META-AF ECHO CANCELLATION FOR IMPROVED KEYWORD SPOTTING

Jonah Casebeer   *Junkai Wu*   *Paris Smaragdis*

University of Illinois at Urbana-Champaign

ABSTRACT

Adaptive filters (AFs) are vital for enhancing the performance of downstream tasks, such as speech recognition, sound event detection, and keyword spotting. However, traditional AF design prioritizes isolated signal-level objectives, often overlooking downstream task performance. This can lead to suboptimal performance. Recent research has leveraged meta-learning to automatically learn AF update rules from data, alleviating the need for manual tuning when using simple signal-level objectives. This paper improves the Meta-AF [1] framework by expanding it to support end-to-end training for arbitrary downstream tasks. We focus on classification tasks, where we introduce a novel training methodology that harnesses self-supervision and classifier feedback. We evaluate our approach on the combined task of acoustic echo cancellation and keyword spotting. Our findings demonstrate consistent performance improvements with both pre-trained and joint-trained keyword spotting models across synthetic and real playback. Notably, these improvements come without requiring additional tuning, increased inference-time complexity, or reliance on oracle signal-level training data.

Index Terms— adaptive filtering, keyword spotting, learning to learn, meta-learning, acoustic echo cancellation

1. INTRODUCTION

Adaptive filters (AFs) are essential for many smart systems, including automatic speech recognition, sound event detection, and keyword spotting [2]. In these applications, AFs serve as key preprocessing steps and are designed to enhance the performance of their downstream components by removing noise, reverberation, and other distortions from their input signals. Traditionally, AFs have been hand-tuned for optimizing signal-level objectives, such as minimizing mean-squared error, using techniques like recursive least squares [3, 4, 5]. However, recent advancements in deep learning have introduced methods for controlling pre-built AF update rules [6, 7, 8], or for learning entirely new update rules from scratch [9, 1]. Existing techniques often do not take into account the downstream task and may not lead to the best performance in practice.

Several pioneering works demonstrated that incorporating downstream task knowledge into AF control, through manual unification schemes, could enhance performance [10, 11, 12]. Such improvements necessitate a complex, task/model-dependent derivation process, as well as feedback from the downstream model to the AF at test time. This line of work inspired many DNN-based approaches. These DNN approaches can be broadly grouped into two categories. The first explicitly decouples preprocessing and downstream operations, treating them as sequential components within a larger DNN pipeline [13, 14, 2, 15]. This is highly modular but typically requires more complex training schemes and engineering overhead [16]. The
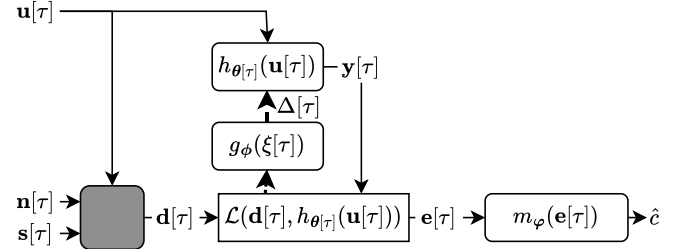
Fig. 1: The echo canceller $h_{\boldsymbol{\theta}[\tau]}$, controlled by optimizer $g_{\boldsymbol{\phi}}$, cancels own-playback and passes its output to keyword classifier $m_{\boldsymbol{\varphi}}$. The optimizer is trained end-to-end with the classifier, improving performance without the need for additional tuning.

second takes an end-to-end strategy, trading modularity for simpler design and training schemes [17, 18, 19, 20].

In this paper, we propose a novel framework that bridges these two DNN-based approaches, allowing for the incorporation of downstream classification task knowledge into AFs without task-dependent design or test-time feedback. We propose a modified training scheme for the Meta-AF framework [1], which we call classification-trained Meta-AF (CT-Meta-AF). CT-Meta-AF is unique in that it does not require oracle signal-level supervision such as impulse responses or clean speech. Instead, it leverages a combination of signal-level self-supervision and automatic classifier feedback. This eliminates the need for manual task/model dependent tuning and is directly compatible with both off-the-shelf and end-to-end trained downstream models.

To evaluate our framework, we apply it to two fundamental tasks for smart devices: acoustic echo cancellation (AEC) and keyword spotting (KWS). AEC is a critical component of nearly all devices with own-playback, and KWS is essential for enabling hands-free operation of smart devices. We construct a challenging joint AEC and KWS dataset by combining two existing datasets [21, 22], and benchmark CT-Meta-AFs against Meta-AF [23] and a Kalman Filter [24], two state-of-the-art approaches for signal-level objectives. Results show that CT-Meta-AF achieves consistent performance gains across both synthetic and real playback scenarios, multiple playback lengths, various task/model configurations, and both pre-trained and joint-trained keyword spotting models.

The contributions of this work are: 1) A versatile framework for seamlessly integrating downstream task knowledge into AF update rules, applicable to various AFs and tasks, 2) Task-specific insights for AEC and KWS, yielding models that surpass previous approaches, 3) A comprehensive empirical exploration of incorporating downstream task knowledge across different classifier models, employing various training schemes, and evaluating on both real-world and synthetic data. We also release all code and weights[1].

2. META-AF BACKGROUND

2.1. Adaptive Filters

In this work, an AF is a time-varying filtering procedure with parameters that are controlled by an online optimization rule in order to solve some objective function. We denote the filter as $h_{\boldsymbol{\theta}[\tau]}$, parameterized by $\boldsymbol{\theta}[\tau]$, and updated using an additive rule,

$$\boldsymbol{\theta}[\tau+1] = \boldsymbol{\theta}[\tau] + \boldsymbol{\Delta}[\tau]. \tag{1}$$

Time-varying updates $\boldsymbol{\Delta}[\tau]$ govern filter adaptation, and are produced by an update rule $g_{\boldsymbol{\phi}}$. Conventional AFs are designed to produce updates which solve a tractable AF loss $\mathcal{L}(\cdot)$, which is typically a function of the AF output and some readily available mixture.

To perform our chosen AF task of AEC, we fit a filter to to remove echo caused by own-playback. The time-domain signal model is $\underline{\mathbf{d}}[t] = \underline{\mathbf{u}}[t] * \underline{\mathbf{w}} + \underline{\mathbf{n}}[t] + \underline{\mathbf{s}}[t]$, where $\underline{\mathbf{n}}$ is noise, $\underline{\mathbf{s}}$ is a keyword, and $\underline{\mathbf{w}}$ is the transfer function applied to own playback $\underline{\mathbf{u}}$. The AF attempts to mimic $\underline{\mathbf{w}}$ via a multi-frame frequency-domain filter applied using overlap-save [25]. We denote this as $h_{\boldsymbol{\theta}[\tau]}$ with frequency coefficients $\boldsymbol{\theta}[\tau] = \hat{\mathbf{w}}[\tau] \in \mathbb{C}^{K \times B}$ where K is frequencies, B is the number of blocks, and blocks advance by $K/2$ samples. The output is echo reduced signal $\underline{\mathbf{e}}[\tau] = h_{\boldsymbol{\theta}[\tau]}(\underline{\mathbf{d}}[\tau], \underline{\mathbf{u}}[\tau])$.

2.2. Meta-Adaptive Filters

The Meta-AF [1] approach showed that AF control rules could be entirely learned with a learning-to-learn strategy [26]. Instead of solving the AF loss $\mathcal{L}(\cdot)$ directly, the procedure solves a meta-loss,

$$\hat{\boldsymbol{\phi}} = \arg\min_{\boldsymbol{\phi}} E_{\mathcal{D}}[\, \mathcal{L}_M(\, g_{\boldsymbol{\phi}}, \mathcal{L}(h_{\boldsymbol{\theta}}, \cdots)\,)\,], \tag{2}$$

where $\mathcal{L}_M(\, g_{\boldsymbol{\phi}}, \mathcal{L}(h_{\boldsymbol{\theta}}, \cdots)\,)$ is the meta-loss that is a function of the AF $h_{\boldsymbol{\theta}[\tau]}$, optimizer neural network $g_{\boldsymbol{\phi}}$, and $E_{\mathcal{D}}$ represents expectation over dataset $\mathcal{D}$. The standard setup solves this via an additive update $\boldsymbol{\theta}[\tau+1] = \boldsymbol{\theta}[\tau] + g_{\boldsymbol{\phi}}(\cdot)$, where $g_{\boldsymbol{\phi}}$ is a complex valued DNN parameterized by $\boldsymbol{\phi}$, applied per frequency k. Typically, $g_{\boldsymbol{\phi}}$ has per-frequency state and inputs,

$$(\boldsymbol{\Delta}_{\mathrm{k}}[\tau], \boldsymbol{\psi}_{\mathrm{k}}[\tau+1]) = g_{\boldsymbol{\phi}}(\boldsymbol{\xi}_{\mathrm{k}}[\tau], \boldsymbol{\psi}_{\mathrm{k}}[\tau]) \tag{3}$$
$$\boldsymbol{\theta}_{\mathrm{k}}[\tau+1] = \boldsymbol{\theta}_{\mathrm{k}}[\tau] + \boldsymbol{\Delta}_{\mathrm{k}}[\tau], \tag{4}$$

where we index across K frequencies using subscript k. The input is a stack of the frequency-domain filter gradient, input, target, error, and output respectively: $\boldsymbol{\xi}_{\mathrm{k}}[\tau] = [\nabla_{\mathrm{k}}[\tau], \mathbf{u}_{\mathrm{k}}[\tau], \mathbf{d}_{\mathrm{k}}[\tau], \mathbf{e}_{\mathrm{k}}[\tau], \mathbf{y}_{\mathrm{k}}[\tau]]$. The outputs are AF update $\boldsymbol{\Delta}_{\mathrm{k}}[\tau]$ and a new internal state $\boldsymbol{\psi}_{\mathrm{k}}[\tau]$. To learn parameters $\boldsymbol{\phi}$, we use backpropagation-through-time and a meta-optimizer (e.g. Adam) over L optimization steps. We set our meta-loss to the self-supervised log-mse meta-loss,

$$\mathcal{L}_M(\bar{\mathbf{e}}) = \ln E[\|\bar{\underline{\mathbf{e}}}[\tau]\|^2], \tag{5}$$

where $\bar{\underline{\mathbf{e}}}[\tau] = \mathrm{cat}(\underline{\mathbf{e}}[\tau], \cdots, \underline{\mathbf{e}}[\tau+L-1]) \in \mathbb{R}^{RL}$, and cat is the concatenation operator. Intuitively, this loss encourages an optimizer to maximally reduce the energy of the AF output without needing oracle echo ($\underline{\mathbf{u}} * \underline{\mathbf{w}}$) or oracle speech ($\underline{\mathbf{s}}$) information.

3. CLASSIFICATION-TRAINED META-AF

In this work, we propose CT-Meta-AF, an extension to the Meta-AF framework [1]. CT-Meta-AF learns an AF optimizer for downstream classification tasks in a data-driven manner, without task/model dependent engineering or the need for oracle signals like clean speech. Specifically, we use an adaptive filter, such as an echo canceller, whose time-varying parameters $\boldsymbol{\theta}[\tau]$ are updated by a meta-learned update rule $g_{\boldsymbol{\phi}}$. The goal is to use the processed signals from the adaptive filter for a downstream classification task, such as keyword spotting, performed by a model $m_{\boldsymbol{\varphi}}$. To this end, we employ an end-to-end training setup that enables joint optimization of the adaptive filter, optimizer, and the classification model. Our approach uses feedback from the classifier during training to learn a custom optimizer, resulting in enhanced performance without requiring explicit coupling or the need for test-time feedback. This approach is also suitable for training on real mixtures.

3.1. Overview

To address the downstream objective, we introduce a modification to the original meta-loss from (5) that incorporates classification feedback, which we call $\mathcal{L}_{CM}$. Specifically, $\mathcal{L}_{CM}$ consists of two components: the classification loss, denoted by $\mathcal{L}_C(c, \hat{c})$, where c is the true class and $\hat{c} = m_{\boldsymbol{\varphi}}(\cdot)$ is the predicted class, and the self-supervised loss, denoted by $\mathcal{L}_M$, as defined in (5). The joint loss,

$$\mathcal{L}_{CM}(\bar{\mathbf{e}}, c, \hat{c}) = \lambda \cdot \mathcal{L}_C(c, \hat{c}) + (1-\lambda) \cdot \mathcal{L}_M(\bar{\mathbf{e}}), \tag{6}$$

uses $\lambda \in [0, 1]$ to control the weighting between classification and self-supervised losses. This approach implicitly couples the representations learned by $g_{\boldsymbol{\phi}}$ and $m_{\boldsymbol{\varphi}}$ without explicitly sharing their parameters, using the idea of "task-splitting" [19, 20].

3.2. Pretrained Classifier

When dealing with complex classification tasks, it is common to use off-the-shelf or pretrained models due to data, compute, or engineering constraints. In this setting, we assume that the classifier, $m_{\boldsymbol{\varphi}}$, is frozen and not trainable. This leads to a modified (2),

$$\hat{\boldsymbol{\phi}} = \arg\min_{\boldsymbol{\phi}} E_{\mathcal{D}}[\, \mathcal{L}_{CM}(\, m_{\boldsymbol{\varphi}}, g_{\boldsymbol{\phi}}, \mathcal{L}(h_{\boldsymbol{\theta}}, \cdots)\,)\,]. \tag{7}$$

This equation incorporates $m_{\boldsymbol{\varphi}}$ without modifying it by using $m_{\boldsymbol{\varphi}}$ as an additional loss term. This method allows custom training of the optimizer for the classifier without the classifier needing to be aware of the custom preprocessing. It customizes the optimizer to the classifier, automatically performing the model unification proposed by Seltzer et al. [10] in a model and task independent fashion.

3.3. Jointly Trained Classifier

When ample data and compute resources are available, joint or end-to-end training is a powerful approach. This involves training both the optimizer $g_{\boldsymbol{\phi}}$ and classifier $m_{\boldsymbol{\varphi}}$ simultaneously using the modified meta-loss function $\mathcal{L}_{CM}$. The parameters of the optimizer and classifier, $\boldsymbol{\phi}$ and $\boldsymbol{\varphi}$, are jointly optimized via

$$\hat{\boldsymbol{\phi}}, \hat{\boldsymbol{\varphi}} = \arg\min_{\boldsymbol{\phi}, \boldsymbol{\varphi}} E_{\mathcal{D}}[\, \mathcal{L}_{CM}(\, m_{\boldsymbol{\varphi}}, g_{\boldsymbol{\phi}}, \mathcal{L}(h_{\boldsymbol{\theta}}, \cdots)\,)\,]. \tag{8}$$

During joint training, classification feedback is used to adjust the parameters of the optimizer, improving its ability to generate outputs that aid in classification. At the same time, the classifier learns to operate on AF outputs. This builds upon the pretrained classifier setup by producing a custom-trained classifier that can better leverage the AF. To speed up training, we can initialize $\boldsymbol{\phi}$ and $\boldsymbol{\varphi}$ with pretrained weights. It is also possible to improve performance by allowing $m_{\boldsymbol{\varphi}}$ to use information beyond the outputs of $h_{\boldsymbol{\theta}}$.

4. EXPERIMENTAL DESIGN

The goal of our setup is to demonstrate that CT-Meta-AF strikes the right balance between flexibility and structure, enabling high-performance while still being a direct replacement for existing approaches. We evaluate the performance of CT-Meta-AF on KWS with device playback. The AEC removes own-playback from recordings by learning the echo path, as shown in Fig. 1. For model details, see Fig. 2. We evaluate three scenarios: one with a pre-trained KWS, one with test-time swapped KWS, and one with a jointly trained KWS, all implemented using Meta-AF [1]. All CT-Meta-AFs use the classifier for training. The jointly trained KWS is optimized end-to-end and trained simultaneously with the AEC.

4.1. AEC Configurations

To perform AEC, we fit an AF to cancel echo. At inference time, we only assume access to the playback, $\underline{\mathbf{u}}$ and echoic mixture, $\underline{\mathbf{d}}$. The AF uses frequency-domain overlap save with a window of $K = 1024$, a hop of $R = 512$, and $B = 4$ blocks. Each block is denoted by ℂFilter in Fig. 2. Each AEC filters the farend $\mathbf{u}[\tau]$ with estimated filter $\hat{\mathbf{w}}[\tau]$ and subtracts the result from $\mathbf{d}[\tau]$. For a review of AEC see Benesty et al. [27]. We compare our approach to four baselines: regular Meta-AF [1], a Kalman filter AEC (Diag. KF) [24], No-Echo, and No-AEC. We grid-search-tune Diag. KF using KWS accuracy. No-Echo simulates running the KWS model in a playback-free environment, while No-AEC performs no echo cancellation and simulates deploying a KWS without AEC.

CT-Meta-AF and Meta-AF use the same higher-order architecture [23] and training scheme, with the only difference being the loss. For CT-Meta-AF, we use $\lambda = 0.5$ in (6), and for regular Meta-AF, we use 0. We set the group size, group hop, and hidden size to 5, 2, and 48, respectively (see Fig. 2 for details). We use Adam with a batch size of 16, a learning rate of $2 \cdot 10^{-4}$, momentum $\beta_1 = 0.99$, and we randomize the truncation length L. Additionally, we apply gradient clipping and reduce the learning rate by half if the validation performance does not improve for 10 epochs, and we stop training after 30 epochs with no improvement. For pretrained KWS experiments, we train ϕ from scratch, and for joint training, we initialize with the best pretrained ϕ. When joint training with Diag. KF or Meta-AF we do not update the pretrained ϕ and just train the KWS. Each model has 32 K complex parameters and trains on one GPU.

4.2. KWS Configurations

We set up KWS as a multi-class classification task, where the keyword $\underline{\mathbf{s}}$ belongs to a single class c. We base our KWS on the model proposed by Cornell et al. [20]. Our version uses 40 log-mel-filter-bank inputs and short-time Fourier transform with a 256 point hop and a 512 point window. The KWS has 3 residual blocks, each with a 1×1 convolution, layer norm, ReLU, a dilated convolution with kernel size 5, layer norm, ReLU, and a final 1×1 convolution. The last residual block is averaged across time and fed to a dense layer with softmax to predict the class distribution. See Fig. 2. We use a binary cross-entropy loss and train the KWS for 50 epochs with a batch size of 128, a learning rate of 10^{-3}, and use the best checkpoint based on validation performance. We pretrain on a dataset without playback. For joint training, due to limited GPU capacity, we use a batchsize of 16, learning rate of 10^{-4}, $\beta_1 = .9$, reduce the learning rate by half after 10 epochs with no improvement, and stop training after 50 epochs with no improvement. All KWS models have ≈ 300K parameters and use ≈ 20MFLOPs per second.
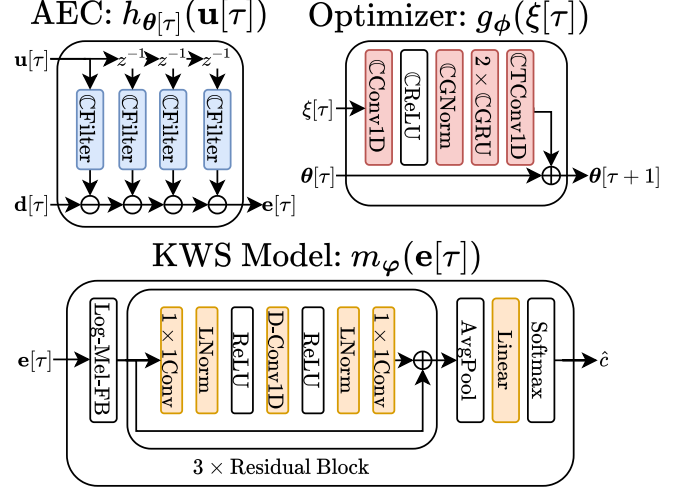

Fig. 2: Detailed view of the AEC, Optimizer, and KWS. The AEC preprocesses the KWS inputs and the optimizer controls the AEC parameters. At test time, only the AEC parameters change

4.3. Dataset & Metrics

We use synthetic playback during training and evaluate our models using both synthetic and real playback data from the Microsoft AEC Challenge [22] and keywords from the Google Speech Commands V2 dataset [21], sampled at 16 KHz. During training, we trim playback to 3 seconds, randomly mix playback with keywords, zero-pad keywords as needed, apply a random shift, and set the signal-to-echo ratio (SER) uniformly at random between -25 dB and 0 dB. Each keyword is used once. To test, we use playback trimmed to either 4 or 12 seconds. We evaluate three datasets: a 35-class dataset and two smaller datasets with 10 and 2 classes, respectively. The smaller datasets contain approximately one-third and one-tenth of the full dataset. We always use the same folds and never mix across folds.

To evaluate performance, we use macro and micro averaged F1, where higher is better. The goal is to reflect the class imbalance in the KWS Dataset. We display results as macro F1 (micro F1).

5. RESULTS & DISCUSSION

Here, we investigate the effect of classification training for Meta-AF on AEC with a downstream KWS. This task simulates a device attempting to classify a spoken keyword while emitting playback. We show results for zero playback (No-Echo), no cancellation (No-AEC), Diag. KF, Meta-AF based AEC without classification training (Meta-AEC), and Meta-AF based AEC with classification training (CT-Meta-AEC). The Diag. KF and Meta-AEC baselines are state-of-the-art traditional and data-driven approaches. All models are trained on synthetic playback with challenging SERs distributed uniformly at random between -25 dB and 0 dB, and real playback is only used for evaluation. We use an off-the-shelf KWS that is not customized for AEC in sections 5.1 and 5.2 but do experiment with retraining the KWS in section 5.3. We evaluate longer playback and compute statistical significance for select models. Across experiments, CT-Meta-AF proved highly stable and required no additional tuning, making it a drop-in replacement for prior approaches.

5.1. Pretrained KWS

First, we test the effectiveness of classification training as described in 3.2. In Table 1, all AEC models share a frozen KWS model pretrained on keywords without playback. This setup mimics using an off-the-shelf KWS without retraining but with access to your own dataset, a common real-world setup. For CT-Meta-AEC training, we use $\lambda = 0.5$ when computing $\mathcal{L}_{CM}$. The F1 degradation from No-Echo to No-AEC shows the importance of AEC.

In Table 1, CT-Meta-AEC significantly outperforms its signal-level Diag. KF and Meta-AEC counterparts, as indicated by better F1 scores. In synthetic playback (left column), CT-Meta-AEC improves over Meta-AEC by some .165 F1. This gap persists in real playback, where CT-Meta-AEC is the top performing model by .091 F1. Interestingly, CT-Meta-AEC achieves a 5.80 dB reduction in echo, while regular Meta-AF attains a 9.57 dB reduction. This observation underscores the limitations of signal-level metrics as effective descriptors of downstream performance, especially in the context of nonlinear processing methods such as DNN-based KWS. Classification-based training effectively addresses this discrepancy and optimizes downstream performance. In a secondary evaluation with 12-second signals, CT-Meta-AEC outperforms Meta-AEC by 0.279 F1 in synthetic playback and 0.119 F1 in real playback.

5.2. Specialization and Sensitivity

Next, we study the specialization capabilities of CT-Meta-AEC in different KWS setups. We use the approach from 3.2 to train three CT-Meta-AEC models: CT-Meta-AEC-35, CT-Meta-AEC-10, and CT-Meta-AEC-2, corresponding to datasets with 35, 10, and 2 keyword classes, respectively. We assess the performance of each CT-Meta-AEC model on all three datasets, but always use a downstream KWS model specifically trained on a matching keyword setup. Table 2 presents the results of our experiments on synthetic playback scenarios, where each row corresponds to a distinct AEC model and each column represents a unique test-time KWS configuration. The first row displays our Meta-AEC baseline. Table entries on the diagonal show matched train/test setups.

In all columns of Table 2, CT-Meta-AEC with a matching KWS performs best. This demonstrates that CT-Meta-AEC is learning a specialized optimizer (the diagonal) which outperforms both generalist optimizer (top row) and optimizers trained on different models (off diagonal). Notably, results above the diagonal show that specialization can yield superior results even in the presence of limited training data. Specifically, Meta-AEC and CT-Meta-AEC-35, which are trained with three times more data than CT-Meta-AEC-10, are surpassed by the latter on 10 class KWS. This suggests that classification training is a data-efficient approach for improving performance on downstream tasks. The results from mismatched train/test setups show that CT-Meta-AECs are good at related tasks, and usually still outperform Meta-AEC.

5.3. Jointly Trained KWS

We investigate joint training of the AEC and KWS to simulate having sufficient data for training a complete system, without oracle signal-level supervision. Both the AEC and KWS models undergo joint training, following the approach in 3.3. The KWS models undergo a pre-training phase on scenes devoid of playback, followed by fine-tuning using the output of their corresponding AEC models. This represents a significant departure from the approach outlined in 3.2, where KWS models were trained without consideration for

Model	Synthetic Playback	Real Playback
No-Echo	.928 (.935)	.931 (.936)
No-AEC	.087 (.098)	.102 (.117)
Diag. KF	.196 (.208)	.165 (.180)
Meta-AEC	.335 (.340)	.226 (.236)
CT-Meta-AEC	**.500 (.508)**	**.317 (.326)**

Table 1: F1 Macro (F1 Micro) with a pretrained KWS and SER $\in [-25, 0]$ dB. No-AEC and No-Echo are lower/upper bounds.

Model	35 Class	10 Class	2 Class
Meta-AEC	.335 (.340)	.476 (.465)	.751 (.754)
CT-Meta-AEC-35	**.500 (.508)**	.513 (.483)	.786 (.788)
CT-Meta-AEC-10	.344 (.348)	**.533 (.529)**	.781 (.781)
CT-Meta-AEC-2	.274 (.281)	.413 (.400)	**.802 (.802)**

Table 2: Specialization of classification-trained Meta-AEC with performance shown as F1 Macro (F1 Micro). A different model with different keywords is swapped in at test time. SER $\in [-25, 0]$ dB.

Model	Synthetic Playback	Real Playback
No-Echo	.922 (.928)	.931 (.936)
No-AEC	.609 (.619)	.583 (.593)
Diag. KF	.778 (.781)	.690 (.698)
Meta-AEC*	.870 (.876)	.746 (.754)
JCT-Meta-AEC*	**.871 (.877)**	**.754 (.762)**

Table 3: F1 Macro (F1 Micro) with a jointly trained KWS and SER $\in [-25, 0]$ dB. No-AEC and No-Echo are lower/upper bounds. An asterix denotes that we averaged three trials.

echo presence. We call this approach JCT-Meta-AEC. For statistical robustness, we ran three separate trials of the Meta models.

In Table 3, we find that the training scheme from 3.3 improves performance. JCT-Meta-AEC outperforms Meta-AEC by .008 F1 (p-value .01) in real playback scenarios, underscoring the efficacy of classification training for real-world performance. While JCT-Meta-AEC outperforms Meta-AEC on synthetic data, the improvement is not significant. However, on the longer 12-second test-set, JCT-Meta-AEC beats Meta-AEC by a margin of 0.057 (p-value .0005). Of note, we did not encounter any training stability issues.

6. CONCLUSION

We proposed a simple yet effective approach for improving the performance of downstream classification tasks by incorporating classifier feedback into learned adaptive filter update rules. Our approach, classification-trained meta-adaptive filtering (CT-Meta-AF) enables end-to-end training without oracle single-level supervision. We evaluated CT-Meta-AF on echo cancellation and keyword spotting and observed consistent performance gains across synthetic and real playback, multiple keyword configurations, and both pre and jointly trained keyword models. CT-Meta-AF yields performance improvements without additional inference costs or manual tuning. We believe our approach has the potential to improve many adaptive filter pipelines, thanks to it's plug-and-play design and promising real-world performance.

7. REFERENCES

[1] Jonah Casebeer, Nicholas J. Bryan, and Paris Smaragdis, "Meta-AF: Meta-Learning for Adaptive Filters," *IEEE Transactions on Audio, Speech, and Language Processing (TASLP)*, 2022.

[2] Reinhold Haeb-Umbach, Shinji Watanabe, Tomohiro Nakatani, Michiel Bacchiani, Bjorn Hoffmeister, Michael L Seltzer, Heiga Zen, and Mehrez Souden, "Speech Processing for Digital Home Assistants: Combining Signal Processing With Deep-Learning Techniques," *IEEE Signal Processing Magazine (SPM)*, 2019.

[3] Bernard Widrow and Samuel D. Stearns, *Adaptive Signal Processing*, Prentice-Hall, 1985.

[4] V. John Mathews, "Adaptive Polynomial Filters," *IEEE Signal Processing Magazine (SPM)*, 1991.

[5] Simon S. Haykin, *Adaptive Filter Theory*, Pearson, 2008.

[6] Jonah Casebeer, Jacob Donley, Daniel Wong, Buye Xu, and Anurag Kumar, "NICE-Beam: Neural Integrated Covariance Estimators for Time-Varying Beamformers," *arXiv:2112.04613*, 2021.

[7] Amir Ivry, Israel Cohen, and Baruch Berdugo, "Deep Adaptation Control for Acoustic Echo Cancellation," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022.

[8] Thomas Haubner, Andreas Brendel, and Walter Kellermann, "End-To-End Deep Learning-Based Adaptation Control for Frequency-Domain Adaptive System Identification," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022.

[9] Jonah Casebeer, Nicholas J Bryan, and Paris Smaragdis, "Auto-DSP: Learning to Optimize Acoustic Echo Cancellers," in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2021.

[10] Michael L Seltzer, Bhiksha Raj, and Richard M Stern, "Likelihood-Maximizing Beamforming for Robust Hands-Free Speech Recognition," *IEEE Transactions on Speech and Audio Processing (TSAP)*, 2004.

[11] Guangji Shi, Parham Aarabi, and Hui Jiang, "Phase-Based Dual-Microphone Speech Enhancement Using a Prior Speech Model," *IEEE Transactions on Audio, Speech, and Language Processing (TASLP)*, 2006.

[12] Xianyu Zhao and Zhijian Ou, "Closely Coupled Array Processing and Model-Based Compensation for Microphone Array Speech Recognition," *IEEE Transactions on Audio, Speech, and Language Processing (TASLP)*, 2007.

[13] Jahn Heymann, Lukas Drude, Christoph Boeddeker, Patrick Hanebrink, and Reinhold Haeb-Umbach, "Beamnet: End-To-End Training of a Beamformer-Supported Multi-Channel Asr System," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017.

[14] Tsubasa Ochiai, Shinji Watanabe, Takaaki Hori, John R Hershey, and Xiong Xiao, "Unified Architecture for Multichannel End-To-End Speech Recognition With Neural Beamforming," *IEEE Journal of Selected Topics in Signal Processing (JSTSP)*, 2017.

[15] Wangyou Zhang, Xuankai Chang, Christoph Boeddeker, Tomohiro Nakatani, Shinji Watanabe, and Yanmin Qian, "End-To-End Dereverberation, Beamforming, and Speech Recognition in a Cocktail Party," *IEEE Transactions on Audio, Speech, and Language Processing (TASLP)*, 2022.

[16] Wangyou Zhang, Christoph Boeddeker, Shinji Watanabe, Tomohiro Nakatani, Marc Delcroix, Keisuke Kinoshita, Tsubasa Ochiai, Naoyuki Kamo, Reinhold Haeb-Umbach, and Yanmin Qian, "End-To-End Dereverberation, Beamforming, and Speech Recognition With Improved Numerical Stability and Advanced Frontend," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021.

[17] Sri Harish Mallidi, Roland Maas, Kyle Goehner, Ariya Rastrow, Spyros Matsoukas, and Björn Hoffmeister, "Device-directed Utterance Detection," *Interspeech*, 2018.

[18] Nathan Howard, Alex Park, Turaj Zakizadeh Shabestary, Alexander Gruenstein, and Rohit Prabhavalkar, "A Neural Acoustic Echo Canceller Optimized Using an Automatic Speech Recognizer and Large Scale Synthetic Data," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021.

[19] Sebastian Braun and Maria Luis Valero, "Task Splitting for Dnn-Based Acoustic Echo and Noise Removal," in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2022.

[20] Samuele Cornell, Thomas Balestri, and Thibaud Sénéchal, "Implicit Acoustic Echo Cancellation for Keyword Spotting and Device-Directed Speech Detection," in *Workshop on Spoken Language Technology (SLT)*. IEEE, 2023.

[21] Pete Warden, "Speech Commands: A Dataset for Limited-Vocabulary Speech Recognition," *arXiv:1804.03209*, 2018.

[22] Ross Cutler, Ando Saabas, Tanel Parnamaa, Marju Purin, Hannes Gamper, Sebastian Braun, Karsten Sorensen, and Robert Aichner, "ICASSP 2022 Acoustic Echo Cancellation Challenge," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.

[23] Junkai Wu, Jonah Casebeer, Nicholas J. Bryan, and Paris Smaragdis, "Meta-Learning for Adaptive Filters with Higher-Order Frequency Dependencies," in *IEEE International Workshop on Acoustic Signal Enhancement (IWAENC)*, 2022.

[24] Fabian Kuech, Edwin Mabande, and Gerald Enzner, "State-Space Architecture of the Partitioned-Block-Based Acoustic Echo Controller," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2014.

[25] J-S Soo and Khee K Pang, "Multidelay Block Frequency Domain Adaptive Filter," *IEEE Transactions on Acoustics, Speech, and Signal Processing (TASSP)*, 1990.

[26] Marcin Andrychowicz, Misha Denil, Sergio Gómez Colmenarejo, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando de Freitas, "Learning to Learn by Gradient Descent by Gradient Descent," in *NeurIPS*, 2016.

[27] Jacob Benesty, Tomas Gänsler, Dennis R Morgan, Steven L. Gay, and M. Mohan Sondhi, *Advances in Network and Acoustic Echo Cancellation*, Springer, 2001.