

Received 6 August 2023; revised 27 September 2023; accepted 27 September 2023. Date of publication 13 March 2024; date of current version 18 June 2024. The review of this article was arranged by Associate Editor Lisandro Lovisolo.

Digital Object Identifier 10.1109/OJSP.2024.3376289

# ICASSP 2023 Acoustic Echo Cancellation Challenge

ROSS CUTLER <sup>1</sup> (Member, IEEE), ANDO SAABAS <sup>1</sup>, TANEL PÄRNAMAA <sup>1</sup>, MARJU PURIN <sup>1</sup>,  
EVGENII INDENBOM <sup>1</sup>, NICOLAE-CĂTĂLIN RISTEA <sup>1</sup>, JEGOR GUŽVIN <sup>1</sup>,  
HANNES GAMPER <sup>1</sup> (Member, IEEE), SEBASTIAN BRAUN <sup>1</sup> (Senior Member, IEEE),  
AND ROBERT AICHNER <sup>1</sup>

<sup>1</sup>Microsoft Corporation, Redmond, WA 98052 USA

CORRESPONDING AUTHOR: ROSS CUTLER (e-mail: ross.cutler@microsoft.com).

This work involved human subjects or animals in its research. The authors confirm that all human/animal subject research procedures and protocols are exempt from review board approval.

**ABSTRACT** The ICASSP 2023 Acoustic Echo Cancellation Challenge is intended to stimulate research in acoustic echo cancellation (AEC), which is an important area of speech enhancement and is still a top issue in audio communication. This is the fourth AEC challenge and it is enhanced by adding a second track for personalized acoustic echo cancellation, reducing the algorithmic + buffering latency to 20 ms, as well as including a full-band version of AECMOS (Purin et al., 2020). We open source two large datasets to train AEC models under both single talk and double talk scenarios. These datasets consist of recordings from more than 10,000 real audio devices and human speakers in real environments, as well as a synthetic dataset. We open source an online subjective test framework and provide an objective metric for researchers to quickly test their results. The winners of this challenge were selected based on the average mean opinion score (MOS) achieved across all scenarios and the word accuracy (WAcc) rate.

**INDEX TERMS** Acoustic echo cancellation, deep learning, speech enhancement, speech quality assessment, subjective test.

## I. INTRODUCTION

With the growing popularity and need for working remotely, the use of teleconferencing systems such as Microsoft Teams, Skype, WebEx, Zoom, etc., has increased significantly. It is imperative to have good quality calls to make the user's experience pleasant and productive. The degradation of call quality due to acoustic echoes is one of the major sources of poor speech quality ratings in voice and video calls. While digital signal processing (DSP) based AEC models have been used to remove these echoes during calls, their performance can degrade when model assumptions are violated, e.g., fast time-varying acoustic conditions, unknown signal processing blocks or non-linearities in the processing chain, or failure of other models (e.g., background noise estimates). This problem becomes more challenging during full-duplex modes of communication where echoes from double talk scenarios

are difficult to suppress without significant distortion or attenuation [2].

With the advent of deep learning techniques, many supervised learning algorithms for AEC have shown better performance compared to their classical counterparts, e.g., [3], [4], [5]. Some studies have also shown good performance using a combination of classical and deep learning methods such as using adaptive filters and recurrent neural networks (RNNs) [5], [6] but only on synthetic datasets. While these approaches are promising, they lack evidence of their performance on real-world datasets with speech recorded in diverse noise and reverberant environments. This makes it difficult for researchers in the industry to choose a good model that can perform well on a representative real-world dataset.

Most AEC publications use objective measures such as echo return loss enhancement (ERLE) [7] and perceptual

**TABLE 1** Pearson Correlation Coefficient (PCC) and Spearman’s Rank Correlation Coefficient (SRCC) Between Objective and Subjective P.808 Results on Single Talk Echo Scenarios (See Section VII).

	PCC	SRCC
ERLE	0.31	0.23
PESQ	0.67	0.57

evaluation of speech quality (PESQ) [8]. ERLE in dB is defined as:

$$\text{ERLE} = 10 \log_{10} \frac{\mathbb{E}[y^2(n)]}{\mathbb{E}[e^2(n)]} \quad (1)$$

where  $y(n)$  is the microphone signal, and  $e(n)$  is the residual echo after cancellation. ERLE is only appropriate when measured in a quiet room with no background noise and only for single talk scenarios (not double talk), where we can use the processed microphone signal as an estimate for  $e(n)$ . PESQ has also been shown to not have a high correlation to subjective speech quality in the presence of background noise [9]. Using the datasets provided in this challenge we show that ERLE and PESQ have a low correlation to subjective tests (Table 1). In order to use a dataset with recordings in real environments, we can not use ERLE and PESQ. A more reliable and robust evaluation framework is needed that everyone in the research community can use, which we provide as part of the challenge.

This AEC challenge is designed to stimulate research in the AEC domain by open-sourcing a large training dataset, test set, and subjective evaluation framework. We provide two new open-source datasets for training AEC models. The first is a real dataset captured using a large-scale crowdsourcing effort. This dataset consists of real recordings that have been collected from over 10,000 diverse audio devices and environments. The second dataset is synthesized from speech recordings, room impulse responses, and background noise derived from [10]. An initial test set was released for the researchers to use during development and a blind test set near the end, which has been used to decide the final competition winners. We believe these datasets are large enough to facilitate deep learning and representative enough for practical usage in shipping telecommunication products (e.g., see [11]).

This is the fourth AEC challenge we have conducted. The first challenge was held at ICASSP 2021 [12], the second at INTERSPEECH 2021 [13], and the third at ICASSP 2022 [14]. These challenges had 49 participants with entries ranging from pure deep models, hybrid linear AEC + deep echo suppression, and DSP methods. While the submitted AECs have consistently been getting better, there is still significant room for improvement as shown in Table 2. The two largest areas for improvement are (1) Single Talk Near End quality, which is affected by background noise, reverberation, and capture device distortions, and (2) Double Talk Other Degradations, which includes missing audio, distortions, and cut-outs. In addition, the overall challenge metric,  $M$  was

**TABLE 2** Amount of Improvement Remaining Based on the ICASSP 2022 AEC Challenge [14].

Area	Headroom
Single talk near end MOS	0.60
Single talk far end MOS	0.19
Double talk echo MOS	0.25
Double talk other MOS	0.57
WAcc	0.19

0.883 out of 1.0 in the ICASSP 2022 challenge, which also shows significant room for improvement.

To improve the challenge and further stimulate research in this area we have made the following changes:

- We included a second track for personalized AEC. Based on the excellent results for personalized noise suppression in the ICASSP 2022 Deep Noise Suppression Challenge [15], we expected significant improvements for the double talk scenario.
- We reduced the algorithmic latency + buffering latency from 40 ms to 20 ms which is necessary for use in real-time collaboration systems. This will make getting the same speech quality done in previous challenges more difficult.
- We provided a full-band version of AECMOS so it can be better used for full-band training and testing. AECMOS is freely available at <https://github.com/microsoft/AEC-Challenge>.

An overview of the four AEC challenges is given in Table 3.

Related work is reviewed in Section II. The challenge description is given in Section III. The training dataset is described in Section IV, and the test set in Section V. We describe a baseline deep neural network-based AEC method in Section VI. The online subjective evaluation framework is discussed in Section VII, and the objective function in Section VIII. The challenge metric is given in Section IX and the challenge rules are described in <https://aka.ms/aec-challenge>. The results and analysis are given in Section X, and conclusions are discussed in Section XI.

## II. RELATED WORK

There are many standards for measuring AEC performance. For objective metrics, IEEE 1329 [2] defines metrics like terminal coupling loss for single talk (TCLwst) and double talk (TCLwdt), which are measured in anechoic chambers. TIA 920 [16] uses many of these metrics but defines required criteria. ITU-T Rec. G.122 [17] defines AEC stability metrics, and ITU-T Rec. G.131 [18] provides a useful relationship of acceptable Talker Echo Loudness Rating and one-way delay time. ITU-T Rec. G.168 [19] provides a comprehensive set of AEC metrics and criteria. However, it is not clear how to combine these dozens of metrics into a single metric, or how well these metrics correlate to subjective quality.

**TABLE 3** Summary of AEC Challenges. BAK and SIG are Measurements of the Background Noise Quality and Speech Signal Quality.

Challenge	Tracks	Datasets	Algorithmic + Buffering Latency	Notes
ICASSP 2021	Real-time	2,500 real environments Synthetic	40 ms	Crowdsourced P.831
INTERSPEECH 2021	Real-time Non-real-time	5,000 real environments Synthetic	40 ms	Made test set more comprehensive Increased subjective test framework accuracy Added AECMOS service
ICASSP 2022	Real-time	7,500 real environments Synthetic	40 ms	Added mobile scenarios Added WAcc Made datasets, test sets full band
ICASSP 2023	Real-time Personalized	10,000 real environments Synthetic	20 ms	Added fullband AECMOS Split near end quality into BAK and SIG

Subjective speech quality assessment is the gold standard for evaluating speech enhancement processing and telecommunication systems, and the ITU-T has developed several recommendations for subjective speech quality assessment. ITU-T P.800 [20] describes lab-based methods for the subjective determination of speech quality. In P.800, users are asked to rate the quality of speech clips on a Likert scale from 1: Poor to 5: Excellent. Many ratings are taken for each clip, and the average score for each clip is the MOS. ITU-T P.808 [21] describes a crowdsourcing approach for conducting subjective evaluations of speech quality. It provides guidance on test material, experimental design, and a procedure for conducting listening tests in the crowd. The methods are complementary to laboratory-based evaluations described in P.800. An open-source implementation of P.808 is described in [22]. ITU-T P.835 [23] provides a subjective evaluation framework that gives standalone quality scores of speech (SIG) and background noise (BAK) in addition to the overall quality (OVRL). An open-source implementation of P.835 is described in [24]. More recent multidimensional speech quality assessment standards are ITU-T P.863.2 [25] and P.804 [26] (listening phase), which measure noisiness, coloration, discontinuity, and loudness. An open-source implementation of P.804 using crowdsourcing is described in [27].

ITU-T Rec. P.831 [28] provides guidelines on how to conduct subjective tests for network echo cancellers in the laboratory. ITU-T Rec. P.832 [8] focuses on the hands-free terminals and covers a broader range of degradations. Cutler et al. [29] provide an open-source crowdsourcing tool extending P.831 and P.832 and include validation studies that show it is accurate compared to expert listeners and repeatable across multiple days and different raters. Purin et al. [1] created an objective metric, AECMOS, based on this tool's results on hundreds of different AEC models. AECMOS has a high correlation to subjective opinion.

While there have been hundreds of papers published on deep echo cancellation since the first AEC challenge, we feel the winners of each challenge are of special note since they have been tested and evaluated using realistic and challenging test sets and subjective evaluations. Table 4 provides

**TABLE 4** AEC Challenge Top 3 Performers.

Challenge	1st	Rank	2nd	Rank	3rd	Rank
ICASSP 2021	[30]	1	[31]	2	[32]	5
INTERSPEECH 2021	[33]	6	[34]	8	[35]	10
ICASSP 2022	[36]	1	[37]	2	[38]	3

the top three papers for each previous AEC challenge. Note that because the performance rankings and paper acceptances were decoupled in ICASSP 2021 and INTERSPEECH 2021, the challenge placement and performance rankings are not identical, and for INTERSPEECH 2021 not well correlated. For ICASSP 2022 and 2023, the top five papers based on the challenge performance were submitted for review, fixing the disparity between paper acceptance and model performance.

### III. CHALLENGE DESCRIPTION

#### A. TRACKS

This challenge included two tracks:

- Non-personalized AEC. This is similar to the ICASSP 2022 AEC Challenge.
- Personalized AEC. This adds speaker enrollment for the near end speaker. A speaker enrollment is a 15-25 s recording of the near end speaker that can be used for adopting the AEC for personalized echo cancellation. For training and model evaluation, the datasets in <https://github.com/microsoft/AEC-Challenge> can be used, which include both echo and near end only clips from users. For the blind test set, the enrollment clips will be provided.

#### B. LATENCY AND RUNTIME REQUIREMENTS

Algorithmic latency is defined by the offset introduced by the whole processing chain including short-time Fourier transform (STFT), inverse STFT, overlap-add, additional lookahead frames, etc., compared to just passing the signal through without modification. It does not include buffering latency. Some examples are:

- A STFT-based processing with window length = 20 ms and hop length = 10 ms introduces an algorithmic delay of window length – hop length = 10 ms.
- A STFT-based processing with window length = 32 ms and hop length = 8 ms introduces an algorithmic delay of window length – hop length = 24 ms.
- An overlap-save-based processing algorithm introduces no additional algorithmic latency.
- A time-domain convolution with a filter kernel size = 16 samples introduces an algorithmic latency of kernel size – 1 = 15 samples. Using one-sided padding, the operation can be made fully “causal”, i.e., left-sided padding with kernel size-1 samples would result in no algorithmic latency.
- A STFT-based processing with window\_length = 20 ms and hop\_length = 10 ms using 2 future frames information introduces an algorithmic latency of (window\_length – hop\_length) + 2\*hop\_length = 30 ms.

Buffering latency is defined as the latency introduced by block-wise processing, often referred to as hop length, frame-shift, or temporal stride. Some examples are:

- A STFT-based processing has a buffering latency corresponding to the hop size.
- An overlap-save processing has a buffering latency corresponding to the frame size.
- A time-domain convolution with stride 1 introduces a buffering latency of 1 sample.

Real-time factor (RTF) is defined as the fraction of time it takes to execute one processing step. For a STFT-based algorithm, one processing step is the hop size. For a time-domain convolution, one processing step is 1 sample.  $RTF = \text{compute time} / \text{time step}$ .

All models submitted to this challenge must meet all of the below requirements:

- 1) To be able to execute an algorithm in real-time, and to accommodate for variance in compute time which occurs in practice, we require  $RTF \leq 0.5$  in the challenge on an Intel Core i5 Quadcore clocked at 2.4 GHz using a single thread.
- 2) Algorithmic latency + buffering latency  $\leq 20$  ms.
- 3) No future information can be used during model inference.

#### IV. TRAINING DATASETS

The challenge includes two open-source datasets, one real and one synthetic. The datasets are available at <https://github.com/microsoft/AEC-Challenge>.

##### A. REAL DATASET

The first dataset was captured using a large-scale crowdsourcing effort. This dataset consists of more than 50,000 recordings from over 10,000 different real environments, audio devices, and human speakers in the following scenarios:

- 1) Far end single talk, no echo path change
- 2) Far end single talk, echo path change
- 3) Near end single talk, no echo path change
- 4) Double talk, no echo path change




FIGURE 1. Sorted near end single talk clip quality (P.808) with 95% confidence intervals.





FIGURE 2. Distribution of reverberation time (RT60) for desktop (left) and mobile (right).


- 5) Double talk, echo path change
- 6) Sweep signal for RT60 estimation

RT60 is the time for an initial signal’s sound pressure level to attenuate 60 dB from its original level. For the far end single talk case, there is only the loudspeaker signal (far end) played back to the users and users remain silent (no near end speech). For the near end single talk case, there is no far end signal and users are prompted to speak, capturing the near end signal. For double talk, both the far end and near end signals are active, where a loudspeaker signal is played and users talk at the same time. Echo path changes were incorporated by instructing the users to move their device around or bring themselves to move around the device. The RT60 distribution for 4387 desktop environments in the real dataset for which impulse response measurements were available is estimated using a method by Karjalainen et al. [39] and shown in Fig. 2. For 1251 mobile environments the RT60 distribution shown was estimated blindly from speech recordings [40]. The RT60 estimates can be used to sample the dataset for training. The near end single talk speech quality is given in Fig. 1.

We use *Amazon Mechanical Turk* as the crowdsourcing platform and wrote a custom HIT application that includes a custom tool that users download and execute to record the six scenarios described above. The dataset includes Microsoft Windows and Android devices. Each scenario includes the microphone and loopback signal (see Fig. 3). Even though our application uses the WASAPI raw audio mode to bypass built-in audio effects, the PC can still include Audio DSP on the receive signal (e.g., equalization and Dynamic Range Compression); it can also include Audio DSP on the send signal, such as AEC and noise suppression.



**FIGURE 3.** Custom recording application recorded the loopback and microphone signals. ADSP is audio digital signal processing, and DRC is dynamic range compression.



**FIGURE 4.** Distribution of reverberation time (RT60) for the synthetic dataset.

For far end signals, we use both clean speech and real-world recordings. For clean speech far end signals, we use the speech segments from the Edinburgh dataset [41]. This corpus consists of short single speaker speech segments (1 to 3 seconds). We used a long short term memory (LSTM) [42] based gender detector to select an equal number of male and female speaker segments. Further, we combined 3 to 5 of these short segments to create clips of length between 9 and 15 seconds in duration. Each clip consists of a single gender speaker. We create a gender-balanced far end signal source comprising of 500 male and 500 female clips. Recordings are saved at the maximum sampling rate supported by the device and in 32-bit floating point format; in the released dataset we down-sample to 48 kHz and 16-bit using automatic gain control to minimize clipping.

For noisy speech far end signals we use 2000 clips from the near end single talk scenario. Clips are gender balanced to include an equal number of male and female voices.

For the far end single talk scenario, the clip is played back twice. This way, the echo canceller can be evaluated both on the first segment, when it has had minimal time to converge, and on the second segment, when the echo canceller has converged and the result is more indicative of a real call scenario.

For the double talk scenario, the far end signal is similarly played back twice, but with an additional silent segment in the middle, when only near end single talk occurs.

For near end speech, the users were prompted to read sentences from a TIMIT [43] sentence list. Approximately 10 seconds of audio is recorded while the users are reading.

For track two (personalized AEC) we include 30 seconds of target speaker for each clip in the test set. In addition, the training and test set from the ICASSP 2022 Deep Noise Suppression Challenge track two [15] can be used.

### B. SYNTHETIC DATASET

The second dataset provides 10,000 synthetic scenarios, each including single talk, double talk, near end noise, far end noise, and various nonlinear distortion scenarios. Each scenario includes a far end speech, echo signal, near end speech, and near end microphone signal clip. We use 12,000 cases (100 hours of audio) from both the clean and noisy speech datasets derived in [10] from the LibriVox project<sup>1</sup> as source clips to sample far end and near end signals. The LibriVox project is a collection of public-domain audiobooks read by volunteers. [10] used the online subjective test framework ITU-T P.808 to select audio recordings of good quality ( $4.3 \leq \text{MOS} \leq 5$ ) from the LibriVox project. The noisy speech dataset was created by mixing clean speech with noise clips sampled from AudioSet [44], Freesound<sup>2</sup> and DEMAND [45] databases at signal to noise ratios sampled uniformly from [0, 40] dB.

To simulate a far end signal, we pick a random speaker from a pool of 1,627 speakers, randomly choose one of the clips from the speaker, and sample 10 seconds of audio from the clip. For the near end signal, we randomly choose another speaker and take 3–7 seconds of audio which is then zero-padded to 10 seconds. The selected far end speakers were 71% male, and 67% of the near end speakers were male. To generate an echo, we convolve a randomly chosen room impulse response from a large Microsoft unreleased database with the far end signal. The room impulse responses are generated by using Project Acoustics technology<sup>3</sup> and the RT60 ranges from 200 ms to 1200 ms. The distribution of RT60 is shown in Fig. 4. In 80% of the cases, the far end signal is processed by a nonlinear function to mimic loudspeaker distortion (the linear-to-nonlinear ratio is 0.25). For example, the transformation can be clipping the maximum amplitude, using a sigmoidal function as in [46], or applying learned distortion functions, the details of which we will describe in a future paper. This signal gets mixed with the near end signal at

<sup>1</sup>[Online]. Available: <https://librivox.org>

<sup>2</sup>[Online]. Available: <https://freesound.org>

<sup>3</sup>[Online]. Available: <https://www.aka.ms/acoustics>




FIGURE 5. Echo canceller test set-up for third party listening test B according to the ITU-T Rec.P.831 (after [28]). *S* is send and *R* is receive.




FIGURE 6. Double talk scenario in the third party listening test. The test participant (marked by “You”) is positioned in the center of the communication.

TABLE 5 AECMOS Performance Using Pearson’s Correlation Coefficient (PCC), Spearman’s Rank Correlation Coefficient SRCC, and Kendall’s Tau-B With a 95% Confidence Interval [49].

Scenario	PCC	SRCC	Tau-b 95
Far end single talk echo DMOS	0.99	0.92	0.83
Near end single talk SIG MOS	0.85	0.70	0.59
Near end single talk BAK MOS	0.93	0.88	0.74
Double talk echo DMOS	0.99	0.96	0.86
Double talk other DMOS	0.96	0.91	0.85

a signal-to-echo ratio uniformly sampled from  $-10$  dB to  $10$  dB. The signal-to-echo ratio is calculated based on the clean speech signal (i.e., a signal without near end noise). The far end and near end signals are taken from the noisy dataset in 50% of the cases. The first 500 clips can be used for validation as these have a separate list of speakers and room impulse responses. Detailed metadata information can be found in the repository.

### V. TEST SET

Two test sets are included, one at the beginning of the challenge and a blind test set near the end. Both consist of 800 real-world recordings, between 30–45 seconds in duration. The test sets include the following scenarios that make echo cancellation more challenging:

- Long - or varying delays, i.e., files where the delay between loopback and mic-in is atypically long or varies during the recording
- Strong speaker and/or mic distortions
- Stationary near end noise
- Non-stationary near end noise

- Recordings with audio DSP processing from the device, such as AEC or noise reduction
- Glitches, i.e., files with “choppy” audio, for example, due to very high CPU usage
- Gain variations, i.e., recordings where far end level changes during the recording (Section IV-A), sampled randomly

### VI. BASELINE AEC METHOD

We adapt a noise suppression model developed in [47] to the task of echo cancellation. Specifically, a recurrent neural network with gated recurrent units takes concatenated log power spectral features of the microphone signal and far end signal as input and outputs a spectral suppression mask. The short-time Fourier transform is computed based on 20 ms frames with a hop size of 10 ms, and a 320-point discrete Fourier transform. We use a stack of two gated recurrent unit layers, each of size 322 nodes, followed by a fully-connected layer with a sigmoid activation function. The model has 1.3 million parameters. The estimated mask is point-wise multiplied by the magnitude spectrogram of the microphone signal to suppress the far end signal. Finally, to resynthesize the enhanced signal, an inverse short-time Fourier transform is used on the phase of the microphone signal and the estimated magnitude spectrogram. We use a mean squared error loss between the clean and enhanced magnitude spectrograms. The Adam optimizer [48] with a learning rate of 0.0003 is used to train the model. The model and the inference code are available in the challenge repository.<sup>4</sup>

### VII. ONLINE SUBJECTIVE EVALUATION FRAMEWORK

We have extended the open source P.808 Toolkit [22] with methods for evaluating echo impairments in subjective tests. We followed the *Third-party Listening Test B* from ITU-T Rec. P.831 [28] and ITU-T Rec. P.832 [8] and adapted them to our use case as well as for the crowdsourcing approach based on the ITU-T Rec. P.808 [21] guidance.

A third-party listening test differs from the typical listening-only tests (according to the ITU-T Rec. P.831) in the way that listeners hear the recordings from the *center* of the connection rather in the former one in which the listener is positioned at one end of the connection [28] (see Fig. 6). Thus, the speech material should be recorded by having this concept in mind. During the test session, we use different

<sup>4</sup>[Online]. Available: <https://github.com/microsoft/AEC-Challenge/tree/main/baseline/icassp2022>

**TABLE 6** Challenge Results. See Table 8 for Descriptions of the Metrics. The Microsoft Models Were Not Officially Entered in the Challenge.

Team	Personalized	ST FE Echo MOS	DT Echo MOS	DT Other MOS	ST NE SIG MOS	ST NE BAK MOS	Overall MOS	CI MOS	Wacc ratio	Final Score
Microsoft-1*	N	4.688	4.703	4.299	4.265	4.412	4.473	0.018	0.797	0.856
ByteAudio_pAEC-18	Y	4.703	4.736	4.357	4.062	4.361	4.444	0.018	0.822	0.854
Microsoft-2*	N	4.695	4.707	4.295	4.155	4.415	4.453	0.019	0.807	0.854
ByteAudio-18	N	4.709	4.770	4.312	3.993	4.380	4.433	0.018	0.823	0.852
KuaiShou-13	N	4.703	4.679	4.087	4.099	4.252	4.364	0.019	0.780	0.831
NWPU-19	N	4.704	4.725	4.169	3.918	4.202	4.344	0.019	0.795	0.829
NWPU-10	N	4.709	4.664	4.124	3.912	4.192	4.320	0.019	0.790	0.823
Nanjing-16	N	4.619	4.640	3.926	3.920	4.149	4.251	0.020	0.755	0.803
NWPU_Elevoc-7	N	4.661	4.526	3.804	3.962	4.250	4.241	0.021	0.767	0.803
NWPU_pAEC-20	Y	4.664	4.599	3.756	3.914	4.115	4.210	0.021	0.750	0.794
NHCUSpeechLab-11	N	4.640	4.612	3.929	3.736	4.244	4.232	0.020	0.690	0.788
Harbin-1	N	4.655	4.504	3.735	3.754	4.119	4.153	0.022	0.762	0.784
CTV_Tencent-pAEC-21	Y	4.003	4.427	4.010	3.825	4.192	4.091	0.023	0.732	0.766
BJTL-5	N	4.697	4.062	3.750	4.065	4.052	4.125	0.020	0.649	0.759
EEO-9	N	4.661	4.253	3.583	3.726	3.871	4.019	0.023	0.718	0.749
baseline	N	4.535	4.283	3.479	3.883	3.887	4.013	0.023	0.649	0.736
ZhongxingTele-2	N	4.567	4.112	3.269	3.828	4.050	3.965	0.024	0.608	0.719
WuhanUni-6	N	4.388	4.238	3.511	3.406	4.031	3.915	0.024	0.663	0.718
CQUPT-14	N	4.420	4.347	3.229	3.850	3.993	3.968	0.024	0.582	0.715
Orange-17	N	3.603	3.969	3.578	3.458	3.697	3.661	0.027	0.673	0.667
CEVA-8	N	4.126	4.268	3.145	3.038	3.531	3.622	0.026	0.636	0.652
Yonsei-24	N	3.988	3.453	3.001	3.457	3.476	3.475	0.028	0.480	0.596
NoiseBusters-15	N	2.641	2.497	3.726	3.999	3.556	3.284	0.030	0.422	0.546
NoiseBusters-25	N	2.638	2.475	3.781	3.923	3.517	3.267	0.031	0.388	0.537

**TABLE 7** Results for Mobile and Desktop Recordings.

	DT Echo DMOS	ST FE Echo DMOS	ST NE SIG MOS	ST NE BAK MOS	DT Other DMOS	WAcc	M
Desktop	4.39	4.44	4.16	4.33	4.03	0.77	0.65
Mobile	4.21	4.32	3.85	4.05	3.63	0.64	0.58

**TABLE 8** Performance Metrics Used in Table 6.

Metric	Description
ST FE Echo MOS	Single talk far end echo MOS
DT Echo MOS	Double talk echo MOS
DT Other MOS	Other double talk distortion MOS
ST NE SIG MOS	Single talk near end speech signal MOS
ST NE BAK MOS	Single talk near end background MOS
WAcc	Word accuracy rate for speech recognition

combinations of single- and multi-scale Absolute Category Ratings depending on the speech sample under evaluation. We distinguish between single talk and double talk scenarios. For the near end single talk, we ask for the overall quality. For the far end single talk and double talk scenario, we ask for an echo annoyance and for impairments of other degradations in two separate questions:

- 1) How would you judge the degradation from the echo?
- 2) How would you judge other degradations (noise, missing audio, distortions, cut-outs)

Both impairments are rated on the degradation category scale (from 1: *Very annoying*, to 5: *Imperceptible*) to obtain degradation mean opinion scores (DMOS). Note that we do not use the Other degradation category for far end single talk for evaluating echo cancellation performance, since this metric mostly reflects the quality of the original far end signal.

However, we have found that having this component in the questionnaire helps increase the accuracy of echo degradation ratings (when measured against expert raters). Without the Other category, raters can sometimes assign degradations due to noise to the Echo category [29].

The setup illustrated in Fig. 5 is used to process all speech samples with all of the AECs under the study. To simplify the rating process for crowd workers, we distinguished between near end and far end single talk as well as the double talk scenarios and tried to simulate them for the test participants. In the case of near end single talk we recorded the AEC output ( $S_{out}$ ). For far end single talk, we added the output of the AEC ( $S_{out}$ ) with a delay of 600 ms to the loopback ( $R_{in}$ ) signal, yielding  $R_{in} +$  delayed  $S_{out}$ . For the listener, this simulates hearing the echo of their own speech (i.e.,  $R_{in}$  as an acoustic sidetone). For double talk the process is similar, but due to there being more speakers, simply adding the delayed AEC output ( $S_{out}$ ) would cause confusion for the test participants. To mitigate this issue, the signals are played in stereo instead, with the loopback signal ( $R_{in}$ ) played in one ear (i.e., acoustic sidetone) and the delayed output of the AEC ( $S_{out}$ ) played in the other. Fig. 6 was used to illustrate the double talk scenario to crowdworkers.

For the far end single talk scenario, we evaluate the second half of each clip to avoid initial degradations from initialization, convergence periods, and initial delay estimation. For the double talk scenario, we evaluate the final third of the audio clip.

**TABLE 9 ANOVA for the Top Challenge Entries. The Pair-Wise P-Values are Shown for the Lower Triangular Matrix Only.**

	ByteAudio-18	KuaiShou-13	NWPU-19	NWPU-10	Nanjing-16	NWPU_Elevoc-7
KuaiShou-13	0.00					
NWPU-19	0.00	0.34				
NWPU-10	0.00	0.00	0.01			
Nanjing-16	0.00	0.00	0.00	0.00		
NWPU_Elevoc-7	0.00	0.00	0.00	0.00	0.96	
NHCUSpeechLab-11	0.00	0.00	0.00	0.00	0.00	0.00

**TABLE 10 Comparison of the Top 5 Teams.**

Place	Track	Paper	Hybrid	Params	RTF	Dataset Hours	Additional Datasets	Score
MS-1	1	[11]	N	4.1 M	0.42	950		0.856
1	1	[51]	Y	1.5 M	0.28	2000	LibriSpeech	0.854
MS-2	1	[11]	N	7.5 M	0.36	950		0.854
1	2	[51]	Y	1.6 M	0.31	2000	LibriSpeech	0.852
2	1	[52]	N	32.6 M	0.19	605		0.831
3	1	[53]	Y	9.6 M	0.35	3000		0.829
4	1	[54]	Y	3.8 M	0.20	1400		0.823
5	1	[55]	N	19.2 M	0.22	600		0.803
PCC to Score			0.10	-0.54	0.67	0.19	0.49	

**TABLE 11 Comparison of Top Performing Models From 2022 and 2023 Challenge.**

	MS-1	2022 Winner [36]
Single talk far end echo MOS	4.65	4.66
Double talk echo MOS	4.68	4.68
Double talk other MOS	4.26	4.28
Algorithmic + buffering latency	20 ms	40 ms

**TABLE 12 Amount of Improvement Remaining to Get Excellent Quality Rated Speech or Perfect WAcc on This Challenge.**

Area	Headroom
Single talk far end echo MOS	0.30
Double talk echo MOS	0.26
Double talk other MOS	0.64
Single talk near end SIG MOS	0.74
Single talk near end BAK MOS	0.64
Overall MOS	0.56
WAcc	0.18

The subjective test framework is available at <https://github.com/microsoft/P.808>. A more detailed description of the test framework and its validation is given in [29].

**VIII. OBJECTIVE METRIC**

We have developed an objective perceptual speech quality metric called AECMOS. It can be used to stack rank different AEC methods based on MOS estimates with high accuracy. It is a neural network-based model that is trained using the

ground truth human ratings obtained using our online subjective evaluation framework. The audio data used to train the AECMOS model is gathered from the numerous subjective tests that we conducted in the process of improving the quality of our AECs as well as the first two AEC challenge results. The performance of AECMOS on AEC models is given in Table 5 compared with subjective human ratings on the 18 submitted models. A more detailed description of AECMOS is given in [1]. Sample code can be found on <https://aka.ms/aec-challenge>.

**IX. CHALLENGE METRIC**

The challenge performance is determined using the average of the five subjective scores described in Section VII and WAcc, all weighted equally; see (2), where  $FE$  is far end single talk,  $NE_{SIG}$  and  $NE_{BAK}$  are P.835 SIG and BAK scores for near end single talk,  $DT_{echo}$  is double talk echo, and  $DT_{other}$  is double talk other.

$$M = \frac{\frac{FE-1}{4} + \frac{NE_{SIG}-1}{4} + \frac{NE_{BAK}-1}{4} + \frac{DT_{echo}-1}{4} + \frac{DT_{other}-1}{4} + WAcc}{6} \tag{2}$$

**X. RESULTS AND ANALYSIS**

The challenge had 20 entries, 17 for the non-personalized track and 3 for the personalized track. In addition, we included two internally developed models based on [11], labeled MS-1 and MS-2. We batched all submissions into three sets:

- Near end single talk files for a MOS test (NE ST MOS).



- Far end single talk files for an Echo and Other degradation DMOS test (FE ST Echo/Other DMOS).
- Double talk files for an Echo and Other degradation DMOS test (DT Echo/Other DMOS).

The results are given in Table 6, and the analysis Of variance (ANOVA) for the top entries is given in Table 9. The 2nd and 3rd places were tied. For the ties, the winners were selected using the lower complexity model.

A high-level comparison of the top-5 entries is given in Table 10. Some observations are given below:

- There is a  $PCC=-0.54$  between the model size and the overall score. For this challenge, smaller models tend to outperform the larger models.
- There is a  $PCC=0.67$  between the RTF and the overall score. More complex models tend to outperform the less complex models.
- There is a  $PCC=0.10$  between if the model was a hybrid and the overall score. Both hybrid and deep models perform well.
- There is a  $PCC=0.19$  between the training dataset size and the overall score. Dataset size was not a significant factor in this challenge.
- There is a  $PCC=0.49$  between using additional datasets and the overall scale. Only one team added additional data (LibriSpeech [50]), though they were the first-place team [51].
- The first-place entry showed that personalized AEC did increase performance, but only by a small amount (improving the final score by 0.002).

Similar to the previous year's AEC challenge, scores for mobile recordings were significantly worse than those for desktop (Table 7). Since the data for both platforms was collected simultaneously using similar task descriptions, it indicates that on average, mobile phones pose a more challenging use case for echo cancellers.

#### A. PERFORMANCE COMPARISON OF THE ICASSP 2022 AND 2023 AEC CHALLENGE

To compare the winning model performance from the ICASSP 2022 AEC Challenge to models from this year's challenge, we apply the top-scoring model MS-1 on 2022 AEC Challenge data and use the online subjective evaluation framework to compare the results. Table 11 shows that MS-1 is statistically the same as the 2022 AEC Challenge winner [36], even though the algorithmic latency + buffering latency for MS-1 is 20 ms and for [36] 40 ms. In studies with the CRUSE [56] noise suppression model which MS-1 is based on, changing the frame size from 20 ms to 40 ms increased DNSMOS OVRL by 0.1. In addition, changing the frame size of MS-1 from 20 ms to 10 ms decreased DNSMOS OVRL by 0.07. Therefore, we conclude that MS-1 should be significantly better than [36] if that model also had an algorithmic + buffering latency of 20 ms.

#### XI. CONCLUSION

This latest AEC challenge induced lower algorithmic latency + buffering latency requirements and added a personalized track. The performance of the top models is exceptional, though it shows there is still a lot of headroom to improve, especially in the double talk other, single near end, and WAcc metrics (see Table 12). We are optimistic that the personalized enrollment data can improve these areas much more than was shown in this challenge, which is a good area for future research. In addition, even lower latency requirements are needed for a telecommunication system to achieve end-to-end latencies of less than 50 ms, which is the just-noticeable difference when latency impacts conversations [57]. End-to-end latencies significantly above 50 ms have been shown to be correlated to lower participation in group meetings [58]. To achieve this goal the algorithmic latency + buffering latency should be less than 5 ms, which is another good area for future work. Finally, improving mobile performance is still an area of research, as the gap between mobile and desktop scenarios is still quite significant.

#### REFERENCES

- [1] M. Purin, S. Sootla, M. Sponza, A. Saabas, and R. Cutler, "AECMOS: A speech quality assessment metric for echo impairment," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2022, pp. 901–905.
- [2] *IEEE 1329-2010 Method for Measuring Transmiss. Perform. of Hands-free Telephone Sets*, Standard IEEE 1329-2010, 2010.
- [3] A. Fazel, M. El-Khamy, and J. Lee, "CAD-AEC: Context-aware deep acoustic echo cancellation," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2020, pp. 6919–6923.
- [4] M. Halimeh and W. Kellermann, "Efficient multichannel nonlinear acoustic echo cancellation based on a cooperative strategy," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2020, pp. 461–465.
- [5] L. Ma, H. Huang, P. Zhao, and T. Su, "Acoustic echo cancellation by combining adaptive digital filter and recurrent neural network," 2020, *arXiv:2005.09237*.
- [6] H. Zhang, K. Tan, and D. Wang, "Deep learning for joint acoustic echo and noise cancellation with nonlinear distortions," in *Proc. Interspeech*, 2019, pp. 4255–4259.
- [7] G. Enzner, H. Buchner, A. Favrot, and F. Kuech, "Acoustic echo control," in *Academic Press Library in Signal Processing*, vol. 4. Amsterdam, The Netherlands: Elsevier, 2014, pp. 807–877.
- [8] *Subjective Perform. Eval. of Hands-free Terminals*, Rec. ITU-T P.832, International Telecommunications Union, Geneva, Switzerland, 2000.
- [9] A. R. Avila, H. Gamper, C. Reddy, R. Cutler, I. Tashev, and J. Gehrke, "Non-intrusive speech quality assessment using neural networks," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2019, pp. 631–635.
- [10] C. K. Reddy et al., "Deep noise suppression challenge: Datasets, subjective testing framework, and challenge results," in *Proc. Interspeech*, 2020, pp. 2492–2496.
- [11] E. Indenbom, N.-C. Ristea, A. Saabas, T. Parnamaa, J. Guzvina, and R. Cutler, "DeepVQE: Real time deep voice quality enhancement for joint acoustic echo cancellation, noise suppression and dereverberation," in *Proc. Interspeech*, 2023, pp. 3819–3823.
- [12] K. Sridhar et al., "ICASSP 2021 Acoustic Echo Cancellation Challenge: Datasets, testing framework, and results," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2021, pp. 151–155.
- [13] R. Cutler et al., "Interspeech 2021 Acoustic Echo Cancellation Challenge," in *Proc. Interspeech*, 2021, pp. 4748–4752.
- [14] R. Cutler et al., "ICASSP 2022 Acoustic Echo Cancellation Challenge," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2022, pp. 9107–9111.

- [15] H. Dubey et al., "ICASSP 2022 Deep Noise Suppression Challenge," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2022, pp. 9271–9275.
- [16] TIA, "TIA-920: Transmission requirements for wideband digital wireline telephones," Standard TIA, 2002.
- [17] *Influence of Nat. Syst. on Stability and Talker Echo in Int. Connections*, Rec. ITU-T G.122, International Telecommunications Union, Geneva, Switzerland, 2012.
- [18] *Talker Echo and Its Control*, Rec. ITU-T G.131, International Telecommunications Union, Geneva, Switzerland, 2003.
- [19] *Digit. Netw. Echo Cancellers*, Rec. ITU-T G.168, International Telecommunications Union, Geneva, Switzerland, 2012.
- [20] *Methods for Subjective Determination of Transmiss. Qual.*, Rec. ITU-T P.800, International Telecommunications Union, Geneva, Switzerland, 1996.
- [21] *Subjective Eval. of Speech Qual. with a Crowdsourcing Approach*, Rec. ITU-T P.808, International Telecommunications Union, Geneva, Switzerland, 2018.
- [22] B. Naderi and R. Cutler, "An open source implementation of ITU-T recommendation P.808 with validation," in *Proc. INTERSPEECH*, 2020, pp. 2862–2866.
- [23] *Subjective Test Methodol. for Evaluating Speech Commun. Syst. that Include Noise Suppression Algorithm*, Rec. ITU-T P.835, International Telecommunications Union, Geneva, Switzerland, 2003.
- [24] B. Naderi and R. Cutler, "Subjective evaluation of noise suppression algorithms in crowdsourcing," in *Proc. INTERSPEECH*, 2021, pp. 2132–2136.
- [25] *Extension of ITU-T P.863 For Multi-Dimensional Assessment of Degradations in Telephony Speech signals up to Full-Band*, Rec. ITU-T P.863, International Telecommunications Union, Geneva, Switzerland, 2022.
- [26] *Subjective Diagn. Test Method for Conversational Speech Qual. Anal.*, Rec. ITU-T P.804, International Telecommunications Union, Geneva, Switzerland, 2017.
- [27] B. Naderi, R. Cutler, and N.-C. Ristea, "Multi-dimensional speech quality assessment in crowdsourcing," 2023, [arXiv:2309.07385](https://arxiv.org/abs/2309.07385).
- [28] *Subjective Perform. Eval. of Netw. Echo Cancellers*, Rec. ITU-T P.831, International Telecommunications Union, Geneva, Switzerland, 1998.
- [29] R. Cutler, B. Naderi, M. Loide, S. Sootla, and A. Saabas, "Crowdsourcing approach for subjective evaluation of echo impairment," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2021, pp. 406–410.
- [30] J.-M. Valin, S. Tanneti, K. Helwani, U. Isik, and A. Krishnaswamy, "Low-complexity, real-time joint neural echo control and speech enhancement based on PercepNet," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2021, pp. 7133–7137.
- [31] Z. Wang, Y. Na, Z. Liu, B. Tian, and Q. Fu, "Weighted recursive least square filter and neural network based residual ECHO suppression for the AEC-Challenge," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2021, pp. 141–145.
- [32] N. L. Westhausen and B. T. Meyer, "Acoustic echo cancellation with the dual-signal transformation LSTM network," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2020, pp. 7138–7142.
- [33] R. Peng, L. Cheng, C. Zheng, and X. Li, "Acoustic echo cancellation using deep complex neural network with nonlinear magnitude compression and phase information," in *Proc. INTERSPEECH*, 2021, pp. 4768–4772.
- [34] S. Zhang, Y. Kong, S. Lv, Y. Hu, and L. Xie, "F-T-LSTM based complex network for joint acoustic echo cancellation and speech enhancement," Jun. 2021, [arXiv:2106.07577](https://arxiv.org/abs/2106.07577).
- [35] L. Pfeifenberger, M. Zoehrer, and F. Pernkopf, "Acoustic echo cancellation with cross-domain learning," in *Proc. INTERSPEECH*, 2021, pp. 4753–4757.
- [36] G. Zhang, L. Yu, C. Wang, and J. Wei, "Multi-scale temporal frequency convolutional network with axial attention for speech enhancement," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2022, pp. 9122–9126.
- [37] H. Zhao et al., "A deep hierarchical fusion network for fullband acoustic echo cancellation," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2022, pp. 9112–9116.
- [38] S. Zhang et al., "Multi-task deep residual echo suppression with echo-aware loss," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2022, pp. 9127–9131.
- [39] M. Karjalainen, P. Antsalo, A. Mäkiranta, T. Peltonen, and V. Välimäki, "Estimation of modal decay parameters from noisy response measurements," *J. Audio Eng. Soc.*, vol. 50, May 2001, Art. no. 5290.
- [40] H. Gamper and I. J. Tashev, "Blind reverberation time estimation using a convolutional neural network," in *Proc. 16th Int. Workshop Acoustic Signal Enhancement*, 2018, pp. 136–140.
- [41] C. Valentini-Botinhao, X. Wang, S. Takaki, and J. Yamagishi, "Speech enhancement for a noise-robust text-to-speech synthesis system using deep recurrent neural networks," in *Proc. Interspeech*, 2016, pp. 352–356.
- [42] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [43] J. S. Garofolo, L. F. Lamel, W. M. Fisher, J. G. Fiscus, D. S. Pallett, and N. L. Dahlgren, "DARPA TIMIT acoustic-phonetic continuous speech corpus CD-ROM TIMIT," Feb. 1993.
- [44] J. F. Gemmeke et al., "Audio set: An ontology and human-labeled dataset for audio events," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2017, pp. 776–780.
- [45] J. Thiemann, N. Ito, and E. Vincent, "The diverse environments multi-channel acoustic noise database (DEMAND): A database of multichannel environmental noise recordings," *J. Acoustical Soc. Amer.*, vol. 133, no. 5, pp. 3591–3591, 2013.
- [46] C. M. Lee, J. W. Shin, and N. S. Kim, "DNN-Based residual echo suppression," in *Proc. 16th Annu. Conf. Int. Speech Commun. Assoc.*, 2015, pp. 1775–1779.
- [47] Y. Xia, S. Braun, C. K. A. Reddy, H. Dubey, R. Cutler, and I. Tashev, "Weighted speech distortion losses for neural-network-based real-time speech enhancement," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2020, pp. 871–875.
- [48] D. P. Kingma and J. Ba, "ADAM: A method for stochastic optimization," in *Proc. Int. Conf. Learn. Representations*, 2015.
- [49] M. G. Kendall, "A new measure of rank correlation," *Biometrika*, vol. 30, no. 1/2, pp. 81–93, 1938.
- [50] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2015, pp. 5206–5210.
- [51] Z. Chen et al., "A progressive neural network for acoustic echo cancellation," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2023, pp. 1–2.
- [52] H. Zhao et al., "A low-latency deep hierarchical fusion network for fullband acoustic echo cancellation," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2023, pp. 1–2.
- [53] Z. Zhang et al., "Two-step band-split neural network approach for full-band residual echo suppression," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2023, pp. 1–2.
- [54] J. Sun, D. Luo, Z. Li, J. Li, Y. Ju, and Y. Li, "Multi-task sub-band network for deep residual echo suppression," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2023, pp. 1–2.
- [55] W. Xu and Z. Guo, "TaylorAECNet: A Taylor style neural network for full-band echo cancellation," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2023, pp. 1–2.
- [56] S. Braun, H. Gamper, C. K. A. Reddy, and I. Tashev, "Towards efficient models for real-time deep noise suppression," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2021, pp. 656–660.
- [57] N. Kitawaki and K. Itoh, "Pure delay effects on speech quality in telecommunications," *IEEE J. Sel. Areas Commun.*, vol. 9, no. 4, pp. 586–593, May 1991.
- [58] R. Bothra et al., "Don't forget the user: It's time to rethink network measurements," in *Proc. 22nd ACM Workshop Hot Topics Netw.*, 2023, pp. 109–116.



**ROSS CUTLER** (Member, IEEE) received the B.S. degree in mathematics, computer science, and physics and the Ph.D. degree in computer science from the University of Maryland, College Park, MD, USA, in 1992 and 2000, respectively. Since 2000, he has been with Microsoft (MS), starting as a Researcher with Microsoft Research. He is currently a Distinguished Engineer with the IC3 Group, MS, where he manages the IC3-AI team of applied scientists and software engineers with a focus on improving Teams/Skype audio/video quality and reliability, and enabling new functionality with AI. He has authored or coauthored more than 80 peer-reviewed conference and journal papers and has more than 100 granted patents in his research interests, which include computer vision, machine learning, signal processing, acoustics, optics, and VoIP.



**ANDO SAABAS** received the Ph.D. degree in computer science from the Tallinn University of Technology, Tallinn, Estonia, in 2008. He is currently a Principal ML Scientist Manager with the IC3-AI Team, Microsoft, focusing on improving call quality and reliability in Skype and Teams. He was with Skype and Microsoft on social network analysis, causal inference, and call quality and reliability improvements.



**JEGOR GUŽVIN** is currently a Senior Software Development Engineer with the IC3-AI Group, Microsoft. His research interests include integrating ML-based solutions into Teams/Skype and optimizing them for the best performance.



**TANEL PÄRNAMAA** received the M.Sc. degree in mathematical statistics with the University of Tartu, Tartu, Estonia, in 2015. He is currently a Machine Learning Scientist with the IC3-AI Group, Microsoft, focusing on improving the real-time communications stack. His education includes an exchange year with the University of Toronto, Toronto, ON, Canada, and Research visits to the Montreal Institute for Learning Algorithms (Mila) and Wellcome Trust Sanger Institute. His research interests include speech enhancement and visual



**HANNES GAMPER** (Member, IEEE) received the master's degrees in sound engineering from the University of Technology and the University of Music and Performing Arts, Graz, Austria, and the Ph.D. degree in media technology from Aalto University, Helsinki, Finland. He is currently a Researcher with the Audio and Acoustics Research Group, Microsoft Research, Redmond, WA, USA.

perception.



**MARJU PURIN** received the Ph.D. degree in mathematics from her Syracuse University, Syracuse, NY, USA, in 2011. She was a Mathematics Professor in United States, for about a decade. In 2020, she joined Microsoft, where she is currently a Machine Learning Scientist with the IC3-AI Team, focusing on improving deep models for real-time communication. She has authored or coauthored papers in deep learning, focusing on echo cancellation and objective metrics for model evaluation. Her research focuses on homological



**SEBASTIAN BRAUN** (Senior Member, IEEE) received the M.Sc. degree in electrical and sound engineering from Technical University (TU), Berlin, Germany, and the University of Arts (KU), Graz, Austria, in 2012, and the Ph.D. degree from International Audio Laboratories, a joint institution of the University of Erlangen-Nuremberg and the Fraunhofer IIS, Erlangen, Germany, in 2018. He is currently a Researcher with the Audio and Acoustics Research Group, Microsoft Research.

algebra.

He has authored or coauthored more than 20 peer-reviewed conference and journal papers. His research interests include audio and acoustic signal processing, speech enhancement and separation, microphone array processing, and spatial and binaural audio processing. He was the recipient of Best Student Paper Award at IWAENC 2014.



**EVGENII INDENBOM** received the M.S. degree in space plasma physics from the Moscow Institute for Physics and Technology, Dolgoprudny, Russia, in 1995. He is currently a Machine Learning Scientist with the IC3-AI Team, Microsoft. He transitioned to software engineering and computer science, with a focus on computer linguistics and natural language processing. Since 2021, he has been with Microsoft on enhancing call quality and reliability with Skype and Teams.



**ROBERT AICHNER** received the Ph.D. degree in electrical engineering in the area of audio signal processing from the University of Erlangen-Nuremberg, Erlangen, Germany. Since 2007, he has been with Microsoft, starting as an audio signal processing developer and then holding multiple roles in program management ranging from a focus on the algorithms in the real-time communications stack to also working in the user interface/user experience area. He is currently a Partner Group Program Manager with the Intelligent Communication and Conversation Cloud (IC3) Group, Microsoft, responsible for the areas of machine learning and call quality. He has authored or coauthored more than 30 conference papers, journal papers, and book chapters.



**NICOLAE-CĂTĂLIN RISTEA** received the Graduation degree (as Valedictorian) from the Faculty of Electronics, Telecommunications and Information Technology, University Politehnica of Bucharest, Bucharest, Romania, in 2019, and the M.Sc. degree in the image processing field from the University Politehnica of Bucharest, where he is currently working toward the Ph.D. degree. In 2021, he joined Microsoft, where he works on the audio quality on the Teams platform. He is also the first author of multiple papers accepted at top-tier

conferences and journals, such as *CVPR*, *INTERSPEECH*, and *IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE*. He was on the organization team for multiple challenges at *INTERSPEECH* and Proceedings of the International Conference on Acoustics, Speech, and Signal Processing. His research interests include AI, computer vision, machine learning, signal processing, and deep learning.