

MIT Open Access Articles

AirMixML: Over-the-Air Data Mixup for Inherently Privacy-Preserving Edge Machine Learning

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Koda, Yusuke, Park, Jihong, Bennis, Mehdi, Vepakomma, Praneeth and Raskar, Ramesh. 2021. "AirMixML: Over-the-Air Data Mixup for Inherently Privacy-Preserving Edge Machine Learning." 2021 IEEE Global Communications Conference (GLOBECOM).

As Published: 10.1109/GLOBECOM46510.2021.9685232

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <https://hdl.handle.net/1721.1/146533>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



AirMixML: Over-the-Air Data Mixup for Inherently Privacy-Preserving Edge Machine Learning

Yusuke Koda[†], Jihong Park[‡], Mehdi Bennis[†], Praneeth Vepakomma[§], and Ramesh Raskar[§]

[†] Centre for Wireless Communications, University of Oulu, 90014 Oulu, Finland

[‡] School of Information Technology, Deakin University, Geelong, VIC 3220, Australia

[§] Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA-02139, USA

Abstract—Wireless channels can be inherently privacy preserving by distorting the received signals due to channel noise, and superpositioning multiple signals over-the-air. By harnessing these natural distortions and superpositions by wireless channels, we propose a novel privacy-preserving machine learning (ML) framework at the network edge, coined *over-the-air mixup ML* (AirMixML). In AirMixML, multiple workers transmit analog-modulated signals of their private data samples to an edge server who trains an ML model using the received noisy-and-superpositioned samples. AirMixML coincides with model training using mixup data augmentation achieving comparable accuracy to that with raw data samples. From a privacy perspective, AirMixML is a differentially private (DP) mechanism limiting the disclosure of each worker’s private sample information at the server, while the worker’s transmit power determines the privacy disclosure level. To this end, we develop a fractional channel-inversion power control (PC) method, α -Dirichlet mixup PC (DirMix(α)-PC), wherein for a given global power scaling factor after channel inversion, each worker’s local power contribution to the superpositioned signal is controlled by the Dirichlet dispersion ratio α . Mathematically, we derive a closed-form expression clarifying the relationship between the local and global PC factors to guarantee a target DP level. By simulations, we provide DirMix(α)-PC design guidelines to improve accuracy, privacy, and energy-efficiency. Finally, AirMixML with DirMix(α)-PC is shown to achieve reasonable accuracy compared to a privacy-violating baseline with neither superposition nor PC.

I. INTRODUCTION

Big data is instrumental in building high-quality machine learning (ML) models. One compelling source of big data is edge devices ranging from phones to internet-of-thing (IoT) sensors. These devices generate a massive amount of data that is spatially dispersed and often privacy-sensitive. Analog federated learning (FL) has a great potential in utilizing such a user-generated data while ensuring data privacy [1]–[4]. The underlying principle behind analog FL is to locally train ML models at edge devices and to aggregate the model parameters at a server through wireless channels using analog-modulated signals. Consequently, the server only receives the ‘over-the-air’ superpositioned signals, within which each worker’s model parameter information is inherently hidden. Furthermore, random channel fading [4] and additive channel noise [2], [3] additionally distort the received signals, thereby improving robustness against model inversion attacks recovering raw data by adversaries [5]. Last but not least, instead of avoiding inter-device interfering signals using orthogonal bandwidth allocations, analog FL harnesses them, significantly reducing the required bandwidth [1].

While interesting existing analog FL frameworks rely on local model training that is not always feasible for battery-limited and memory-limited edge devices, particularly for deep neural network models. This mandates the development of extremely lightweight edge ML frameworks without on-device training while retaining the benefits of communication-efficiency and data privacy from analog transmissions. Spurred by this, we propose a novel privacy-preserving edge ML

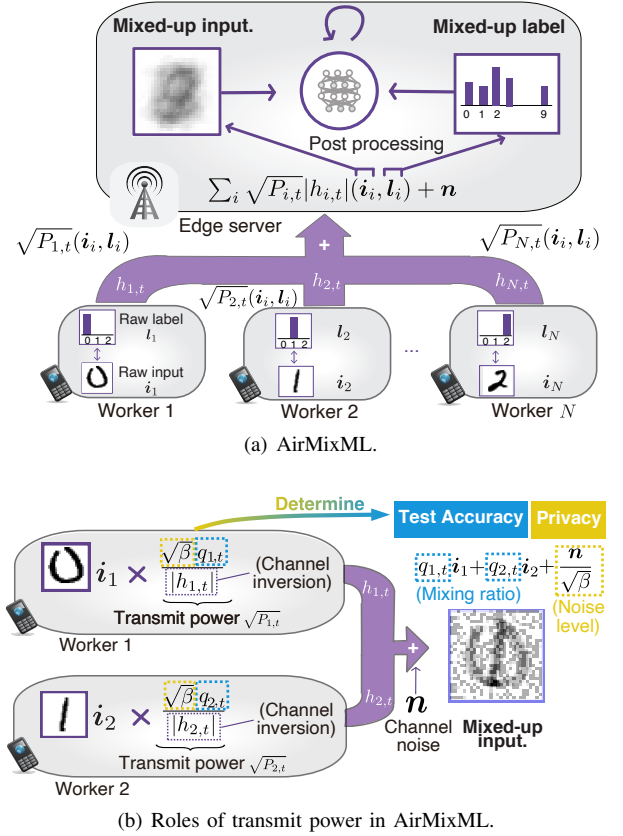


Fig. 1. Illustration of (a) over-the-air mixup machine learning (AirMixML) for privacy preserving edge learning and (b) roles of transmit power in AirMixML in a two worker case. The edge server receives mixed-up samples from multiple workers under additive noise. In transmit power $P_{i,t}$, the term $q_{i,t}$ determines mixup ratios, and scaling factor β determines the level of differential privacy achieved via additive noise.

framework without on-device training, while harnessing wireless channel superpositioning and noise additive properties, coined *over-the-air mixup ML* (AirMixML).

As illustrated in Fig. 1(a), the key new element of AirMixML is to offload model training from edge devices to an edge server who collects private data samples from edge devices using analog-modulated signals superpositioned *over-the-air*, as opposed to analog FL collecting model parameters [1]–[4]. From an ML perspective, AirMixML coincides with model training using *mixup* augmented data [6], i.e., linearly superpositioned data samples, which is known to achieve comparable accuracy to that with raw data samples. From a privacy perspective, AirMixML guarantees *differential privacy* (DP) [7], i.e., one cannot accurately infer about whether a sample is included in a device’s local dataset. The key to

success is controlling transmit power that affects both ML accuracy and data privacy, as depicted in Figs. 1(b). Indeed, the transmit power $P_{i,t}$ of the i -th device at time t controls the data sample mixing ratio, and adjusts the relative noise level, i.e., $\sqrt{P_{i,t}}|h_{i,t}|$ and \mathbf{n} in the middle of Fig 1(a), respectively.

In this respect, we propose a novel transmit power control (PC) method, termed α -Dirichlet mixup PC (DirMix(α)-PC) where the transmit power is set as $\sqrt{P_{i,t}} = \sqrt{\beta}q_{i,t}/|h_{i,t}|$. Here, we fractionally cancel out the channel fading $|h_{i,t}|$ by multiplying $\sqrt{\beta}/|h_{i,t}|$, i.e., fractional channel inversion with a global power scaling factor β . Then, we change the sample mixing ratio by multiplying $q_{i,t}$ sampled from a Dirichlet distribution with the dispersion ratio α . For an infinite α , each device's local power uniformly contributes to the superpositioned signal, yielding the equal mixing ratio, and otherwise the mixing ratio becomes imbalanced, which may improve the ML accuracy.

Contributions. The main contributions of this work are summarized as follows.

- AirMixML is the first privacy-preserving edge ML framework harnessing over-the-air signal superposition and an additive channel noise without on-device training, in stark contrast to existing analog FL methods rooted in on-device training [2]–[4].
- This is the first work analyzing DP with mixup with unequal mixing ratios (see **Proposition 1**), as opposed to prior works with equal mixing ratios [8], [9].
- We derive a closed-form expression of the power scaling factor β with the channel inversion to guarantee a target DP level, revealing the dependence of β on AirMixML system parameters, e.g., data dimension, channel noise, the number of scheduled workers (see **Proposition 2** and **Remark 1**).
- Based on the derived expression, we provide a guideline for setting a mixing ratio $q_{i,t}$, i.e., the mixing ratio should be set proportionally with the target DP level (see **Remark 2** and **Table I**).

Related Works. Privacy-preserving edge ML has been recently investigated. For example, in FL [10], [11], devices perform training of ML models and sequentially exchange their model parameters with an edge server. Channel superposition of analog modulation is proposed to achieve scalability [1], [4], [12], [13] and DP [2]–[4]. However, these approaches still require on-device training of entire ML models. In parallel, DP for mixup has been studied [8], [9], where equal mixing ratios are assumed. Apart from these, we investigate DP in a Dirichlet mixup scenario where mixup ratios are not necessarily identical across workers.

II. OVER-THE-AIR MIXUP FOR PRIVACY-PRESERVING EDGE ML

Consider N workers that transmit their private input sample $\mathbf{i}_i \in [0, 1]^{d_X}$ and associated one-hot label $\mathbf{l}_i \in \{0, 1\}^{d_Y}$ to an edge server. In AirMixML, each worker harnesses channel superposition to transmit the pair of input-label, $(\mathbf{i}_i, \mathbf{l}_i)$ without revealing their raw data $(\mathbf{i}_i, \mathbf{l}_i)$ to curious eavesdroppers including a honest-but-curious edge server. In what follows, we detail the transmission model and server-side ML procedure in AirMixML.

A. Transmission Model

We consider that transmission time is slotted as in a time division multiple access (TDMA) channel, and the above input-label pair is transmitted within each time slot $t = 1, 2, \dots, T$. We assume a block fading channel, where the channel coefficient is constant over, at least, a time slot, and

we let $h_{i,t} \in \mathbb{C}$ denote the channel coefficient between the edge server and worker $i \in \mathcal{N} := \{1, 2, \dots, N\}$. Within each time slot, randomly-scheduled workers $\mathcal{N}_t \subseteq \mathcal{N}$ transmit all analog-modulated symbols of each element in \mathbf{i}_i and \mathbf{l}_i ; namely, they transmit $d_X + d_Y$ symbols in each time slot. Considering a symbol level synchronization among each worker, simultaneous co-channel transmission, and phase shift cancellation in $h_{i,t}$, for $d = 1, 2, \dots, d_X + d_Y$, the d th received symbol in time slot t is given by

$$y_t^{(d)} = \sum_{i \in \mathcal{N}_t} \sqrt{P_{i,t}} |h_{i,t}| s_i^{(d)} + n_t^{(d)}, \quad (1)$$

where $P_{i,t}$ denotes the transmit power at worker i at time slot t , and $n_t^{(d)}$ denotes the channel noise. We consider an additive white Gaussian noise (AWGN) with $n_t^{(d)} \sim \mathcal{CN}(0, \sigma_n^2)$, where σ_n^2 denotes the noise power. Lastly, $s_i^{(d)}$ denotes the d th transmitted symbol at worker i . Without loss of generality, we consider that for $d \in \{1, 2, \dots, d_X\}$, $s_i^{(d)} = i_i^{(d)}$, whereas for $d \in \{d_X + 1, d_X + 2, \dots, d_X + d_Y\}$, $s_i^{(d)} = l_i^{(d - d_X)}$, where $i_i^{(d)} \in [0, 1]$ and $l_i^{(d)} \in \{0, 1\}$ denote the d th element in \mathbf{i}_i and \mathbf{l}_i , respectively.

Formatting $y_t^{(d)}$ in a vectorized form yields *mixup samples* formed as the linear aggregation of input samples and labels among the scheduled workers, which are leveraged in the subsequent learning procedure at the edge server. First, by formatting the real-part of $y_t^{(1)}, \dots, y_t^{(d_X)}$ in a vectorized form $\mathbf{y}_{I,t}$, we obtain:

$$\mathbf{y}_{I,t} = \sum_{i \in \mathcal{N}_t} \sqrt{P_{i,t}} |h_{i,t}| \mathbf{i}_i + \mathbf{n}_{I,t}, \quad (2)$$

where $\mathbf{n}_{I,t} \sim \mathcal{N}(0, (\sigma_n^2/2) \mathbf{I}_{d_X})$. Note that for $n \in \mathbb{N}$, \mathbf{I}_n denotes the unit matrix of size $n \times n$. In (2), we can find that $\mathbf{y}_{I,t}$ exactly means the superposed input samples among the workers. Similarly, by formatting the real-part of $y_t^{(d_X+1)}, \dots, y_t^{(d_X+d_Y)}$ in a vectorized form $\mathbf{y}_{L,t}$, we obtain:

$$\mathbf{y}_{L,t} = \sum_{i \in \mathcal{N}_t} \sqrt{P_{i,t}} |h_{i,t}| \mathbf{l}_i + \mathbf{n}_{L,t}, \quad (3)$$

where $\mathbf{n}_{L,t} \sim \mathcal{N}(0, (\sigma_n^2/2) \mathbf{I}_{d_Y})$. Similarly to (2), in (3), $\mathbf{y}_{L,t}$ refers to the superposed labels among the workers.

B. Machine Learning Procedure at Edge Server with Mixed Up Samples

After obtaining superposed input-label pairs $(\mathbf{y}_{I,t}, \mathbf{y}_{L,t})_{t=1}^T$, the edge server performs an ML procedure. Let $f(\cdot, \cdot; \boldsymbol{\theta})$ denote the loss function, and the objective of the edge server is to minimize the loss function by using the superposed input-label pairs. This problem boils down to the optimization problem below.

$$\underset{\boldsymbol{\theta}}{\text{minimize}} \quad \sum_{t=1}^T f(\mathbf{y}_{I,t}/b_t, \mathbf{y}_{L,t}/b_t; \boldsymbol{\theta}), \quad (4)$$

where $b_t := \sum_{i \in \mathcal{N}_t} \sqrt{P_{i,t}} |h_{i,t}|$ is the coefficient that normalizes both mixed-up input samples and labels.

III. DIFFERENTIALLY PRIVATE DIRICHLET MIXUP AND POWER CONTROL

A. Dirichlet Mixup

First, we introduce the Dirichlet mixup where the mixup ratios are set as realizations drawn from a Dirichlet distribution. Let the aforementioned weights for the channel superposition be denoted by $\lambda_t := (\sqrt{P_{i,t}} |h_{i,t}|)_{i \in \mathcal{N}_t}$. More




Fig. 2. Visualization of the impact of the dispersion parameter α on mixup ratios, i.e., each element in q_t .

formally, Dirichlet mixup is characterized by the following formula:

$$\lambda_t \propto q_t, \quad q_t \sim \text{Dir}(\alpha p_t), \quad (5)$$

for $t = 1, 2, \dots, T$, where $p_t = (p_{i,t})_{i \in \mathcal{N}_t}$, and $\sum_{i \in \mathcal{N}_t} p_{i,t} = 1$. The distribution p_t is called prior distribution and is generally set as the uniform distribution, i.e., $p_{i,t} = 1/|\mathcal{N}_t|$. In (5), α is the dispersion parameter. Note that each entry in q_t is more than or equal to zero, and the sum of them equals one. The motivation for introducing this Dirichlet mixup is to flexibly consider the best mixing ratio q_t under privacy constraints provided in Section III-C. This is done by achieving a mid-level model mixup ratio between the following two extreme baselines:

Baseline 1. non-Mixup: This baseline does not yield a mixture but results in receiving a sample from one worker. This is done via setting the mixup ratio of the worker to one and setting those of the residual workers to zero. As shown in the subsequent section, this is preferable without any privacy constraints to achieve higher test accuracy.

Baseline 2. Equal mixup: This baseline mixes the samples of all scheduled workers equally by setting the mixup ratios as p_t . As shown in the subsequent section, this is preferable under stringent privacy constraints to achieve higher test accuracy.

It should be noted that our proposed Dirichlet mixup subsumes these two baselines as extreme cases, and hence, we can systematically target mid-level mixup ratios. This is because the parameter α characterizes the concentration of the mass in q_t into the entries as shown in Fig. 2. More concisely, as $\alpha \rightarrow 0$, the mass in q_t concentrates on one entry, implying that the mixup ratio of one entry becomes one whereas those of residual entries become zero. This exactly coincides with the non-Mixup baseline. Meanwhile, as $\alpha \rightarrow \infty$, the mass in q_t spreads into all entries, implying that q_t becomes closer to the uniform distribution p_t , which exactly coincides with the equal mixup baseline. Hence, by setting α as non-extreme values, we can flexibly target mid-level mixup ratios.

B. DirMix(α)-PC: Transmit Power Control Strategy for Dirichlet Mixup

After sampling q_t in each time slot, each worker performs a transmit power control so that the weights of the channel superposition of all workers λ_t become proportional to q_t . This is done via channel inversion, where each worker cancels its channel coefficient. We consider that the edge server surveys channel coefficients $|h_{i,t}|$ for each scheduled worker and informs each worker of the channel coefficient along with the targeted mixup ratio $q_{i,t}$. Given this, each worker performs transmit power control as follows:

$$P_{i,t} = \beta q_{i,t}^2 / |h_{i,t}|^2, \quad (6)$$

where $\beta \geq 0$ is the constant among all scheduled workers and scales transmit power not to exceed a transmit power constraint for each worker, which we call scaling factor, hereinafter. In

this transmit power control, the normalized mixed-up sample $\mathbf{y}_{I,t}/b_t$ and $\mathbf{y}_{L,t}/b_t$ are respectively given by:

$$\sum_{i \in \mathcal{N}_t} q_{i,t} \mathbf{i}_i + \mathbf{n}_{I,t} / \sqrt{\beta}, \quad (7)$$

$$\sum_{i \in \mathcal{N}_t} q_{i,t} \mathbf{l}_i + \mathbf{n}_{L,t} / \sqrt{\beta}, \quad (8)$$

where the normalized weights for training samples exactly equal $q_{i,t}$. We term this power control policy DirMix(α)-PC.

C. DP(ϵ)-DirMix(α)-PC: Differentially Private Dirichlet Mixup Transmit Power Control

DP quantifies the impact of each worker's data on perturbed results calculated from all workers' data (e.g., the average of the data perturbed by a random noise). This calculation is referred to as *randomized mechanism*. In our case of AirMixML, exposing the mixed-up samples in (7) and (8) is also a randomized mechanism, and hence, we can quantify the privacy level of DirMix(α)-PC via DP. The definition of DP is as follows:

Definition 1. (Differential Privacy [7]) Let d and d' denote adjacent datasets, implying that d' can be formed by removing one data from d . A randomized mechanism \mathcal{M} is (ϵ, δ) -differentially private if for any pair of adjacent dataset d and d' and any sort of possible outcome $\mathcal{S} \subseteq \text{Range}(\mathcal{M})$, we obtain

$$\mathbb{P}(\mathcal{M}(d) \in \mathcal{S}) \leq e^\epsilon \mathbb{P}(\mathcal{M}(d') \in \mathcal{S}) + \delta, \quad (9)$$

where $\text{Range}(\mathcal{M})$ denotes the set comprised by all possible outcomes of a mechanism \mathcal{M} . The values ϵ and δ indicate the similarity in the distribution of the outcomes of the mechanism calculated by a dataset d , i.e., the mixup from all workers in our case, and those calculated by a dataset d' , i.e., the mixup from all workers, but removing one worker. The lower ϵ and δ result in similar distributions of the two outcomes of $\mathcal{M}(d)$ and $\mathcal{M}(d')$ and hence indicate the higher privacy.

Power Control Strategy for Differential Privacy. The scaling factor β is controlled to satisfy a required DP level ϵ and δ . This is done by deriving the following proposition that quantifies the DP level of exposing the mixed-up samples for $t = 1, 2, \dots, T$ in (7) and (8).

Proposition 1. Let r define $|\mathcal{N}_t|/|\mathcal{N}|$ and be constant over all time slots. Exposing the mixed samples in (7) and (8) for $t = 1, 2, \dots, T$ is (ϵ, δ) -differentially private such that:

$$\epsilon = \min_{\gamma \in \{2, 3, \dots\}} \sum_{t=1}^T \epsilon'_t(\gamma) + \frac{\ln(1/\delta)}{\gamma-1}, \quad (10)$$

where $\epsilon'_t(\gamma) =$

$$\frac{1}{\gamma-1} \ln \left(1 + r^2 \binom{\gamma}{2} \min\{4(\exp(\epsilon_t(2)) - 1), 2\exp(\epsilon_t(2))\} \right. \\ \left. + 4 \sum_{j=3}^{\gamma} r^j \binom{\gamma}{j} \sqrt{C_t(2\lfloor j/2 \rfloor) \cdot C_t(2\lceil j/2 \rceil)} \right), \quad (11)$$

$$C_t(x) = \sum_{i=0}^x (-1)^i \binom{x}{i} \exp(i-1)\epsilon_t(i), \quad \text{and} \quad (12)$$

$$\epsilon_t(\gamma) = \frac{\gamma}{2} \cdot \frac{\max_{i \in \mathcal{N}_t} q_{i,t}^2 (d_X + d_Y)}{(\sigma_n^2/2\beta)}. \quad (13)$$

Proof. Our proof leverages recent developments in Rényi DP (RDP) defined as Definition 3 in [14]. If a randomized mechanism is proven to be $(\gamma, \epsilon'(\gamma))$ -RDP, it is also $(\epsilon'(\gamma) + \ln(1/\delta)/(\gamma-1), \delta)$ -DP for any $\delta > 0$ and γ . Hence, any RDP mechanisms can be translated into a DP mechanism

by minimizing $\epsilon'(\gamma) + \ln(1/\delta)/(\gamma-1)$ with respect to γ , which is the reason for (10).

Given this, the problem boils down to proving that the joint mechanisms of (7) and (8) iterated for $t = 1, 2, \dots, T$ are $(\gamma, \sum_{t=1}^T \epsilon'_t(\gamma))$ -RDP, which is completed as follows: First, as shown in Appendix, releasing one mixup samples for each time-slot is shown to be $(\gamma, \epsilon'_t(\gamma))$ -RDP. Finally, the proof is completed by leveraging the composition theorem (Proposition 1 in [14]), which states that: Simultaneously revealing consequences from $(\gamma, \epsilon'_{(1)}(\gamma))$ and $(\gamma, \epsilon'_{(2)}(\gamma))$ -RDP mechanisms is $(\gamma, \epsilon'_{(1)}(\gamma) + \epsilon'_{(2)}(\gamma))$ -RDP. \square

By solving (10) with respect to β , we can target any desired privacy level, while it is highly intractable. Hence, we use the following loose bound of the privacy level.

Corollary 1. The above privacy level in (10) is bounded by:

$$\epsilon \leq \sum_{t=1}^T \ln \left(1 + r^2 \min \left\{ 4 \left(e^{\frac{\max_{i,t} q_{i,t}^2 (d_X + d_Y)}{(\sigma^2/2\beta)}} - 1 \right), 2e^{\frac{\max_{i,t} q_{i,t}^2 (d_X + d_Y)}{(\sigma^2/2\beta)}} \right\} \right) + \ln(1/\delta). \quad (14)$$

Proof. Substituting $\gamma = 2$ into (10), we obtain (14). \square

Using the loose upper bound, the power scaling factor β to target the desired privacy level is as follows:

Proposition 2 (Guideline on β). Exposing the mixed samples via DirMix(α)-PC is (ϵ, δ) -differentially private, if the power scaling factor β satisfies the following conditions.

$$\beta = \begin{cases} \frac{(\sigma_n^2/2) \ln \frac{e^{\frac{\epsilon + \ln \delta}{T}} - 1}{2r^2}}{\max_{i \in \mathcal{N}_t} q_{i,t}^2 (d_X + d_Y)}, & \text{if } \epsilon \geq T \ln(1 + 4r^2) - \ln \delta \\ \frac{(\sigma_n^2/2) \ln \frac{e^{\frac{\epsilon + \ln \delta}{T}} - (1 - 4r^2)}{4r^2}}{\max_{i \in \mathcal{N}_t} q_{i,t}^2 (d_X + d_Y)}, & \text{otherwise} \end{cases} \quad (15)$$

Proof. For $\epsilon_1 \leq \epsilon_2$, a randomized mechanism with (ϵ_1, δ) -DP is also (ϵ_2, δ) -DP. Hence, based on (14), the solution of the following equation with respect to β is sufficient to yield (ϵ, δ) -DP.

$$\frac{\epsilon - \ln(1/\delta)}{T} = \ln \left(1 + r^2 \min \left\{ 4 \left(e^{\frac{\max_{i,t} q_{i,t}^2 (d_X + d_Y)}{(\sigma^2/2\beta)}} - 1 \right), 2e^{\frac{\max_{i,t} q_{i,t}^2 (d_X + d_Y)}{(\sigma^2/2\beta)}} \right\} \right). \quad (16)$$

Remark 1 (Guideline on Transmit Power and (ϵ, δ) -DP). In (15), the scaling factor β monotonously decreases when both ϵ and δ decrease. This indicates that the transmit power should be smaller when a more stringent privacy level is required (Recall that a smaller (ϵ, δ) means a higher privacy level.) Note that β also monotonously decreases as the noise variance σ_n decreases, indicating that the transmit power should be carefully set according to this noise variance.

Remark 2 (Guideline on α). From (15), we can provide a guideline to determine α by focusing on the term $\max_{i \in \mathcal{N}_t} q_{i,t}^2$. As the parameter α determines how much the probability mass in q_t concentrates on one entry, a larger α indicates smaller $\max_{i \in \mathcal{N}_t} q_{i,t}^2$ (See Fig. 2). Hence, if a strict privacy level is required, the parameter α should be set as a larger value to enhance the signal-to-noise ratio and thereby to enhance the model performance, which is validated in Section IV-B. Meanwhile, when the privacy constraint is not severe, and a

Algorithm 1 AirMixML with DP(ϵ)-DirMix(α)-PC

```

1: Initialize: model parameter in edge server  $\theta$ , targeted privacy
   level  $(\epsilon, \delta)$ , dispersion parameter  $\alpha$ 
2: Mixup in over-the-air:
3: for Each time slot  $t \in \{1, 2, \dots, T\}$  do
4:   Edge Server:
5:     Randomly select workers  $\mathcal{N}_t$ 
6:     Probe channel coefficients  $(h_{i,t})_{i \in \mathcal{N}_t}$ 
7:     Sample mixup coefficients  $q_t$  according to (5)
8:     Randomly assign each element of  $q_t$  to workers
9:     Distribute  $q_{i,t}$ ,  $h_{i,t}$ , and  $\sigma_n$  to worker  $i \in \mathcal{N}_t$ 
10:   Workers  $i \in \mathcal{N}_t$ :
11:     Determine power scaling factor  $\beta$  as in (15)
12:     Perform channel inversion as in (6)
13:     Send  $\sqrt{P_{i,t}} \mathbf{z}_i$  and  $\sqrt{P_{i,t}} \mathbf{l}_i$  to edge server
14:   end for
15: Train ML model:
16: Edge Server:
17:   Optimize model parameter as in (4)

```

sufficiently large signal-to-noise ratio is ensured regardless of $\max_{i \in \mathcal{N}_t} q_{i,t}^2$, α should be set as a smaller value to avoid to intensively mix samples, thereby enhancing the model performance, which is also validated in Section IV-B.

We term this power control policy DP(ϵ)-DirMix(α)-PC. Algorithm 1 summarizes the overall procedure of AirMixML with DP(ϵ)-DirMix(α)-PC.

IV. NUMERICAL RESULTS

To study the effectiveness of AirMixML and DP(ϵ)-DirMix(α)-PC, we consider the following workers, wireless environmental, and edge ML training settings.

Workers Setting. We uniformly distribute workers in a 500 m \times 500 m square, each of which possesses one training sample. As training samples, we consider two datasets referred to as dataset S and dataset L, where the input and label samples have small and large dimensionalities, respectively. For dataset S, we leverage the Iris dataset [15], where each data sample consists of four features of widths/heights of the sepal/petal and a class label indicating iris species. In dataset S, $d_X = 4$ and $d_Y = 3$. For dataset L, we deploy 2000 workers, each of which holds one training sample randomly chosen from 100 samples. For dataset L, we leverage a MNIST dataset [16], where each data sample consists of a hand-written digit with 28×28 pixels and a class label indicating the written number. For dataset L, $d_X = 784$ and $d_Y = 10$. When using dataset L, we deploy 60000 workers that hold one training sample randomly chosen from 60000 samples.

Wireless Environmental Setting. As a channel model, we consider a large-scale path-loss and small-scale fading, where the channel coefficient $|h_{i,t}|$ is given by: $|h_{i,t}| = \sqrt{\beta_U} d^{-n/2} |g_{i,t}|$, where β_U , d , n , $|g_{i,t}|$ are the path-loss with the unit distance, distance, path-loss exponent, and the small-scale fading coefficient, respectively. For small-scale fading, we consider Rician and Rayleigh fading, where $|g_{i,t}|$ follows a Rician distribution with a Rician factor K and Rayleigh distribution with unit-variance, respectively. Note that the smaller Rician factor yields more intensive fluctuation in $|g_{i,t}|$, and when $K \rightarrow 0$, the Rician fading is equivalent to the Rayleigh fading. The unit path-loss, noise variance σ_n^2 , and maximum transmit power P_{\max} are set as -32 dB, -114 dBm, and 23 dBm, respectively. The workers perform analog mixup for $T = 1000$ and $T = 100000$ time slots when using datasets S and L, respectively.




Fig. 3. Impact of path-loss exponent on model performance in DirMix(α)-PC, NonMix-PC, and EqualMix-PC.

Edge ML Training Settings. When using dataset S, the ML model at the edge server consists of two fully connected layers with 32 and 16 units. When using dataset L, the ML model at the edge server consists of two convolutional layers and two fully connected layers with 100 units. The convolutional layers have 32 and 48 filters with the size of 5×5 , each of which is followed by a max-pooling operation with the pooling dimension of 2×2 . The stride of filtering is one whereas that of the pooling is two. Both models are trained by solving the optimization problem in (4) with the categorical cross-entropy loss. The training is done via the Adam optimizer [17] with the learning rate of 1.0×10^{-3} , the decay parameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The batch sizes are 32 for dataset S and 64 for dataset L. The elapsed training epochs are 500 for dataset S and 10 for dataset L.

Based on the aforementioned settings, we investigate the performance of AirMixML with DP(ϵ)-DirMix(α)-PC in terms of test accuracy and energy consumption, under different channel conditions and privacy requirements in the following subsections, respectively.

A. Impact of Wireless Channel Conditions

First, we demonstrate the feasibility of the proposed AirMixML and DirMix(α)-PC in various channel conditions using the dataset L. To focus on this objective, we set the scaling factor β as the maximum value under maximum transmit power constraint without the DP constraint, which is: $P_{\max} \min_{i \in \mathcal{N}_t} |h_{i,t}|^2 / q_{i,t}^2$. We term the non-mixup and equal-mixup baselines as NonMix-PC and EqualMix-PC, respectively.

Impact of Path-Loss Exponents. In Fig. 3, we show the impact of the path loss exponent on the model performance in NonMix-PC, EqualMix-PC, and DirMix(α)-PC with $\alpha = 5$. Note that this is without fading effects. DirMix(α)-PC outperforms the EqualMix-PC baseline in both the path loss exponents of $n = 2$ and $n = 4$ with a accuracy loss of less than 2% relative to NonMix-PC. Meanwhile, the model performance becomes poorer as the path loss exponent increases. This is attributable to the fact that a larger path loss results in a smaller channel coefficient $|h_{i,t}|$, which mandates workers to set a lower scaling factor β . Recalling (7) and (8), the lower scaling factor β indicates a larger noise relative to the mixtures, which decreases the model performance.

Impact of Fading. In Fig. 4(a), we show the impact of the fading effect on the model performance in DirMix(α)-PC with $\alpha = 5$ with the path loss exponent of $n = 4$. From Fig. 4(a), the fading effect harms accuracy, and as the Rician factor K increases, the accuracy becomes poorer. This is because similarly to the path loss exponent, the fading effect

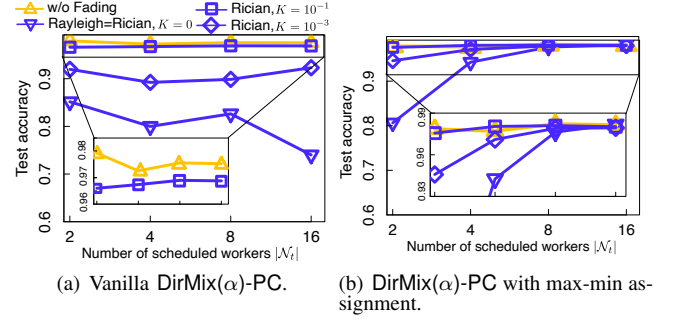


Fig. 4. Impact of fading effect on test accuracy in DirMix(α)-PC with $\alpha = 5$. (a) DirMix(α)-PC with a random assignment of q_t . (b) DirMix(α)-PC with an assignment of q_t maximizing $\min_{i \in \mathcal{N}_t} |h_{i,t}|^2 / q_{i,t}^2$.

for a larger K introduces more intensive fluctuation in the channel coefficient $|h_{i,t}|$, where the workers are required to set a lower scaling factor β . Particularly, under the Rayleigh fading channel, the channel coefficient $|h_{i,t}|$ can take a near-zero value, which severely degrades β . This yields noisy mixed-up samples in (7) and (8), which leads to poorer model performance.

How to Overcome Fading Effects? By setting β as: $P_{\max} \min_{i \in \mathcal{N}_t} |h_{i,t}|^2 / q_{i,t}^2$, we can compensate for the deep fade of $|h_{i,t}|$. More specifically, by assigning a smaller element q_t into a worker with a deep fade of $|h_{i,t}|$, the scaling factor β could be larger. Based on this, we examined a “max-min assignment”, where the elements in q_t are assigned to workers so that $\min_{i \in \mathcal{N}_t} |h_{i,t}|^2 / q_{i,t}^2$ could be maximized. Fig. 4(b) validates the effectiveness of this max-min assignment, demonstrating that the test accuracy under fading is approximately identical to that of without fading.

B. Impact of Privacy Requirements

In Tables I and II, we respectively show the test accuracy and total energy consumption for various DP constraints ϵ and dispersion parameter α in the Dirichlet distribution. We calculate the total energy consumption by: $\tau \sum_{t=1}^T \sum_{i \in \mathcal{N}_t} P_{i,t}$, where $\tau = 1$ ms is the length of each time slot. Note that in this evaluation, we set $n = 2$ and do not consider a small-scale fading effect, i.e., $|g_{i,t}| = 1$.

Impact of Dispersion Parameter α on Accuracy. From Table I, we obtain the following two insights: First, in a lower DP level, i.e., under tight privacy constraints, a larger α results in a better test accuracy for both IRIS and MNIST datasets. This coincides with Remark 2, which stated that we should set α with a large value in severe privacy constraints. Second, in a higher DP level ϵ , i.e., under looser privacy constraints, a smaller α results in better test accuracy. This also coincides with Remark 2.

Note that for dataset S, i.e., Iris dataset, we achieve 92% accuracy even for a stringent privacy level $\epsilon = 5$. Meanwhile, for dataset L, i.e., MNIST dataset, the test accuracy is around 10% for $\epsilon = 10$ for any dispersion parameters, which is the worst accuracy for $d_Y = 10$. This is because from (15), scaling factor β monotonously decreases as d_X and d_Y decrease, which results in a higher noise variance. To solve this, effective data compression methods to reduce the data dimension is required, which is deferred to future work.

Impact of Dispersion Parameter α on Energy Footprints. From Table II, we can see that a smaller α yields a smaller total energy consumption. This is because a smaller α leads to a concentration of the mass in q_t in one entry, which results in

TABLE I
TEST ACCURACY VS. DP LEVEL IN DP(ϵ)-DIRMIX(α)-PC

DP Level ($\delta = 0.01$)	Test Accuracy for Dataset S, i.e., Iris (%)					
	$ \mathcal{N}_t = 4$			$ \mathcal{N}_t = 8$		
	$\alpha = 1$	10	10^5	1	10	10^5
$\epsilon = 5$	74.0	70.4	87.6	68.0	71.6	92.0
$\epsilon = 10$	71.2	82.0	93.6	71.6	81.5	90.8
$\epsilon = 10^2$	83.6	83.6	92.7	78.3	88.7	90.4
$\epsilon = 10^4$	95.1	76.8	80.0	91.1	84.4	76.0
max. power	100.0	95.5	91.5	98.7	91.9	89.5

DP Level ($\delta = 0.01$)	Test Accuracy for Dataset L, i.e., MNIST (%)					
	$ \mathcal{N}_t = 64$			$ \mathcal{N}_t = 128$		
	$\alpha = 10^2$	10^3	10^7	10^2	10^3	10^7
$\epsilon = 10$	11.3	11.3	11.3	11.3	11.3	11.3
$\epsilon = 10^2$	11.3	11.3	58.5	11.3	11.3	68.7
$\epsilon = 10^5$	11.3	64.1	76.4	12.6	48.2	80.6
$\epsilon = 10^8$	90.2	88.0	89.9	88.2	88.0	87.3
max. power	90.25	89.6	90.21	87.9	88.4	87.3

a larger $\max_i q_{i,t}$ and results in a lower β and transmit power as shown in (15). Hence, according to Remark 2, for looser DP constraint, i.e., larger ϵ , we should set α as a lower value, thereby achieving both lower energy consumption and higher test accuracy.

V. CONCLUSION

To perform privacy-preserving ML model training without any on-device training, we proposed AirMixML, where the training data is superposed over wireless channels. AirMixML addresses the issue of how to control transmit power with the goal of enhancing the training performance while guaranteeing DP constraints. To answer this, we proposed DP(ϵ)-DirMix(α)-PC and optimized both the dispersion parameter α and scaling factor β . An interesting direction is to consider mixup in digital over-the-air computation, where quantization introduces another desired nuisance to ensure privacy.

APPENDIX

Equations (10)–(12) are the same as those in the Theorem 1 in [8], and the difference is in the expression of $\epsilon_t(\gamma)$ in (13). Generally, the left-hand-side of (13) is defined as an upper limit of Rényi divergence [14] between the outcomes of the mechanism having adjacent inputs, i.e., $\mathcal{M}(d)$ and $\mathcal{M}(d')$. Hence, we prove that this upper limit of Rényi divergence in the mechanisms (7) and (8) corresponds to (13) focusing on the difference of the mixup ratios from [8], where we consider general mixup ratios $q_t = (q_{i,t})_{i \in \mathcal{N}_t}$ whereas [8] considers equal mixup ratios $q_t = (1/|\mathcal{N}_t|, \dots, 1/|\mathcal{N}_t|)$.

The mechanisms in (7) and (8) consist of two “Gaussian mechanisms”, which refer to the procedure that takes sums of input data and adds Gaussian noise. According to Eq. (10) in [8], the upper limit $\epsilon_t(\gamma)$ is given by:

$$\epsilon_t(\gamma) = \frac{\gamma}{2\sigma_1^2} \sup_{d_1, d'_1} \|\mu_{d_1} - \mu_{d'_1}\|^2 + \frac{\gamma}{2\sigma_2^2} \sup_{d_2, d'_2} \|\mu_{d_2} - \mu_{d'_2}\|^2,$$

where $\mu_{(\cdot)}$ is the sum of the input data and σ_1^2 and σ_2^2 is the noise variance of the first and second mechanism, respectively. The proof is completed by deriving $\epsilon_t(\gamma)$ while retaining the generality of mixup ratios q_t , which is different from Eq. (10) in [8]. We obtain:

$$\begin{aligned} \epsilon_t(\gamma) &= \frac{\gamma}{2(\sigma_n^2/2\beta)} \left(\sup_{j \in \mathcal{N}_t} \left\| \sum_{i \in \mathcal{N}_t} q_{i,t} \mathbf{i}_i - \sum_{i \in \mathcal{N}_t \setminus \{j\}} q_{i,t} \mathbf{i}_i \right\|^2 \right. \\ &\quad \left. + \sup_{j \in \mathcal{N}_t} \left\| \sum_{i \in \mathcal{N}_t} q_{i,t} \mathbf{l}_i - \sum_{i \in \mathcal{N}_t \setminus \{j\}} q_{i,t} \mathbf{l}_i \right\|^2 \right) \\ &= \frac{\gamma}{2(\sigma_n^2/2\beta)} \left(\sup_{j \in \mathcal{N}_t} \|q_{j,t} \mathbf{i}_j\|^2 + \sup_{j \in \mathcal{N}_t} \|q_{j,t} \mathbf{l}_j\|^2 \right) = \frac{\gamma}{2} \frac{(d_X + d_Y)}{(\sigma_n^2/2\beta)} \max_{i \in \mathcal{N}_t} q_{i,t}^2. \end{aligned}$$

TABLE II
ENERGY FOOTPRINTS VS. DP LEVEL IN DP(ϵ)-DIRMIX(α)-PC

DP Level ($\delta = 0.01$)	Energy Footprints for Dataset S, i.e., Iris (μJ)					
	$ \mathcal{N}_t = 4$			$ \mathcal{N}_t = 8$		
	$\alpha = 1$	10	10^5	1	10	10^5
$\epsilon = 5$	0.0912	0.137	0.291	0.0615	0.105	0.375
$\epsilon = 10$	0.152	0.230	0.487	0.125	0.215	0.765
$\epsilon = 10^2$	0.220	0.333	0.705	0.196	0.338	1.201
$\epsilon = 10^4$	2.61	3.94	8.35	2.70	4.64	16.4
max. power	0.246 J	0.476 J	0.348 J	0.257 J	0.411 J	0.817 J

DP Level ($\delta = 0.01$)	Energy Footprints for Dataset L, i.e., MNIST (μJ)					
	$ \mathcal{N}_t = 64$			$ \mathcal{N}_t = 128$		
	$\alpha = 10^2$	10^3	10^7	10^2	10^3	10^7
$\epsilon = 10$	0.295	0.547	0.775	0.288	0.687	1.16
$\epsilon = 10^2$	0.419	0.777	1.09	0.448	1.06	1.81
$\epsilon = 10^5$	1.04	1.93	2.73	1.25	2.98	5.07
$\epsilon = 10^8$	363	673	952	469	1119	1902
max. power	319 J	460 J	498 J	428 J	790 J	950 J

Note that if we consider an equal mixup, i.e., $q_{i,t} = 1/|\mathcal{N}_t|$ as done in [8], the above expression becomes identical to Eq. (10) in [8].

REFERENCES

- [1] G. Zhu, Y. Wang, and K. Huang, “Broadband analog aggregation for low-latency federated edge learning,” *IEEE Trans. Wireless Commun.*, vol. 19, no. 1, pp. 491–506, Oct. 2019.
- [2] Y. Koda, K. Yamamoto, T. Nishio, and M. Morikura, “Differentially private aircomp federated learning with power adaptation harnessing receiver noise,” in *Proc. IEEE GLOBECOM 2020*, Taipei, Taiwan, Dec. 2020, pp. 1–6.
- [3] D. Liu and O. Simeone, “Privacy for free: Wireless federated learning via uncoded transmission with adaptive power control,” *IEEE J. Sel. Areas Commun.*, vol. 39, no. 1, pp. 170–185, Jan. 2021.
- [4] A. Elgabli, J. Park, C. B. Issaid, and M. Bennis, “Harnessing wireless channels for scalable and privacy-preserving federated learning,” *arXiv preprint arXiv:2007.01790*, Nov. 2020.
- [5] M. Fredrikson, S. Jha, and T. Ristenpart, “Model inversion attacks that exploit confidence information and basic countermeasures,” in *Proc. ACM CCS 2015*, Denver, Colorado, USA, Oct. 2015, pp. 1322–1333. [Online]. Available: <https://doi.org/10.1145/2810103.2813677>
- [6] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, “mixup: Beyond empirical risk minimization,” *arXiv preprint arXiv:1710.09412*, Apr. 2018.
- [7] C. Dwork, A. Roth *et al.*, “The algorithmic foundations of differential privacy,” *Found. Trends Theor. Comput. Sci.*, vol. 9, no. 3–4, pp. 211–407, Aug. 2014.
- [8] K. Lee, H. Kim, K. Lee, C. Suh, and K. Ramchandran, “Synthesizing differentially private datasets using random mixing,” in *Proc. IEEE ISIT 2019*, Paris, France, Jul. 2019, pp. 542–546.
- [9] E. Borgnia, J. Geiping, V. Cherepanova, L. Fowl, A. Gupta, A. Ghiasi, F. Huang, M. Goldblum, and T. Goldstein, “DP-instahide: Provably defusing poisoning and backdoor attacks with differentially private data augmentations,” *arXiv preprint arXiv:2103.02079*, Apr. 2021.
- [10] H. B. McMahan, E. Moore, D. Ramage, S. Hampson *et al.*, “Communication-efficient learning of deep networks from decentralized data,” in *Proc. AISTATS 2017*, Fort Lauderdale, FL, USA, Apr. 2017, pp. 1–11.
- [11] H. B. McMahan *et al.*, “Advances and open problems in federated learning,” *Found. Trends Mach. Learn.*, vol. 14, no. 1, Jul. 2021.
- [12] K. Yang, T. Jiang, Y. Shi, and Z. Ding, “Federated learning via over-the-air computation,” *IEEE Trans. Wireless Commun.*, vol. 19, no. 3, pp. 2022–2035, Mar. 2020.
- [13] T. Sery, N. Shlezinger, K. Cohen, and Y. C. Eldar, “Over-the-air federated learning from heterogeneous data,” *arXiv preprint arXiv:2009.12787*, Oct. 2020.
- [14] I. Mironov, “Rényi differential privacy,” in *Proc. IEEE CSF 2017*, Santa Barbara, CA, USA, Jun. 2017, pp. 263–275.
- [15] “The Iris dataset.” [Online]. Available: https://scikit-learn.org/stable/auto_examples/datasets/plot_iris_dataset.html
- [16] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [17] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, “On the importance of initialization and momentum in deep learning,” in *Proc. ICML 2013*, Atlanta, GA, USA, Jun. 2013, pp. 1139–1147.