory usage due to the sequential execution of Delta and on-device CL.

- *Cloud-Side*. The analytical solution for optimal cloud-side data sampling can be computed within 2.56–7.15 ms using a single 10-core Intel CPU with a memory footprint of 0.12–7.8 MB. This high computational efficiency allows for parallel cloud-side operations for numerous devices simultaneously.
- *Device-cloud Communication*. For each context, the communication overhead includes the uploading of device-side directory weight, which consists of only several vectors (≤ 1 KB), and the downloading of cloud-side enriched data, which requires a total of 30.4/2.89/23.5/6.43 KB for IC/HAR/ AR/TC tasks under default settings.

6.3 Component-Wise Analysis

We further delve into the functionality and sensitivity of each key component within Delta framework.

Device-Side Data Soft Matching. To illustrate the importance of soft matching strategy, we assess the performance of Delta using various strategies to address sub-objective (2a), including gradient descent (*GD*), hard matching (*argmax*) and soft matching (*softmax*) with varying temperatures τ . Figure 8(a) indicates that *GD* underperforms across most tasks, while *softmax* consistently outperforms *argmax*. The reasons are twofold: 1) *GD* is susceptible to getting trapped in local optima and leads to the overfitted directory weight; 2) *argmax* fails to exploit the similarities between one device-side sample and multiple cloud-side clusters, which is essential when the cloud-side data is finely clustered. We also note that the optimal τ differs by task due to varying feature distributions, and we set $\tau = 1.0$ for stable performance.

Device-Side Data Size. Figure 8(b) shows the impact of user data amount on Delta and baseline performances, where we present testing loss instead of accuracy for clearer comparison. Delta demonstrates relatively high robustness,




Figure 7: System overheads of Delta.

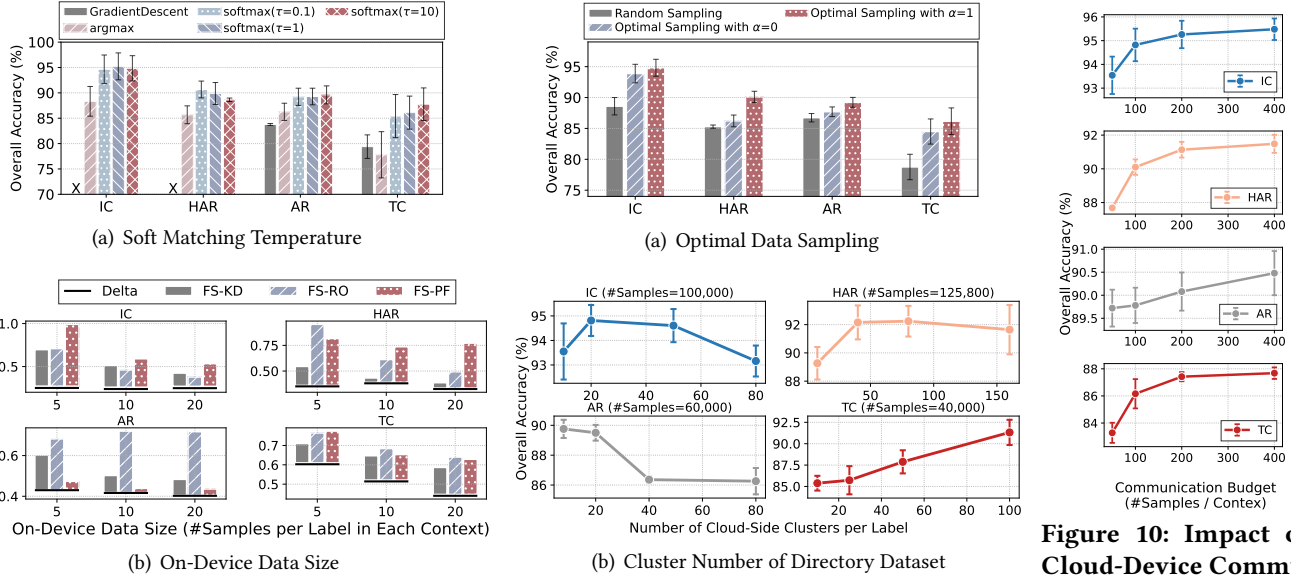


Figure 8: Device-side Sensitivity Analysis. Figure 9: Cloud-side Sensitivity Analysis.

which is attributed to 1) the effective solution of the device-side sub-problem with scarce on-device data through our soft matching strategy, and 2) the substantial performance improvement brought by the abundant cloud-side enriched data compared to additional device-side user data. Also, we observe that baselines show greater sensitivity to on-device data quantity, highlighting the critical role of on-device data enrichment and further motivates our work.

Cloud-Side Directory Dataset. Figure 9(b) plots the performance of Delta with varying numbers of cloud-side data clusters for directory dataset construction, where we replace the cloud-side data sampling scheme with random sampling to isolate the effects of directory dataset. We observe that a slight increase in cluster number can improve Delta’s performance by making directory dataset more representative and aligning the cloud-side sub-objective (2b) more closely with the overall objective (2). However, an excessively large cluster number can result in numerous similar clusters, leading to the selection of redundant data for enrichment. For stable performance, we set cluster number per label to 20.

Cloud-Side Optimal Data Sampling. To evaluate the importance of cloud-side optimal data sampling, we assess

Delta’s performance with different sampling schemes, including random sampling, optimal sampling for solely new context ($\alpha = 0$) and optimal sampling considering all contexts ($\alpha = 1$). Figure 9(a) indicates that optimal data sampling for only new context improves overall model accuracy by 5.3/0.9/1.0/5.7% for IC/HAR/AR/TC tasks. Considering past contexts further enhances accuracy by 0.9/3.9/1.5/1.7%. Notably, the most significant improvements are observed in IC and TC tasks, as the visual and textual data we used are more diverse, making random sampling less stable and effective.

Device-Cloud Communication Budget. We further evaluate Delta’s performance with varying sizes of cloud-side enriched data to simulate different communication budgets. Figure 10 shows that Delta’s performance improves significantly as the enriched data size per context increases from 50 to 100, and then stabilizes with larger data sizes. This robustness highlights Delta’s applicability for real-world devices with diverse network conditions.

7 Related Work

Cloud-Side Continual Learning aims to train ML models over non-stationary data streams to acquire new contextual knowledge without forgetting past contexts. This approach

Figure 10: Impact of Cloud-Device Communication Budget.

is inspired by the capability of biological neural networks to modulate synaptic memory and plasticity in response to dynamic inputs [72, 82]. Existing solutions include: 1) stabilizing previously-learned synaptic changes by penalizing parameter changes of the past optimal model [34, 82]; 2) expanding and pruning synaptic connections to form new synaptic memory via creating additional parameter space for new contexts and re-normalizing them with past contexts [48, 60]; 3) consolidating synaptic memory by storing the important data of past contexts and replaying them during learning new contexts [42, 44], where the data importance can be measured by representativeness [57, 62], diversity [81] or uncertainty [5]. Previous studies [26, 37, 39] have found that data replay methods provides the best trade-off between model performance and system efficiency, and thus our experiments are mainly conducted in this case. *Delta framework serves as a plug-in component to enrich on-device data and enhance performance for all these methods.*

Device-Side Continual Learning focuses on optimizing the utilization of hardware resources to implement cloud-side CL algorithms on resource-constrained devices. This includes saving storage cost through data quantization techniques [29, 56], reducing memory overhead through context-aware parameter sparsity [38, 74], accelerating data loading via hierarchical memory management [39, 46], and accelerating computation with adaptive computing resource [35, 36, 40, 71]. However, most of these works overlook the data bottleneck on mobile device (scarce, personal and unpredictable user data), and thus *Delta is complementary to existing on-device CL works focusing on hardware bottleneck.*

On-Device Data Augmentation is a powerful technique to improve model training performance by generating diverse data from existing user data, such as leveraging geometric and color space transformation and random erasing for visual images [65], using techniques grounded in physical principles for IMU signal [75], as well as employing language rule-based transformations and synonym replacement for textual data [8]. However, a significant limitation of data augmentation is that each data modality and task necessitates specifically designed augmentation techniques to accommodate unique data characteristics, making the data augmentation process cumbersome and inefficient. *Delta serves as a generally solution to complement these works by directly expanding the on-device available data.*

8 Discussion

Privacy Consideration. In *Delta* framework, the information uploaded by devices includes the directory weights, which excludes any raw user data and protects privacy like FL [51]. Unlike FL, where the transmitted model updates inherently encode specific features of training data, *Delta*'s

transmitted weights only indicate the similarity between user data and directory dataset (e.g. likelihood of weather conditions rather than pixels in IC task, probability of device placement rather than specific IMU signals in HAR task), which reveals rough context information and makes the recovery or identification of raw data more challenging. To further enhance privacy, secure aggregation techniques like secure multi-party computation [21] and homomorphic encryption [4] can be integrated into the communication and computation processes in *Delta*.

Comparison with FL. The intuitions behind *Delta* framework and FL paradigm are distinct. FL aims to leverage device-side data to develop a global model that can generalize well across diverse user contexts, i.e. *global knowledge aggregation*. In contrast, *Delta* utilizes cloud-side data to enhance the personalization of local models for individual user contexts, i.e. *local knowledge augmentation*. As a result, the applicability of FL is primarily limited by device-side constraints, including the vast number of devices, high participation rates, cross-device data heterogeneity and tolerance for communication overheads. *Delta*, on the other hand, seeks to shift the limitations to the cloud, assuming that cloud server can collect abundant public data to match different users. This aligns with the recent success of training billion-scale models over sufficiently diverse datasets for various tasks. Additionally, when confronted with extremely rare user contexts, *Delta* could still identify the most helpful and relevant cloud-side data-subsets to provide data foundation for existing model or algorithm-based augmentation methods. In conclusion, FL and *Delta* are applicable for different scenarios and could potentially be complementary.

9 Conclusion

In this work, we explore the potential of leveraging cloud-side abundant data resource to address the data bottleneck in on-device CL. We formalize the data enrichment problem and propose *Delta*, a private, efficient and effective cloud-assisted data enrichment framework for on-device CL. On extensive experiments, *Delta* shows superior model performance and system efficiency across various mobile computing tasks, data modalities and model structures.

Acknowledgments

This work was supported in part by National Key R&D Program of China (No. 2022ZD0119100), in part by China NSF grant No. 62322206, 62132018, U2268204, 62025204, 62272307, 62372296. The opinions, findings, conclusions, and recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agencies or the government. The authors thank the anonymous reviewers and the shepherd for their insightful feedbacks.

References

- [1] 2024. Apple Intelligence Preview - Apple. <https://www.apple.com/apple-intelligence/>.
- [2] 2024. Common Crawl maintains a free, open repository of web crawl data that can be used by anyone. <https://commoncrawl.org/>.
- [3] 2024. Google Smart Lens - Search What You See. <https://lens.google/>.
- [4] Abbas Acar, Hidayet Aksu, A Selcuk Uluagac, and Mauro Conti. 2018. A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys (CSUR)* 51, 4 (2018), 1–35.
- [5] Rahaf Aljundi, Eugene Belilovsky, Tinne Tuytelaars, Laurent Charlin, Massimo Caccia, Min Lin, and Lucas Page-Caccia. 2019. Online continual learning with maximal interfered retrieval. *Advances in Neural Information Processing Systems (NeurIPS)* 32 (2019).
- [6] Apple. 2023. Legal - Siri Suggestions, Search; Privacy. <https://www.apple.com/legal/privacy/data/en/siri-suggestions-search/>.
- [7] Jie Bao, Yu Zheng, and Mohamed F Mokbel. 2012. Location-based and preference-aware recommendation using sparse geo-social networking data. In *International Conference on Advances in Geographic Information Systems (SIGSPATIAL)*. 199–208.
- [8] Markus Bayer, Marc-André Kaufhold, and Christian Reuter. 2022. A survey on data augmentation for text classification. *ACM Computing Surveys (CSUR)* 55, 7 (2022), 1–39.
- [9] Romil Bhardwaj, Zhengxu Xia, Ganesh Ananthanarayanan, Junchen Jiang, Yuanhao Shu, Nikolaos Karianakis, Kevin Hsieh, Paramvir Bahl, and Ion Stoica. 2022. Ekya: Continuous learning of video analytics models on edge compute servers. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. 119–135.
- [10] Matic Broz. 2023. How many pictures are there (2024): Statistics, trends, and forecasts. <https://photutorial.com/photos-statistics/>.
- [11] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. 2020. Dark experience for general continual learning: a strong, simple baseline. In *Advances in Neural Information Processing Systems (NeurIPS)*. 15920–15930.
- [12] Dongqi Cai, Yaozong Wu, Shangguang Wang, Felix Xiaozhu Lin, and Mengwei Xu. 2023. Efficient federated learning for modern nlp. In *International Conference on Mobile Computing and Networking (MobiCom)*. 1–16.
- [13] Han Cai, Chuang Gan, Ligeng Zhu, and Song Han. 2020. Tinytl: Reduce memory, not parameters for efficient on-device learning. In *Advances in Neural Information Processing Systems (NeurIPS)*. 11285–11297.
- [14] Chengliang Chai, Jiabin Liu, Nan Tang, Guoliang Li, and Yuyu Luo. 2022. Selective data acquisition in the wild for model charging. *Proceedings of the VLDB Endowment (VLDB)* 15 (2022), 1466–1478.
- [15] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc'Aurelio Ranzato. 2019. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486* (2019).
- [16] Robert C. Daley and Peter G. Neumann. 1965. A general-purpose file system for secondary storage. In *Proceedings of the 1965 fall joint computer conference, part I, AFIPS 1965 (Fall, part I), Las Vegas, Nevada, USA, November 30 - December 1, 1965*. 213–229.
- [17] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2024. Qlora: Efficient finetuning of quantized llms. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [19] Jiahua Dong, Lixu Wang, Zhen Fang, Gan Sun, Shichao Xu, Xiao Wang, and Qi Zhu. 2022. Federated class-incremental learning. In *The IEEE / CVF Computer Vision and Pattern Recognition Conference (CVPR)*. 10164–10173.
- [20] Robert M French. 1999. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences* 3 (1999), 128–135.
- [21] Oded Goldreich. 1998. Secure multi-party computation. *Manuscript. Preliminary version* 78, 110 (1998), 1–108.
- [22] Chen Gong. [n. d.]. Supplementary Material. https://drive.google.com/drive/folders/1wZ5PHYLPocMURKUuy3x6DkZvEgOp1qc?usp=drive_link.
- [23] Chen Gong, Zhenzhe Zheng, Yunfeng Shao, Bingshuai Li, Fan Wu, and Guihai Chen. 2024. ODE: An Online Data Selection Framework for Federated Learning With Limited Storage. *IEEE/ACM Transactions on Networking (TON)* (2024).
- [24] Chen Gong, Zhenzhe Zheng, Fan Wu, Yunfeng Shao, Bingshuai Li, and Guihai Chen. 2023. To store or not? online data selection for federated learning with limited storage. In *ACM Web Conference (WWW)*. 3044–3055.
- [25] Douglas M Hawkins. 2004. The problem of overfitting. *Journal of Chemical Information and Computer Sciences* 44 (2004), 1–12.
- [26] Tyler L Hayes and Christopher Kanan. 2022. Online continual learning for embedded devices. *Conference on Lifelong Learning Agents (2022)*.
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *The IEEE / CVF Computer Vision and Pattern Recognition Conference (CVPR)*. 770–778.
- [28] Dan Hendrycks and Thomas Dietterich. 2019. Benchmarking Neural Network Robustness to Common Corruptions and Perturbations. *International Conference on Learning Representations (ICLR)* (2019).
- [29] Michael Hersche, Geethan Karunaratne, Giovanni Cherubini, Luca Benini, Abu Sebastian, and Abbas Rahimi. 2022. Constrained few-shot class-incremental learning. In *The IEEE / CVF Computer Vision and Pattern Recognition Conference (CVPR)*. 9057–9067.
- [30] Apple Intelligence. 2024. Siri - Apple. <https://www.apple.com/siri/>.
- [31] Angelos Katharopoulos and François Fleuret. 2018. Not all samples are created equal: Deep learning with importance sampling. In *International Conference on Machine Learning (ICML)*. 2525–2534.
- [32] Zixuan Ke, Bing Liu, Nianzu Ma, Hu Xu, and Lei Shu. 2021. Achieving forgetting prevention and knowledge transfer in continual learning. In *Advances in Neural Information Processing Systems (NeurIPS)*. 22443–22456.
- [33] Mehrdad Khani, Ganesh Ananthanarayanan, Kevin Hsieh, Junchen Jiang, Ravi Netravali, Yuanhao Shu, Mohammad Alizadeh, and Victor Bahl. 2023. {RECL}: Responsive {Resource-Efficient} continuous learning for video analytics. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. 917–932.
- [34] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. In *Proceedings of the National Academy of Sciences (PNAS)*. 3521–3526.
- [35] Dhireesha Kudithipudi, Anurag Daram, Abdullah M Ziyarah, Fatima Tuz Zohora, James B Aimone, Angel Yanguas-Gil, Nicholas Soures, Emre Neftci, Matthew Mattina, Vincenzo Lomonaco, et al. 2023. Design principles for lifelong learning AI accelerators. *Nature Electronics* 6 (2023), 807–822.
- [36] Young D Kwon, Jagmohan Chauhan, Hong Jia, Stylianos I Venieris, and Cecilia Mascolo. 2023. LifeLearner: Hardware-Aware Meta Continual Learning System for Embedded Computing Platforms. In *ACM Conference on Embedded Networked Sensor Systems (SenSys)*.
- [37] Young D Kwon, Jagmohan Chauhan, Abhishek Kumar, Pan Hui HKUST, and Cecilia Mascolo. 2021. Exploring system performance of continual learning for mobile and embedded sensing applications. In *ACM/IEEE Symposium on Edge Computing (SEC)*. 319–332.
- [38] Young D Kwon, Rui Li, Stylianos I Venieris, Jagmohan Chauhan, Nicholas D Lane, and Cecilia Mascolo. 2023. Tinytrain: Deep neural network training at the extreme edge. *arXiv preprint arXiv:2307.09988*

- (2023).
- [39] Soobee Lee, Minindu Weerakoon, Jonghyun Choi, Minjia Zhang, Di Wang, and Myeongjae Jeon. 2022. CarM: Hierarchical episodic memory for continual learning. In *Proceedings of the ACM/IEEE Design Automation Conference (DAC)*. 1147–1152.
- [40] Clayton Frederick Souza Leite and Yu Xiao. 2022. Resource-efficient continual learning for sensor-based human activity recognition. *ACM Transactions on Embedded Computing Systems* 21, 6 (2022), 1–25.
- [41] Chenning Li, Xiao Zeng, Mi Zhang, and Zhichao Cao. 2022. PyramidFL: A fine-grained client selection framework for efficient federated learning. In *International Conference on Mobile Computing And Networking (MobiCom)*. 158–171.
- [42] Zhizhong Li and Derek Hoiem. 2017. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* 40 (2017), 2935–2947.
- [43] Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroan Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zhou. 2020. XGLUE: A New Benchmark Dataset for Cross-lingual Pre-training, Understanding and Generation. *arXiv abs/2004.01401* (2020).
- [44] David Lopez-Paz and Marc Aurelio Ranzato. 2017. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [45] Mingqi Lv, Ling Chen, and Gencai Chen. 2013. Mining user similarity based on routine activities. *Information Sciences* 236 (2013), 17–32.
- [46] Xinyue Ma, Suyeon Jeong, Minjia Zhang, Di Wang, Jonghyun Choi, and Myeongjae Jeon. 2023. Cost-effective On-device Continual Learning over Memory Hierarchy with Miro. In *International Conference on Mobile Computing and Networking (MobiCom)*. 1–15.
- [47] Mohammad Malekzadeh, Richard G Clegg, Andrea Cavallaro, and Hamed Haddadi. 2019. Mobile sensor data anonymization. In *ACM/IEEE Conference on Internet of Things Design and Implementation (IoTDI)*. 49–58.
- [48] Arun Mallya, Dillon Davis, and Svetlana Lazebnik. 2018. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *European Conference on Computer Vision (ECCV)*. 67–82.
- [49] Pratik Mazumder, Pravendra Singh, and Piyush Rai. 2021. Few-shot lifelong learning. In *Association for the Advancement of Artificial Intelligence (AAAI)*. 2337–2345.
- [50] Michael McCloskey and Neal J Cohen. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of Learning and Motivation*. Vol. 24. 109–165.
- [51] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas. 2017. Communication-efficient learning of deep networks from decentralized data. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*. 1273–1282.
- [52] NVIDIA. 2023. Jetson Nano Developer Kit. <https://developer.nvidia.com/embedded/jetson-nano-developer-kit>.
- [53] Official Journal of the European Union. 2018. General Data Protection Regulation. <https://gdpr-info.eu/>.
- [54] Huawei Harmony OS. 2023. Data Donation. <https://developer.huawei.com/consumer/en/doc/hmscore-guides/event-donate-awareness-0000001505674356>.
- [55] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. 2020. Gdumb: A simple approach that questions our progress in continual learning. In *The European Conference on Computer Vision (ECCV)*. 524–540.
- [56] Leonardo Ravaglia, Manuele Rusci, Davide Nadalini, Alessandro Capotondi, Francesco Conti, and Luca Benini. 2021. A tinyml platform for on-device continual learning with quantized latent replays. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* 11 (2021), 789–802.
- [57] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. 2017. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*. 2001–2010.
- [58] Jorge-L Reyes-Ortiz, Luca Oneto, Albert Samà, Xavier Parra, and Davide Anguita. 2016. Transition-aware human activity recognition using smartphones. *Neurocomputing* 171 (2016), 754–767.
- [59] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Sathesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* 115 (2015), 211–252.
- [60] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. 2018. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning (ICML)*. 4548–4557.
- [61] Guangyuan Shi, Jiaxin Chen, Wenlong Zhang, Li-Ming Zhan, and Xiao-Ming Wu. 2021. Overcoming catastrophic forgetting in incremental few-shot learning by finding flat minima. In *Advances in Neural Information Processing Systems (NeurIPS)*. 6747–6761.
- [62] Dongsub Shim, Zheda Mai, Jihwan Jeong, Scott Sanner, Hyunwoo Kim, and Jongseong Jang. 2021. Online class-incremental continual learning with adversarial shapley value. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. 9630–9638.
- [63] Jaemin Shin, Yuanchun Li, Yunxin Liu, and Sung-Ju Lee. 2022. Fed-balancer: Data and pace control for efficient federated learning on heterogeneous clients. In *International Conference on Mobile Systems, Applications and Services (MobiSys)*. 436–449.
- [64] Muhammad Shoaib, Stephan Bosch, Ozlem Durmaz Incel, Hans Scholten, and Paul JM Havinga. 2014. Fusion of smartphone motion sensors for physical activity recognition. *Sensors* 14, 6 (2014), 10146–10176.
- [65] Connor Shorten and Taghi M Khoshgoftaar. 2019. A survey on image data augmentation for deep learning. *Journal of big data* 6, 1 (2019), 1–48.
- [66] Daniel L Silver, Qiang Yang, and Lianghao Li. 2013. Lifelong machine learning systems: Beyond learning algorithms. In *Association for the Advancement of Artificial Intelligence (AAAI)*.
- [67] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [68] Allan Stisen, Henrik Blunck, Sourav Bhattacharya, Thor Siiger Prentow, Mikkel Baun Kjærsgaard, Anind Dey, Tobias Sonne, and Mads Møller Jensen. 2015. Smart devices are different: Assessing and mitigating mobile sensing heterogeneities for activity recognition. In *ACM Conference on Embedded Networked Sensor Systems (SenSys)*. 127–140.
- [69] Xiaoyu Tao, Xiaopeng Hong, Xinyuan Chang, Songlin Dong, Xing Wei, and Yihong Gong. 2020. Few-shot class-incremental learning. In *The IEEE / CVF Computer Vision and Pattern Recognition Conference (CVPR)*. 12183–12192.
- [70] Sebastian Thrun and Tom M Mitchell. 1995. Lifelong robot learning. *Robotics and Autonomous Systems* 15 (1995), 25–46.
- [71] Huangshi Tian, Minchen Yu, and Wei Wang. 2018. Continuum: A platform for cost-aware, low-latency continual learning. In *Proceedings of the ACM Symposium on Cloud Computing (SoCC)*. 26–40.
- [72] Liyuan Wang, Xingxing Zhang, Hang Su, and Jun Zhu. 2024. A comprehensive survey of continual learning: theory, method and application. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* (2024).

- [73] Pete Warden. 2018. Speech commands: A dataset for limited-vocabulary speech recognition. *arXiv preprint arXiv:1804.03209* (2018).
- [74] Minhui Xie, Kai Ren, Youyou Lu, Guangxu Yang, Qingxing Xu, Bihai Wu, Jiazhen Lin, Hongbo Ao, Wanhong Xu, and Jiwu Shu. 2020. Kraken: memory-efficient continual learning for large-scale real-time recommendations. In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. 1–17.
- [75] Huatao Xu, Pengfei Zhou, Rui Tan, and Mo Li. 2023. Practically Adopting Human Activity Recognition. In *International Conference on Mobile Computing and Networking (MobiCom)*. 1–15.
- [76] Yikai Yan, Chaoyue Niu, Renjie Gu, Fan Wu, Shaojie Tang, Lifeng Hua, Chengfei Lyu, and Guihai Chen. 2022. On-Device Learning for Model Personalization with Large-Scale Cloud-Coordinated Domain Adaption. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*. 2180–2190.
- [77] Jianbo Yang, Minh Nhut Nguyen, Phyo Phyo San, Xiaoli Li, and Shonali Krishnaswamy. 2015. Deep convolutional neural networks on multi-channel time series for human activity recognition.. In *International Joint Conference on Artificial Intelligence (IJCAI)*, Vol. 15. 3995–4001.
- [78] Jie Yin, Qiang Yang, and Jeffrey Junfeng Pan. 2008. Sensor-based abnormal human-activity detection. *IEEE Transactions on Knowledge and Data Engineering* 20 (2008), 1082–1090.
- [79] Xue Ying. 2019. An overview of overfitting and its solutions. In *Journal of Physics: Conference Series*, Vol. 1168. 022022.
- [80] Jaehong Yoon, Wonyong Jeong, Giwoong Lee, Eunho Yang, and Sung Ju Hwang. 2021. Federated continual learning with weighted inter-client transfer. In *International Conference on Machine Learning (ICML)*. 12073–12086.
- [81] Jaehong Yoon, Divyam Madaan, Eunho Yang, and Sung Ju Hwang. 2022. Online Coreset Selection for Rehearsal-based Continual Learning. In *International Conference on Learning Representations (ICLR)*.
- [82] Friedemann Zenke, Ben Poole, and Surya Ganguli. 2017. Continual learning through synaptic intelligence. In *International Conference on Machine Learning (ICML)*. 3987–3995.
- [83] Linglan Zhao, Jing Lu, Yunlu Xu, Zhazhan Cheng, Dashan Guo, Yi Niu, and Xiangzhong Fang. 2023. Few-Shot Class-Incremental Learning via Class-Aware Bilateral Distillation. In *The IEEE / CVF Computer Vision and Pattern Recognition Conference (CVPR)*. 11838–11847.
- [84] ZipDo. 2023. Essential Apple Siri Statistics In 2024. <https://zipdo.co/statistics/apple-siri/>.