Machine Unlearning for Equalizer Classifiers in 6G Wireless Communication Systems

Quynh Tu Ngo, *Senior Member, IEEE,* Ying He, *Senior Member, IEEE,* Hasini Abeywickrama, Guangsheng Yu, Eryk Dutkiewicz, *Senior Member, IEEE*, Bathiya Senanayake, and Manik Attygalle, *Senior Member, IEEE*

***Abstract*—This letter proposes a new direction for wireless signal processing by applying machine unlearning (MU) to neural-network-based equalizers in 6G receivers. MU enables a trained equalizer to mitigate the influence of corrupted training samples without full model retraining. A new neural equalizer is designed to better capture nonlinear radio-frequency impairments, inter-symbol interference, and channel-memory effects characteristic of 6G systems. We also develop MU-enabled equalization schemes and evaluate their performance under various data corruption scenarios. Simulation results show that MU effectively restores classification accuracy while reducing computational cost, making it a practical and efficient approach for preserving equalizer robustness in dynamic and imperfect 6G environments.**

***Index Terms*—Equalizer classifier, 6G, Machine Unlearning, Distilling, Fine Tuning, Retraining, SISA**

I. Introduction

Future 6G receivers are expected to rely extensively on neural-network-based equalizers to cope with increasingly complex channel conditions and non-linear impairments. However, the performance of these data-driven equalizers is highly sensitive to corrupted, outdated, or low-quality training samples, conditions that naturally arise in dynamic wireless environments. Once corrupted data enter the training pipeline, the resulting model degradation typically requires full retraining, leading to high computational cost, long latency, and energy inefficiency. These challenges motivate the need for new mechanisms that can efficiently remove the impact of flawed data without rebuilding the model from scratch. This letter explores such a direction by introducing machine unlearning (MU) to neural equalizers.

MU aims to selectively remove or suppress the influence of specific training samples while preserving the remainder of the learned model. MU has recently attracted attention in machine learning due to its benefits in privacy preservation, model maintenance, and robustness against corrupted data [1]. In wireless communication research, MU has been explored in several emerging applications, including interference cancellation through the cleansing of corrupted latent representations without retraining [2], zero-shot MU for privacy-preserving dynamic spectrum access in IoT networks [3], and adversarial MU approaches that eliminate backdoors in beam-selection models while preserving normal performance [4]. These studies demonstrate the potential of MU for reducing retraining overhead and maintaining model reliability, but its application remains limited to high-level learning tasks and has not yet been extended to physical-layer signal processing.

This letter presents the first application of MU to neural-network-based equalizers. Unlike existing SISA-based MU studies in wireless communications, which predominantly focus on learning tasks such as classification and regression (e.g., modulation recognition or channel estimation) [5] rather than sequence-based signal detection problems, we tailor MU to neural equalization, where strong temporal correlation, channel memory, and inter-symbol interference fundamentally couple the training samples. We develop MU-enabled equalization schemes that efficiently mitigate the impact of corrupted training data and restore equalizer performance with only partial the computational cost of full retraining. This work aims to establish practical and empirically validated MU baselines for neural equalizers under realistic 6G channel conditions, rather than deriving new theoretical forgetting guarantees. The contributions of this work are fourfold: (i) we formulate, for the first time, a SISA-based MU framework specifically for neural equalizers in 6G receiver architectures, and design two task-aware unlearning strategies, namely SISA–Distillation–Based MU and SISA–Fine–Tune–Based MU, that account for the sequential and structured nature of equalizer training data; (ii) we propose a new neural equalizer, namely RES-CNN-AMP, an enhanced residual shrinkage convolutional network architecture capable of capturing nonlinear radio-frequency impairments, inter-symbol interference, and channel-memory effects anticipated in 6G communication systems; (iii) we evaluate MU performance under several data corruption scenarios representative of practical 6G wireless environments; and (iv) we demonstrate through simulations that the proposed MU schemes achieve performance comparable to full retraining while reducing computational complexity, and that the proposed equalizer consistently outperforms the conventional baseline.

II. Neural Equalization Framework

Consider a point-to-point orthogonal frequency-division multiplexing (OFDM)-based link incorporating a cyclic prefix and Doppler effects, representative of 5G New Radio

Q. T. Ngo, Y. He, H. Abeywickrama, G. Yu and E. Dutkiewicz are with the Faculty of Engineering and IT, University of Technology Sydney, Australia. Email: {quynhtu.ngo, ying.he, hasini.abeywickrama, guangsheng.yu, eryk.dutkiewicz}@uts.edu.au

B. Senanayake and M. Attygalle are with Defence Science and Technology Group, Australia. Email: {bathiya.senanayake, manik.attygalle}@defence.gov.au

The Commonwealth of Australia (represented by the Department of Defence) supports this research through a Defence Science Partnerships agreement.

waveforms and projected 6G channel dynamics. The received discrete-time signal can be expressed as

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n}, \tag{1}$$

where $\mathbf{x} \in \mathbb{C}^N$ contains N transmitted symbols drawn from an M-ary quadrature amplitude modulation (M-QAM) constellation, $\mathbf{H} \in \mathbb{C}^{N\times N}$ is the channel matrix, and $\mathbf{n} \sim \mathcal{CN}(0, \sigma^2\mathbf{I})$ represents additive white Gaussian noise (AWGN) with zero mean and variance σ^2. Ideal channel estimates $\hat{\mathbf{H}}$ are assumed to isolate equalizer performance. The equalizer recovers $\mathbf{x}$ from $(\mathbf{y}, \hat{\mathbf{H}})$ under channel distortion and noise:

$$\hat{\mathbf{x}} = \mathcal{F}_\theta\left(\mathbf{y}, \hat{\mathbf{H}}\right), \tag{2}$$

where $\hat{\mathbf{x}}$ denotes the estimated symbol vector and $\mathcal{F}_\theta(\cdot)$ denotes the equalization network with parameters θ.

A. Proposed RES-CNN-AMP Equalizer

To model 6G receiver architectures, we propose RES-CNN-AMP, an enhanced version of residual shrinkage convolutional neural network (RES-CNN), which is commonly applied at the receiver side in 6G signal processing. Unlike conventional RES-CNN equalizers that directly process per-symbol in-phase (I) and quadrature (Q) samples, RES-CNN-AMP incorporates amplitude information for each symbol to better capture nonlinear radio-frequency effects anticipated in 6G communication systems. In addition, RES-CNN-AMP internally reconstructs a temporal context window around each received symbol. This window-based processing enables the network to capture inter-symbol interference and channel memory effects without requiring any changes to the dataset or preprocessing pipeline.

For the n-th received symbol with IQ components (I_n, Q_n), RES-CNN-AMP computes an additional amplitude channel,

$$A_n = \sqrt{I_n^2 + Q_n^2 + \epsilon}, \quad n = 1, 2, \ldots, N, \tag{3}$$

where A_n represents the instantaneous symbol amplitude and $0 < \epsilon \ll 1$ is a constant to ensure numerical instability. The model then constructs a temporal window of length $2C + 1$ centered at symbol n, where C is the context size, producing a structured input tensor $\mathbf{X}_n \in \mathbb{R}^{3\times(2C+1)}$ containing the stacked I, Q, and amplitude sequences in the neighborhood of symbol n.

The RES-CNN-AMP architecture consists of three stages: (1) *Initial Convolutional Feature Extraction:* a 1D convolution with batch normalization and ReLU activation transforms the three-channel input into a higher-dimensional feature space, extracting low-level temporal features across the window; (2) *Residual Dilated Convolutional Blocks:* several residual blocks, each containing two convolutional layers with batch normalization, are applied sequentially. The blocks incorporate increasing dilation factors to expand the receptive field and capture long-range dependencies within the symbol window. Skip connections stabilize training and mitigate vanishing gradients; and (3) *Global Feature Aggregation and Classification:* the resulting feature map is aggregated using adaptive average pooling and passed through fully connected layers to produce logits over the M possible constellation symbols. This enables end-to-end symbol classification for arbitrary M-QAM schemes.

B. Training Objective

The RES-CNN-AMP equalizer is trained as a multi-class classifier over the M-QAM constellation. Let $\mathcal{S} = \{s_1, s_2, \ldots, s_M\} \subset \mathbb{C}$ denote the constellation set, and let $t_n \in \{1, \ldots, M\}$ be the ground-truth symbol index corresponding to the n-th training instance. For each input window $\mathbf{X}_n$, the equalizer outputs logits $\{z_{n,1}, \ldots, z_{n,M}\}$, which are converted to probabilities using a softmax activation.

The model parameters are optimized by minimizing the cross-entropy loss:

$$\mathcal{L}_{\text{CE}} = -\sum_{i=1}^{N} \log\left(\frac{\exp(z_{n,t_n})}{\sum_{m=1}^{M}\exp(z_{n,m})}\right), \tag{4}$$

where $z_{n,m}$ is the logit corresponding to class m for the n-th sample, and t_n is the index of the ground-truth symbol $x_n = s_{t_n}$.

At inference time, the predicted symbol $\tilde{x}_n \in \mathcal{S}$ is obtained by selecting the class with the highest logit value:

$$\tilde{x}_n = s_{m^*}, \quad m^* = \underset{m\in\{1,\ldots,M\}}{\operatorname{argmax}} z_{n,m}. \tag{5}$$

C. Computational Complexity

The computational complexity of the proposed RES-CNN-AMP equalizer is higher than that of the baseline RES-CNN due to the inclusion of an additional amplitude channel and the use of a temporal context window of length $2C+1$. In the baseline RES-CNN, each symbol is processed independently using a two channel input (I_n, Q_n), yielding a per-symbol complexity of

$$\mathcal{O}\left(FK + B_R F^2 K + FM\right), \tag{6}$$

where F denotes the number of convolutional feature maps, K is the kernel size, and B_R is the number of residual blocks. In contrast, RES-CNN-AMP forms a structured input tensor $\mathbf{X}_n \in \mathbb{R}^{3\times(2C+1)}$, introducing an $\mathcal{O}(C)$ preprocessing overhead and expanding the spacial dimension of all convolutional operations. The initial convolution thus requires $\mathcal{O}(F \cdot 3K(2C + 1))$, while each residual block incurs $\mathcal{O}(F^2K(2C+1))$ leading to a total per-symbol complexity of

$$\mathcal{O}\left(C + 3FK\,(2C + 1) + B_R F^2 K\,(2C + 1) + FM\right). \tag{7}$$

Since the convolutional blocks dominate the total cost, RES-CNN-AMP increases complexity by approximately a factor of $2C + 1$ relative to RES-CNN.

Note that this analysis focuses on theoretical operation counts and does not account for memory access patterns or hardware-specific optimizations, which may affect actual runtime on different platforms.

III. Machine Unlearning in Neural Equalization

We propose two MU strategies for the RES-CNN-AMP equalizer using the sharded, isolated, sliced, and aggregated (SISA) training framework [6]. SISA allows localized updates, as only shards containing affected samples are retrained, while others remain unchanged. Different from prior SISA-based MU studies in wireless learning that consider independent and identically distributed samples [5], the proposed strategies address the structured, temporally correlated nature of equalization data, and integrate unlearning directly into the receiver signal processing chain. Let $\mathcal{D} = \{(\mathbf{y}_n, \hat{\mathbf{H}}_n, t_n)\}_{n=1}^{N}$ denote the full training dataset being partitioned into S shards $\{\mathcal{D}_s\}_{s=1}^{S}$, each further divided into mini-batches for training. Let $\mathcal{D}_{forget} \subset \mathcal{D}$ denote the samples to be erased, and $\mathcal{D}_{retain} = \mathcal{D} \setminus \mathcal{D}_{forget}$ the retained dataset.

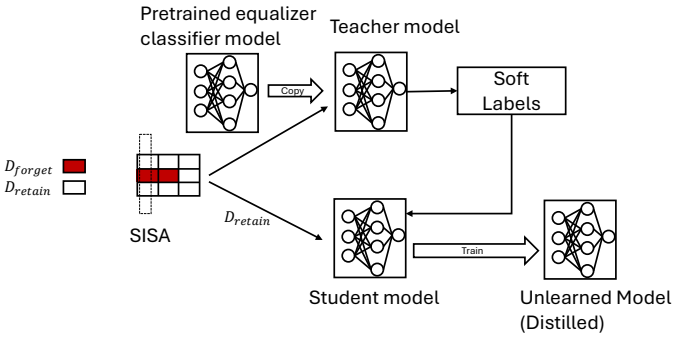

(a) SISA-distillation-based machine unlearning framework

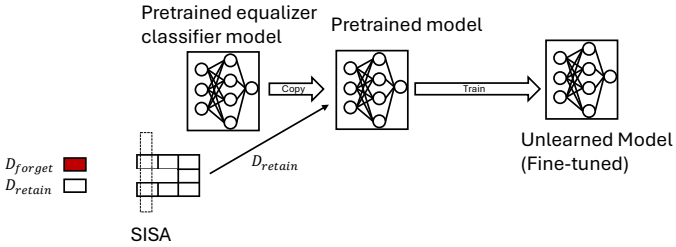

(b) SISA-fine-tune-based machine unlearning framework

Fig. 1: Proposed MU schemes for RES-CNN-AMP equalizer.

A. SISA-Distillation-Based Machine Unlearning (S-DtMU)

In S-DtMU (Fig. 1a), a student network $\mathcal{F}_S$ learns to mimic a pretrained teacher network $\mathcal{F}_T$ trained on $\mathcal{D}$, while ignoring $\mathcal{D}_{forget}$. For each mini-batch $(y, \hat{\mathbf{H}}, t) \sim \mathcal{D}_{retain}$, the student network is trained using a hybrid loss function:

$$\mathcal{L}_{\text{Dt}} = (1-\lambda)\,\mathcal{L}_{\text{CE}}\left(\mathcal{F}_S\left(\mathbf{y}, \hat{\mathbf{H}}\right), t\right) + \lambda T^2 \mathcal{L}_{\text{KD}}\left(\mathcal{F}_S\left(\mathbf{y}, \hat{\mathbf{H}}\right), \mathcal{F}_T\left(\mathbf{y}, \hat{\mathbf{H}}\right)\right), \tag{8}$$

where $\mathcal{L}_{\text{KD}}(z_S, z_T) = \text{KL}\big(\sigma(z_S/T) \,\|\, \sigma(z_T/T)\big)$ is the Kullback–Leibler divergence between the softened logits z_S and z_T with $\sigma(\cdot)$ being the softmax function, $T > 0$ is the distillation temperature, and $\lambda \in [0, 1]$ balances the two objectives.

Within SISA, shards s with $s \cap \mathcal{D}_{forget} = \varnothing$ retain the teacher parameters directly in the student, avoiding redundant retraining. Only shards containing retained samples undergo gradient updates, enabling parallelizable and localized unlearning. S-DtMU procedure is described in Algorithm 1.

Algorithm 1 SISA-Distillation-Based Machine Unlearning

Require: Teacher model $\mathcal{F}_T$ (trained on $\mathcal{D}$), retained set $\mathcal{D}_{\text{retain}}$, temperature T, trade-off λ, epochs E, learning rate η, shards S
Ensure: Student model $\mathcal{F}_S$ unlearned from $\mathcal{D}_{\text{forget}}$

```
Initialize F_S with F_T parameters
for each shard s ∈ S do
    if s ∩ D_forget = ∅ then
        Copy parameters of F_T → F_S; skip shard training
    else
        B ← mini-batches of s ∩ D_retain
        for epoch = 1, ..., E do
            for (y, Ĥ, t) ∈ B do
                z_T ← F_T(y, Ĥ); z_S ← F_S(y, Ĥ)
                Compute L_KD, L_CE, and L_Dt
                Update: θ_S ← θ_S − η∇_{θ_S} L_Dt
            end for
        end for
    end if
end for
return F_S
```

B. SISA-Fine-Tune-Based Machine Unlearning (S-FtMU)

In S-FtMU (Fig. 1b), the pretrained equalizer $\mathcal{F}_{\theta_{\text{old}}}$ is updated directly on $\mathcal{D}_{\text{retain}}$ while suppressing residual knowledge of $\mathcal{D}_{\text{forget}}$. The objective reduces to the standard cross-entropy loss:

$$\mathcal{L}_{\text{Ft}} = \mathcal{L}_{\text{CE}}\left(\mathcal{F}_\theta\left(\mathbf{y}, \hat{\mathbf{H}}\right), t\right), \tag{9}$$

where θ denotes the parameters of the fine-tuned network. Within the SISA framework, shards that do not contain samples from $\mathcal{D}_{\text{forget}}$ are left unchanged, while only relevant shards are fine-tuned on the retained data. Fine-tuning can optionally reinitialize classification heads to accelerate forgetting of $\mathcal{D}_{\text{forget}}$, as detailed in Algorithm 2.

Remark 1. *The assumption of a priori known $\mathcal{D}_{forget}$ allows isolation of machine unlearning evaluation. In practice, $\mathcal{D}_{forget}$ can be inferred via signal quality monitoring or anomaly detection on IQ samples, which may produce false positives or negatives. Evaluating MU under such imperfect detection is left for future work.*

IV. Numerical Results

A. Simulation Setup

We generate two controlled 6G-oriented orthogonal frequency-division multiplexing (OFDM) datasets containing noisy, wideband, per-symbol IQ samples and their corresponding QAM labels. The datasets follow a 1024-subcarrier OFDM waveform with a 7% cyclic prefix, 64-QAM and 256-QAM modulation, and a sampling rate of 122.88 MHz, consistent with high-mobility 6G numerologies [7]. A three-tap Shadowed Rician fading channel is applied, together with a maximum Doppler shift of 1,500 Hz to emulate LEO satellite and airborne mobility conditions [8]. AWGN is added at an SNR of 20 dB. From this waveform, the per-symbol received

Algorithm 2 SISA-Fine-Tune-Based Machine Unlearning

Require: Pretrained model $\mathcal{F}_P$ with parameters θ_{old} (trained on $\mathcal{D}$), retained set $\mathcal{D}_{\text{retain}}$, epochs E, shards S
Ensure: Fine-tuned model $\mathcal{F}_{P'}$ with parameters θ unlearned from $\mathcal{D}_{\text{forget}}$

```
Initialize θ ← θ_old; optionally reinitialize classifier head
for each shard s ∈ S do
    if s ∩ D_forget = ∅ then
        Skip shard
    else
        B ← mini-batches of s ∩ D_retain
        for epoch = 1, ..., E do
            for (y, Ĥ, t) ∈ B do
                z ← F_P'(y, Ĥ)
                Compute L_CE and L_Ft
                Update: θ ← θ − η∇_θ L_Ft
            end for
        end for
    end if
end for
return F_P'
```

constellation points are extracted after OFDM demodulation, producing IQ samples and their associated labels.

We define data corruption as feature-level distortion caused by channel and RF-induced distribution shift, while assuming correct symbol labels. Let a clean training sample be represented as

$$\hat{s}_i = f\left(x_i; \mathbf{H}_0\right) + n_i, \tag{10}$$

where x_i denotes the transmitted QAM symbol, $\mathbf{H}_0$ represents the nominal channel distribution (including shadowed Rician fading and Doppler effects) [9], and n_i is the AWGN. Corrupted samples are generated by replacing $\mathbf{H}_0$ with a shifted channel distribution $\mathbf{H}_c$, yielding

$$\tilde{s}_i = f\left(x_i; \mathbf{H}_c\right) + n_i, \tag{11}$$

where $\mathbf{H}_c$ captures increased channel memory, nonlinear RF distortion, or mobility-induced variations. Importantly, symbol labels remain unchanged, making this a feature-level corruption model that reflects realistic wireless impairments rather than label noise or adversarial injection. Similar channel-induced distribution shifts are widely used in wireless ML to model nonstationary propagation and hardware effects in high-mobility systems [1], [2]. The number of impacted classes scales with modulation order, 26 for 64-QAM and 102 for 256-QAM, capturing the fact that nonlinearities and high-mobility effects disproportionately affect subsets of constellation points rather than the entire symbol alphabet. For each dataset, the training set consists of 200,000 labeled symbols, and the testing set has 20,000 symbols. The hyper-parameters of RES-CNN-AMP equalizer are summarized in Table I.

We evaluate the proposed RES-CNN-AMP equalizer against the baseline RES-CNN across multiple MU strategies. Specially, we compare the performance of each model under the following training/unlearning approaches:

- *Pretrain*: conventional training without MU;
- *S-Pretrain*: SISA-based training;
- *S-DtMU*: proposed SISA-distillation-based MU training;
- *S-FtMU*: proposed SISA-fine-tune-based MU training;
- *DtMU*: distillation-based MU training;
- *FtMU*: fine-tune-based MU training;
- and full retraining from scratch.

TABLE I: Hyper-parameters of RES-CNN-AMP equalizer

Hyper-parameter	Value
Window length	31
Conv1D input → output channels	$3 \rightarrow 128$
Kernel size	5
Padding	2
Activation	ReLU
Normalization	BatchNorm1d
No. residual blocks	4
Channels per residual block	128
Fully connected layer 1	$128 \rightarrow 256$
Fully connected layer 2	$256 \rightarrow$ 64/256

TABLE II: Training time (second) comparison on an OFDM 256-QAM dataset with 102 impacted classes and 30% impacted labels per class

	RES-CNN-AMP	RES-CNN
Pretrain	103.71	100.58
DtMU	63.34	57.43
FtMU	50.13	50.21
S-Pretrain	49.27	50.16
S-DtMU	24.35	22.04
S-FtMU	16.54	15.37
Retrain	92.59	90.75

B. Simulation Results

Figs. 2 and 3 present the classification accuracy of each equalizer model under different MU strategies. Three types of accuracy are reported: (i) *overall*, measured across the entire test set; (ii) *impacted*, measured only on classes whose training samples contained imperfections; and (iii) *unimpacted*, measured on classes unaffected by data imperfections.

As expected, RES-CNN-AMP consistently outperforms the baseline RES-CNN across all MU approaches. Pretrained models without MU, trained on the full dataset including imperfect samples, exhibit the lowest accuracy, particularly on the impacted classes. Retraining from scratch using only the clean subset achieves the highest accuracy, serving as an upper bound for MU performance.

Among the MU strategies, SISA-distillation-based MU achieves accuracy nearly identical to full retraining across all categories, with only a marginal drop on the impacted classes due to residual knowledge transfer. This superior performance can be attributed to the distillation process, which provides smoother and more informative gradients during shard-level updates, thereby improving optimization stability and preserving global decision boundaries in the SISA framework. In contrast, SISA-based fine-tuning improves accuracy relative to the baseline but underperforms compared to S-DtMU, as local gradient updates on limited clean data may be insufficient to fully overwrite representations associated with forgotten samples. Finally, SISA without additional adaptation offers

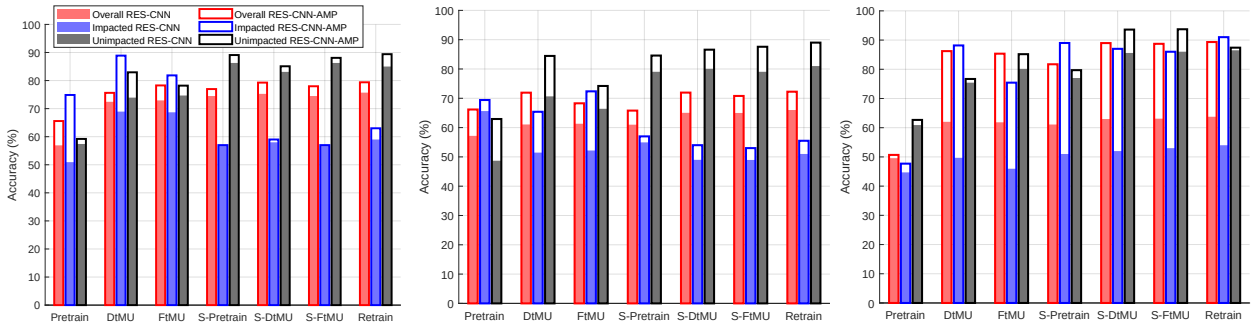

(a) 30% impacted labels per class (b) 50% impacted labels per class (c) 70% impacted labels per class

Fig. 2: Comparison of classification accuracy between RES-CNN and RES-CNN-AMP equalizers under various training strategies on an OFDM 64-QAM dataset with 26 impacted classes.

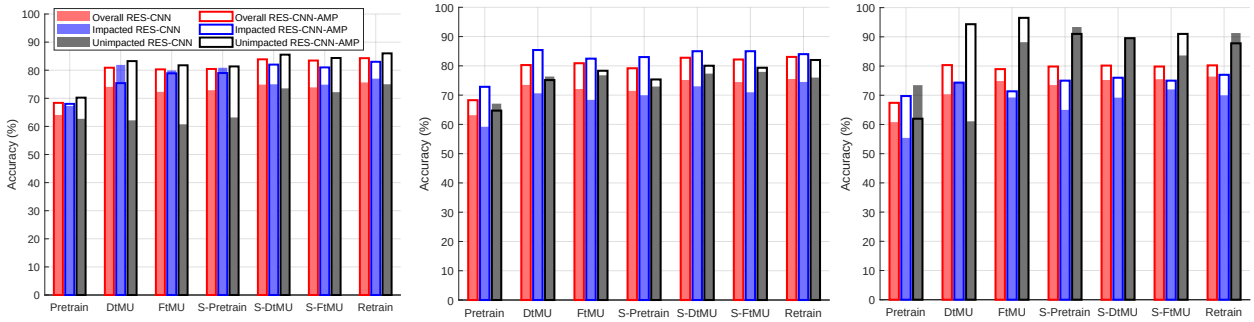

(a) 30% impacted labels per class (b) 50% impacted labels per class (c) 70% impacted labels per class

Fig. 3: Comparison of classification accuracy between RES-CNN and RES-CNN-AMP equalizers under various training strategies on an OFDM 256-QAM dataset with 102 impacted classes.

moderate improvements and a favorable trade-off with reduced training cost, as shown in Table II, making it particularly effective when only a subset of data must be forgotten.

RES-CNN-AMP incurs higher computational complexity than RES-CNN due to amplitude-aware and window-based processing, leading to increased training time and modestly higher inference cost. The resulting accuracy gains therefore reflect a trade-off between computational overhead and improved modeling of nonlinear and memory effects.

Remark 2. *Results assume ideal channel estimates to isolate neural equalizer and MU behavior. In practice, CSI errors may slightly reduce equalization performance and the separability of clean and corrupted samples. Since MU operates on data shards and model slices rather than the channel estimator, its relative unlearning gains and reduced retraining cost are expected to hold under moderate CSI inaccuracies. Robustness evaluation under noisy or outdated CSI is left for future work.*

V. Conclusion

In this letter, we presented the first application of MU to neural equalizers. We developed MU-enabled schemes that effectively mitigate the impact of corrupted training data, restoring equalizer performance at a fraction of the cost of full retraining. We also proposed RES-CNN-AMP, which effectively captures nonlinear RF impairments, inter-symbol interference, and channel-memory effects in 6G receivers. Simulation results under representative 6G data corruption scenarios demonstrate that MU achieves accuracy comparable to retraining while significantly reducing computational cost, and that RES-CNN-AMP consistently outperforms the baseline RES-CNN equalizer.

References

[1] T. T. Nguyen and et. al., "A survey of machine unlearning," *ACM Trans. Intelligent Syst. Technol.*, vol. 16, no. 5, pp. 1–46, 2025.

[2] E. Guven and G. K. Kurt, "Machine unlearning for uplink interference cancellation," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, 2024, pp. 4890–4895.

[3] F. Li and et. al., "Secure dynamic spectrum access in IoT based on machine unlearning," in *Proc. Int. Conf. Artificial Intelligence Inf. Commun.*, 2025, pp. 0131–0136.

[4] Z. Zhang and et. al., "Poison neural network-based mmWave beam selection and detoxification with machine unlearning," *IEEE Trans. Commun.*, vol. 71, no. 2, pp. 877–892, 2022.

[5] S. B. R. Chowdhury and et. al., "Towards scalable exact machine unlearning using parameter-efficient fine-tuning," in *Neurips Safe Generative AI Workshop 2024*.

[6] L. Bourtoule and et. al., "Machine unlearning," in *2021 IEEE Symposium Security Privacy*, 2021, pp. 141–159.

[7] M. Sayed, H. Zakaria, and A. Abdelhady, "Enhancing flexibility and system performance in 6G and beyond," *Digital Commun. Netw.*, vol. 11, no. 4, pp. 975–991, 2025.

[8] Q. T. Ngo and et. al., "A fast fuzzy DRL-based joint beam design and power allocation for multi-beam GEO-LEO coexisting satellite networks," *IEEE Trans. Wireless Commun.*, vol. 24, no. 10, pp. 1558–2248, 2025.

[9] Q. T. Ngo, Y. He, B. Jayawickrama, and et. al., "Quantum RL with Classical Policy Deployment for Resource Allocation in Multi-Beam GEO-LEO Satellite Networks," *IEEE Internet Things J.*, pp. 1–15, 2026.