

ICASSP 2022 ACOUSTIC ECHO CANCELLATION CHALLENGE

Ross Cutler, Ando Saabas, Tanel Parnamaa, Marju Purin, Hannes Gamper, Sebastian Braun, Karsten Sørensen, Robert Aichner

Microsoft Corporation, Redmond, USA
firstname.lastname@microsoft.com

ABSTRACT

The ICASSP 2022 Acoustic Echo Cancellation Challenge is intended to stimulate research in acoustic echo cancellation (AEC), which is an important area of speech enhancement and still a top issue in audio communication. This is the third AEC challenge and it is enhanced by including mobile scenarios, adding speech recognition word accuracy rate as a metric, and making the audio 48 kHz. 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 online objective metric service 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 rate.

Index Terms: acoustic echo cancellation, deep learning, single talk, double talk, subjective test

1. 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 [1].

With the advent of deep learning techniques, many supervised learning algorithms for AEC have shown better performance compared to their classical counterparts, e.g., [2, 3, 4]. 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) [4, 5] 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.

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 5).

	PCC	SRCC
ERLE	0.31	0.23
PESQ	0.67	0.57

Most AEC publications use objective measures such as *echo return loss enhancement* (ERLE) [6] and *perceptual evaluation of speech quality* (PESQ) [7]. ERLE in dB is defined as:

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

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 [8]. 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 [9]. An initial test set will be released for the researchers to use during development and a blind test set near the end, which will be 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.

This is the third AEC challenge we have conducted. The first challenge was held at ICASSP 2021 [10] and the second at INTERSPEECH 2021 [11]. These challenges had 31 participants with entries ranging from pure deep models, hybrid linear AEC + deep echo suppression, and DSP methods. The results show that the deep and hybrid models far outperformed DSP methods, with the latest winners being both pure deep and hybrid models. However, there is still

much room for improvement. To improve the challenge and further stimulate research in this area we have made the following changes:

- The dataset has increased from 5,000 devices and environments to 10,000 to provide additional training data.
- Mobile phone scenarios are now included, which are an important area that is even more challenging than desktop or notebook computers. 50% of the blind test set were mobile devices, and 50% were desktop devices.
- The Microsoft Speech Recognizer’s Word Accuracy rate (WAcc) is used as a metric in the challenge, as many scenarios include speech recognition and the AEC should not degrade WAcc. $WAcc = 1 - \text{Word Error Rate}$.
- The test sets are now 48 kHz, which is an important requirement for many scenarios.

The training dataset is described in Section 2, and the test set in Section 3. We describe a DNN-based AEC method in Section 4. The online subjective evaluation framework is discussed in Section 5, and the objective service in Section 6. The challenge metric is given in Section 7 and the challenge rules are described in <https://aka.ms/aec-challenge>.

2. TRAINING DATASETS

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

2.1. 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
5. Double talk, echo path change
6. Sweep signal for RT60 estimation

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. [12] and shown in Figure 1. For 1251 mobile environments the RT60 distribution shown was estimated blindly from speech recordings [13].

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 Figure 2). 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 (DRC)); it can also include Audio DSP on the send signal, such as AEC and noise suppression.

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 [14]. This corpus consists of short single speaker speech segments (1 to 3 seconds). We used a *long short term memory* (LSTM) 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 to 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 [15] sentence list. Approximately 10 seconds of audio is recorded while the users are reading.

2.2. 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 [9] from the LibriVox project¹ as source clips to sample far end and near end signals. The LibriVox project is a collection of public domain audiobooks read by volunteers. [9] 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 [16], Freesound² and DEMAND [17] 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. Of the selected far end and near end speakers, 71% and 67% are male, respectively. To generate an echo, we convolve a randomly chosen room impulse response from a large internal database with the far end signal. The room impulse responses are generated by using Project Acoustics technology³ and the RT60 ranges from 200 ms to 1200 ms. In 80% of the cases, the far end signal is processed by

¹<https://librivox.org>

²<https://freesound.org>

³<https://www.aka.ms/acoustics>

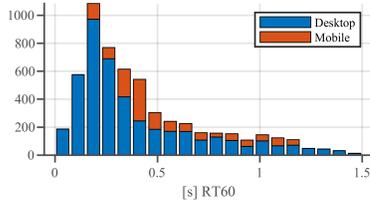


Fig. 1. Distribution of reverberation time (RT60).

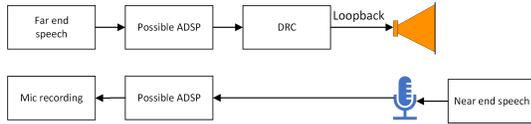


Fig. 2. The custom recording application recorded the loopback and microphone signals.

a nonlinear function to mimic loudspeaker distortion. For example, the transformation can be clipping the maximum amplitude, using a sigmoidal function as in [18], 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 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.

3. 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 datasets 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 (2.1), sampled randomly

4. BASELINE AEC METHOD

We adapt a noise suppression model developed in [19] 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 with 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 with a learning rate of 0.0003 is used to train the model. The model and the inference code is available in the challenge repository.⁴

5. ONLINE SUBJECTIVE EVALUATION FRAMEWORK

We have extended the open source P.808 Toolkit [20] with methods for evaluating the echo impairments in subjective tests. We followed the *Third-party Listening Test B* from ITU-T Rec. P.831 [21] and ITU-T Rec. P.832 [22] and adapted them to our use case as well as for the crowdsourcing approach based on the ITU-T Rec. P.808 [23] 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 [21]. Thus, the speech material should be recorded by having this concept in mind. During the test session, we use different 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 questions⁵. 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.

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 roughly the final third of the audio clip.

The subjective test framework with an AEC extension is available at <https://github.com/microsoft/P.808>. A more detailed description of the test framework and its validation is given in [24].

6. AZURE SERVICE OBJECTIVE METRIC

We have developed an objective perceptual speech quality metric called AECMOS. It can be used to stack rank different AEC methods

⁴<https://github.com/microsoft/AEC-Challenge/tree/main/baseline/icassp2022>

⁵Question 1: How would you judge the degradation from the echo? Question 2: How would you judge other degradations (noise, missing audio, distortions, cut-outs)?

Table 2. AECMOS PCC and SRCC

Scenario	PCC	SRCC
Far end single talk echo DMOS	0.828	0.719
Near end single talk MOS	0.843	0.856
Double talk echo DMOS	0.882	0.766
Double talk other DMOS	0.929	0.913

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 2 compared with subjective human ratings on the 18 submitted models. We note that this model had not seen any mobile nor fullband data during training. The next version of AECMOS will have mobile and fullband data in its training data. A more detailed description of AECMOS is given in [25]. Sample code and details of the evaluation API can be found on <https://aka.ms/aec-challenge>.

7. CHALLENGE METRIC

The challenge performance is determined using the average of the four subjective scores described in Section 5 and WAcc, all weighted equally. Specifically:

$$M = \frac{(FE_{ST-1}) + (NE_{ST-1}) + (DT_{echo-1}) + (DT_{other-1}) + WAcc}{5}$$

where FE_{ST} is far end single talk, NE_{ST} is near end single talk, DT_{echo} is double talk echo, and DT_{other} is double talk other.

8. RESULTS

We received 18 submissions for the challenge. Each team submitted processed files from the blind test set (see Section 3). 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 shown in Figure 3. The score differences in near end, echo, double talk, and WAcc highlight the importance of evaluating all scenarios, since in many cases, performance in one scenario comes at a cost in another scenario. The PCC of WAcc and the mean subjective scores is 0.85, which helps motivate why WAcc needs to be measured.

For the top performing teams, we ran an ANOVA test to determine statistical significance (see <https://aka.ms/aec-challenge>). The 2nd and 3rd, and 5th and 6th places were tied. For the ties, the winners were selected using the lower complexity model.

A high-level comparison of the top 5 performing models is given in Table 3. Real-time factor is the run-time / frame time on an Intel Core i5 quad core 2.4 GHz CPU or equivalent. The 1st place model [26] also won the ICASSP 2022 DNS Challenge [27], providing the only model that didn't induce SIG [28] distortion in that challenge. It is a hybrid model but is unique in that it uses the linear AEC only to condition the DNN, not filtering the audio. Three of the top 5

Team	ST NE MOS	ST FE Echo DMOS	DT Echo DMOS	DT Other DMOS	Subjective	WAcc	M
19	4.40	4.81	4.75	4.43	4.60	0.81	0.88
4	4.32	4.81	4.74	4.25	4.53	0.80	0.87
13	4.24	4.80	4.71	4.26	4.50	0.80	0.86
29	4.23	4.70	4.67	4.07	4.42	0.82	0.84
22	4.32	4.76	4.65	4.03	4.44	0.77	0.82
20	4.19	4.74	4.66	3.97	4.39	0.68	0.82
9	4.24	4.37	4.60	4.14	4.34	0.73	0.82
8	4.22	4.77	4.25	3.88	4.28	0.74	0.81
21	4.06	4.52	4.57	4.00	4.29	0.79	0.81
17	4.07	4.76	4.23	3.87	4.23	0.77	0.80
10	4.12	4.69	4.22	3.84	4.22	0.79	0.80
32	4.23	4.43	4.28	3.76	4.17	0.76	0.78
7	4.18	4.59	4.35	3.54	4.16	0.71	0.76
baseline	4.15	4.56	4.12	3.56	4.10	0.65	0.75
27	4.01	4.33	4.01	3.68	4.01	0.66	0.74
33	4.16	3.98	4.08	3.63	3.96	0.70	0.74
23	4.15	4.65	3.82	3.50	4.03	0.75	0.74
24	4.09	2.97	3.49	3.85	3.60	0.66	0.64
11	3.89	3.02	2.46	3.39	3.19	0.61	0.44

Fig. 3. AEC challenge results

	ST NE MOS	ST FE Echo DMOS	DT Echo DMOS	DT Other DMOS	WAcc	M
Desktop	4.13	4.61	4.48	4.20	0.80	0.83
Mobile	4.23	4.38	4.09	3.45	0.65	0.74

Fig. 4. Results for desktop and mobile recordings

teams use linear AEC's and DNN's, and all 5 use the STFT domain. In addition, all 5 models perform noise suppression in addition to AEC. There is a wide range of model sizes and complexities, and it wasn't necessary to use external datasets to do well in the challenge.

When comparing the results between mobile and desktop recordings (Figure 4), we observe relatively similar scores for the near end single talk category, but significantly lower scores for mobile in echo categories, especially for double talk. The difference is highest in the double talk degradation category, where the score for mobile recordings is lower by 0.75 MOS. One reason for this is that in mobile devices, the loudspeaker is closer to the microphone, so the signal-to-echo ratio in these devices is lower on average.

9. CONCLUSIONS

While the results of this challenge continue to improve over previous challenges, there is still significant room for improvement, especially with the mobile scenario. We hope this challenge, dataset, test set, and test framework stimulate research in this important area of speech enhancement.

Table 3. Comparison of the top 5 teams

Place	Paper	Hybrid	Params	Real-time Factor	Additional Datasets
1	[26]	Y*	1.5 M	0.60	N
2	[29]	Y	17.4 M	0.10	Y
3	[30]	Y	4.8 M	0.20	Y
4	[31]	Y	55.5 M	0.30	N
5	[32]	N	4.3 M	0.02	Y

10. REFERENCES

- [1] “IEEE 1329-2010 Standard: Method for Measuring Transmission Performance of Handsfree Telephone Sets,” *IEEE*, 2010.
- [2] A. Fazel, M. El-Khamy, and J. Lee, “CAD-AEC: Context-Aware Deep Acoustic Echo Cancellation,” in *ICASSP*, May 2020, pp. 6919–6923, ISSN: 2379-190X.
- [3] M. Halimeh and W. Kellermann, “Efficient Multichannel Non-linear Acoustic Echo Cancellation Based on a Cooperative Strategy,” in *ICASSP*, 2020.
- [4] L. Ma, H. Huang, P. Zhao, and T. Su, “Acoustic Echo Cancellation by Combining Adaptive Digital Filter and Recurrent Neural Network,” in *arXiv:2005.09237*, May 2020, arXiv: 2005.09237.
- [5] H. Zhang, K. Tan, and D. Wang, “Deep Learning for Joint Acoustic Echo and Noise Cancellation with Nonlinear Distortions,” in *INTERSPEECH 2019*. Sept. 2019, pp. 4255–4259, ISCA.
- [6] G. Enzner, H. Buchner, A. Favrot, and F. Kuech, “Acoustic Echo Control,” in *Academic Press Library in Signal Processing*, vol. 4, pp. 807–877. Elsevier, 2014.
- [7] “ITU-T Recommendation P.862: Perceptual evaluation of speech quality (PESQ): An objective method for end-to-end speech quality assessment of narrow-band telephone networks and speech codec,” 2001.
- [8] A. R. Avila, H. Gamper, C. Reddy, R. Cutler, I. Tashev, and J. Gehrke, “Non-intrusive Speech Quality Assessment Using Neural Networks,” in *ICASSP*, Brighton, United Kingdom, May 2019, pp. 631–635, IEEE.
- [9] C. K. Reddy, V. Gopal, R. Cutler, E. Beyrami, R. Cheng, H. Dubey, S. Matushevych, R. Aichner, A. Aazami, S. Braun, P. Rana, S. Srinivasan, and J. Gehrke, “The INTERSPEECH 2020 Deep Noise Suppression Challenge: Datasets, Subjective Testing Framework, and Challenge Results,” in *INTERSPEECH 2020*. Oct. 2020, pp. 2492–2496, ISCA.
- [10] K. Sridhar, R. Cutler, A. Saabas, T. Parnamaa, M. Loide, H. Gamper, S. Braun, R. Aichner, and S. Srinivasan, “ICASSP 2021 Acoustic Echo Cancellation Challenge: Datasets, Testing Framework, and Results,” in *ICASSP*, 2021.
- [11] R. Cutler, A. Saabas, T. Parnamaa, M. Loide, S. Sootla, M. Purin, H. Gamper, S. Braun, K. Sorensen, R. Aichner, and S. Srinivasan, “INTERSPEECH 2021 Acoustic Echo Cancellation Challenge,” in *INTERSPEECH*, June 2021.
- [12] M. Karjalainen, P. Antsallo, A. Mäkivirta, T. Peltonen, and V. Välimäki, “Estimation of Modal Decay Parameters from Noisy Response Measurements,” *Journal of the Audio Engineering Society*, vol. 50, May 2001.
- [13] H. Gamper and I. J. Tashev, “Blind Reverberation Time Estimation Using a Convolutional Neural Network,” in *2018 16th International Workshop on Acoustic Signal Enhancement (IWAENC)*, Sept. 2018, pp. 136–140.
- [14] 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 *INTERSPEECH*, 2016, vol. 08-12-Sept, pp. 352–356, ISSN: 19909772.
- [15] 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.
- [16] J. F. Gemmeke, D. P. W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, “Audio Set: An ontology and human-labeled dataset for audio events,” in *ICASSP*, New Orleans, LA, Mar. 2017, pp. 776–780, IEEE.
- [17] J. Thiemann, N. Ito, and E. Vincent, “The Diverse Environments Multi-channel Acoustic Noise Database (DEMAND): A database of multichannel environmental noise recordings,” Montreal, Canada, 2013.
- [18] C. M. Lee, J. W. Shin, and N. S. Kim, “DNN-Based Residual Echo Suppression,” 2015.
- [19] 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 *ICASSP*, Barcelona, Spain, May 2020, pp. 871–875, IEEE.
- [20] B. Naderi and R. Cutler, “An Open Source Implementation of ITU-T Recommendation P.808 with Validation,” *INTERSPEECH*, pp. 2862–2866, Oct. 2020.
- [21] “ITU-T P.831: Subjective performance evaluation of network echo cancellers ITU-T P-series Recommendations,” 1998.
- [22] “ITU-T Recommendation P.832: Subjective performance evaluation of hands-free terminals,” *International Telecommunication Union*, 2000.
- [23] “ITU-T P.808: Subjective evaluation of speech quality with a crowdsourcing approach,” Tech. Rep., 2018.
- [24] R. Cutler, B. Naderi, M. Loide, S. Sootla, and A. Saabas, “Crowdsourcing approach for subjective evaluation of echo impairment,” in *ICASSP*, 2020.
- [25] M. Purin, S. Sootla, M. Sponza, A. Saabas, and R. Cutler, “AECMOS: A speech quality assessment metric for echo impairment,” in *ICASSP*, 2022.
- [26] G. Zhang, L. Yu, C. Wang, and J. Wei, “Multi-scale temporal frequency convolutional network with axial attention for speech enhancement,” *ICASSP*, 2022.
- [27] H. Dubey, V. Gopal, R. Cutler, A. Aazami, S. Matushevych, S. Braun, S. E. Eskimez, M. Thakker, T. Yoshioka, H. Gamper, and R. Aichner, “ICASSP 2022 Deep Noise Suppression Challenge,” in *ICASSP*, 2022.
- [28] B. Naderi and R. Cutler, “Subjective Evaluation of Noise Suppression Algorithms in Crowdsourcing,” in *INTERSPEECH*, 2021.
- [29] H. Zhao, N. Li, R. Han, L. Chen, X. Zheng, C. Zhang, L. Guo, and B. Yu, “A deep hierarchical fusion network for fullband acoustic echo cancellation,” in *ICASSP*, 2022.
- [30] S. Zhang, Z. Wang, J. Sun, Y. Fu, B. Tian, Q. Fu, and L. Xie, “Multi-task deep residual echo suppression with echo-aware loss,” in *ICASSP*, 2022.
- [31] X. Sun, C. Cao, Q. Li, L. Wang, and F. Xiang, “Explore relative and context information with transformer for joint acoustic echo cancellation and speech enhancement,” in *ICASSP*, 2022.
- [32] F. Cui, L. Guo, W. Li, P. Gao, and Y. Wang, “Multi-scale refinement network based acoustic echo cancellation,” in *ICASSP*, 2022.