

# TEXTUAL ECHO CANCELLATION

Shaojin Ding Ye Jia Ke Hu Quan Wang

Google LLC, USA

{ shaojinding, jiaye, huk, quanw }@google.com

## ABSTRACT

In this paper, we propose Textual Echo Cancellation (TEC) — a framework for cancelling the text-to-speech (TTS) playback echo<sup>1</sup> from overlapping speech recordings. Such a system can largely improve speech recognition performance and user experience for intelligent devices such as smart speakers, as the user can talk to the device while the device is still playing the TTS signal responding to the previous query. We implement this system by using a novel sequence-to-sequence model with multi-source attention that takes both the microphone mixture signal and source text of the TTS playback as inputs, and predicts the enhanced audio. Experiments show that the textual information of the TTS playback is critical to enhancement performance. Besides, the text sequence is much smaller in size compared with the raw acoustic signal of the TTS playback, and can be immediately transmitted to the device or ASR server even before the playback is synthesized. Therefore, our proposed approach effectively reduces Internet communication and latency compared with alternative approaches such as acoustic echo cancellation (AEC).

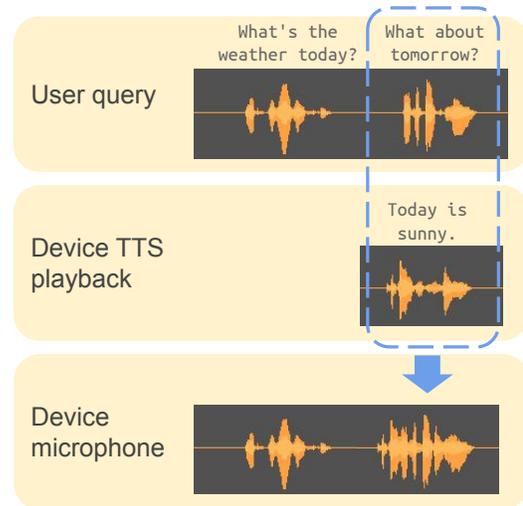
**Index Terms**— echo cancellation, multi-source attention, sequence-to-sequence model

## 1. INTRODUCTION

Intelligent devices with speech interaction features have become popular in recent years, such as mobile devices and smart home speakers. In a typical user interaction, the user first issues a query to the device, then the device responds with synthesized speech; after hearing the response, the user may issue the next query. However, in some scenarios, the user may want to self-correct the previous query or impatiently issue a new query before the device finishes playing the synthesized response. When the user and the device talk at the same time, the acoustic echoes become a challenge for accurate speech recognition. For example, the user could first ask “What’s the weather today”. While the smart speaker plays synthesized response “Today is sunny”, the user may interrupt impatiently with “What about tomorrow”, as illustrated in Fig. 1. In this case, it is very difficult for the device to correctly recognize the query “What about tomorrow” due to the overlaps between query and TTS playback echo.

One of the most straightforward and well-developed approaches to solve this problem is acoustic echo cancellation (AEC) [1, 2, 3, 4, 5, 6, 7]. Conventional signal-processing based AEC approaches [1, 2, 6, 7] usually use adaptive filtering to estimate the echo path between the speaker and the microphone, and then the mixture signal from microphone is combined with the estimated echo path to

<sup>1</sup>TTS playback denotes the synthesized TTS voice, and TTS playback echo denotes the reverberated TTS voice captured by the microphone.



**Fig. 1.** Acoustic echoes caused by TTS playback overlapping with user query.

produce the enhanced signal. More recently, model based AEC approaches [3, 4, 5] have been shown to significantly boost the performance. These models take the microphone mixture signal along with the echo signal as the input and are trained to predict a mask, which is then applied to the microphone mixture signal to produce the enhanced signal.

However, for many practical applications, conventional AEC models are subject to a number of restrictions. First, under the intelligent device settings, there is no way to acquire the actual TTS playback echo since the room configurations are unknown and vary from user to user. As an approximation of the echo, we can directly use the TTS playback. However, there exists a mismatch between the actual echo and the TTS playback, which makes the AEC performance vary in real world conditions. Second, as speech synthesis is a computationally intensive task, TTS playback is usually streamed from a TTS server to the user device to achieve lower latency. However, most of existing AEC systems depend on entire TTS playback when producing the enhanced signal. If these systems run on the device, they cannot start running until the end of TTS streaming, which may introduce significantly high latency to ASR. If AEC and ASR are implemented on servers, then AEC would require the TTS service to stream the TTS playback to the AEC server as side input, which introduces additional Internet traffic.

Other potential solutions to this problem are speech separation [8, 9, 10, 11, 12] and speaker extraction (*a.k.a.* voice filtering) [13, 14, 15]. Speech separation models can directly separate two or more sources from the microphone mixture signal. However, these models require an output channel selection step after the sepa-

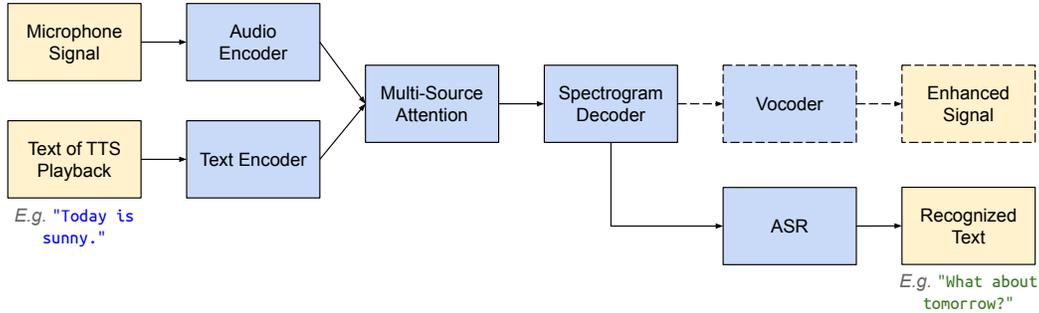


Fig. 2. Diagram of the textual echo cancellation framework.

ration to correctly keep the user’s query by employing an additional speaker verification system. On the other hand, for example, the VoiceFilter system proposed in [15] separates the voice of a target speaker from multi-speaker signals, by making use of a reference signal from the target speaker. VoiceFilter assumes we have speech samples from the target user, such that we can build an embedding vector of the user’s voice characteristics, and use this embedding vector as an auxiliary input to remove any signal that does not belong to the target user. However, if the user does not provide audio samples to enroll on the device, VoiceFilter is not feasible. Also, if the TTS playback voice is similar to the voice of the user, it is challenging for VoiceFilter to remove the TTS from the microphone mixture signal.

Previous studies also explored the use of text information in speech enhancement [16, 17, 18]. However, our approach differs from these works in several aspects. [16] is the most relevant work to ours since they consider the use of text information in a speech enhancement system based on deep neural network. However, their model requires an extra algorithms [16] to align speech signal and text, and these alignment algorithms are usually not very robust to background noise, as illustrated in [19, 20]. By contrast, our method achieves the alignment by the attention mechanism in the seq2seq model, which avoids the need of alignment algorithms. The other two studies share the idea of using text information in speech enhancement, but our approach is significantly different in either model ([17] uses non-negative matrix factorization) and how to incorporate text information ([18] uses a auxiliary speech recognizer).

To address these limitations, we propose a novel framework named as **Textual Echo Cancellation (TEC)**<sup>2</sup>. Instead of using the TTS playback or the user’s speaker embedding as side input, we use the source text of the TTS playback. Comparing to the TTS playback, the source text is much smaller in size. Consequently, it can be efficiently transmitted between servers, or from server to device. This will avoid extra Internet communications and latency. The proposed approach is implemented using seq2seq modeling with multi-source attention [21, 22, 23], as shown in Fig. 2. Our model takes two sequences as input: The Mel-spectrogram of the mixture audio recorded from the microphone, and the source text corresponding to the TTS playback. An audio encoder and a text encoder are used to extract representations for the two input sequences, respectively. A multi-source attention attends to both the encoded speech and the encoded text, and outputs a fixed-dimensional context vector at each decoding step. Finally, the decoder consumes the context vectors and autoregressively predicts a Mel-spectrogram corresponding to the enhanced signal. The main contributions of our work are outlined as below:

- We propose a novel seq2seq model for echo cancellation, which takes microphone mixture signals along with the source text of the TTS playback as the inputs, and produces enhanced audio signal. Compared to conventional echo cancellation approaches, our model avoids the need of TTS playback signal, reducing Internet communication and latency for smart speaker devices.
- We utilize a multi-source attention mechanism in our proposed seq2seq model. It generates a fix-length context vector during each decoding step based on the hidden representations extracted from microphone mixture signals and source text, thus incorporating information from multiple sources.
- We conduct extensive experiments on LibriTTS and VCTK datasets to evaluate the proposed model. We implement a conventional signal-processing AEC baseline and two seq2seq AEC baselines, and we compare the proposed model against them under two different conditions that mimic practical use cases of smart speaker devices.

The rest of the paper is organized as follows. In Section 2, we will give detailed description of different components of the proposed framework. In Section 3, we describe our experimental setup, including the data, model parameters, metrics, and results. Finally, we draw the conclusions and point out potential future work in Section 4.

## 2. TEXTUAL ECHO CANCELLATION

### 2.1. Overview

A diagram illustrating the textual echo cancellation framework is shown in Fig 2. Suppose we have a feature sequence  $\mathbf{x} \in \mathbb{R}^{T_x \times D_m}$  for the microphone mixture signal, an embedding sequence  $\mathbf{y} \in \mathbb{R}^{T_y \times D_e}$  representing the source text of the TTS playback, and a feature sequence  $\mathbf{z} \in \mathbb{R}^{T_z \times D_m}$  for the user’s actual clean speech. Here  $T_x$ ,  $T_y$ , and  $T_z$  are the lengths of the sequences  $\mathbf{x}$ ,  $\mathbf{y}$ ,  $\mathbf{z}$ , respectively,  $D_m$  is the number of Mel-filterbanks (e.g. 128 Mel-filterbanks in this work), and  $D_e$  is the dimension of the text embedding (e.g. 512-dimensional pre-trained phoneme embedding in this work). Our model consists of three modules: (1) an audio encoder, (2) a text encoder, and (3) an autoregressive decoder with a multi-source attention mechanism.

First, the audio encoder takes the microphone feature sequence as the input and produces a hidden representation:

$$\mathbf{h}_x = \text{Encoder}_{\text{audio}}(\mathbf{x}) \quad (1)$$

Similarly, the text encoder takes the the source text sequence as the input and produces another hidden representation:

<sup>2</sup>Audio samples are available at <https://google.github.io/speaker-id/publications/TEC>

$$\mathbf{h}_y = \text{Encoder}_{\text{text}}(\mathbf{y}) \quad (2)$$

Finally, the decoder autoregressively predicts the Mel-spectrogram of the enhanced signal using the attention context computed based on the two encoder outputs:

$$\hat{\mathbf{z}}^t = \text{Decoder}(\hat{\mathbf{z}}^{t-1}, \mathbf{h}_x, \mathbf{h}_y) \quad (3)$$

The predicted Mel-spectrogram can be directly consumed by an ASR model or other downstream components such as vocoders. We will describe each module with details in the following subsections. The hyperparameters of each module are shown in Table 1.

## 2.2. Audio encoder

The audio encoder converts a Mel-spectrogram sequence to a hidden representation sequence. Following [24], we use an encoder composed of two convolutional layers, one bi-directional convolutional LSTM layer (Bi-CLSTM) [25, 26], and three bi-directional LSTM (Bi-LSTM) layers. A convolutional layer has 32 kernels, each of which has a shape of  $3 \times 3$  in time  $\times$  frequency and a stride of  $2 \times 2$ , followed by ReLU activations and batch normalization [27], capturing the local temporal context information. Meanwhile, a convolutional layer reduces the time resolution by a factor of 2, which also reduces the computational cost in the following layers. The Bi-CLSTM and Bi-LSTM layers extract high-level frequency-wise features and capture long-term temporal context information. Each of these layers have 256 units in each direction, followed by ReLU activations and batch normalization. Additionally, the Bi-CLSTM layer has  $1 \times 3$  kernels with a stride of  $1 \times 1$ . As a result, the final output sequence of the audio encoder has a dimension of 512 and is four times shorter compared with the input sequence.

## 2.3. Text encoder

The text encoder converts text sequences (represented by phonemes) to hidden representation sequences. Each of the input phoneme is first represented by a pre-trained 512-dimensional embedding. Then the embedding sequence is passed through three convolutional layers and one Bi-LSTM layer, following [28]. Each convolutional layer has 512 kernels, and each kernel has a shape of  $5 \times 1$  and stride of  $1 \times 1$ , followed by ReLU activations and batch normalization. Each kernel in the convolutional layers spans 5 phonemes, modeling the local context information (*e.g.*, N-grams). The Bi-LSTM layer has 256 units in each direction, followed by ReLU activations and batch normalization, resulting in a 512-dimensional text encoder output sequence.

## 2.4. Decoder

The decoder is an autoregressive recurrent neural network coupled with a multi-source attention mechanism (see Section 2.5), as illustrated in Fig. 3. It takes the encoded sequences produced by the audio and text encoders as the inputs and generates the 128-dimensional enhanced Mel-spectrogram as a prediction of the user’s clean speech signal. We follow the same decoder architecture as in Tacotron 2 [28]. During each decoding step  $t$ , the prediction from the previous decoding step  $\hat{\mathbf{z}}^{t-1}$  is fed to a pre-net containing two fully-connected layers of 256 neurons along with ReLU activations:

$$\mathbf{q}^t = \text{PreNet}(\hat{\mathbf{z}}^{t-1}) \quad (4)$$

which is essential for learning attentions [28]. Then the multi-source attention mechanism computes a 128-dimensional attention context

**Table 1.** Hyperparameters of the TEC network.

Spectral analysis	frame length: 50 ms; frame shift: 12.5 ms; 128 Mel-filterbanks	
Audio encoder	Conv layers $\times$ 2	32 $3 \times 3$ kernel with $2 \times 2$ stride; ReLU; batch norm
	Bi-CLSTM $\times$ 1	256 units per direction; $1 \times 3$ kernel with $1 \times 1$ stride
	Bi-LSTM $\times$ 3	256 units per direction
Text encoder	Text embedding	Pre-trained model; 512-dim
	Conv layers $\times$ 3	512 $5 \times 1$ kernel with $1 \times 1$ stride; ReLU; batch norm
	Bi-LSTM $\times$ 1	256 units per direction
Attention	Multi-source attention	GMM attention for each source; 128-dim attention context
Decoder	PreNet	fully-connected layer $\times$ 2 256 neurons; ReLU
	LSTM $\times$ 2	256 units
	Linear (Mel)	fully-connect layer $\times$ 1 128 neurons; no activation
	Linear (stop token)	fully-connect layer $\times$ 1 2 neurons; no activation
	PostNet	Conv layers $\times$ 5 512 $5 \times 1$ kernel with $1 \times 1$ stride; TanH; batch norm

vector  $\mathbf{c}^t$  using the pre-net output, the attention context during the previous step, and the two encoded sequences from the two encoders:

$$\mathbf{c}^t = \text{MultiSourceAttention}(\mathbf{q}^t, \mathbf{c}^{t-1}, \mathbf{h}_x, \mathbf{h}_y) \quad (5)$$

Next, the pre-net output and the attention context vector are concatenated and passed through two uni-directional LSTM layers with 256 units. The LSTM outputs are then concatenated again with the attention context vector and fed to a linear transformation of 128 units, resulting in a predicted Mel-spectrogram frame for the user’s clean speech:

$$\hat{\mathbf{z}}_{\text{pre}}^t = \text{Linear}\left(\text{LSTM}(\mathbf{q}^t, \mathbf{c}^t), \mathbf{c}^t\right) \quad (6)$$

As the generated Mel-spectrogram from the seq2seq model is not a frame-synchronous estimate of  $\mathbf{z}^t$ , we also need the network to predict if the autoregressive generating process should stop at each decoding step, *i.e.*, a 0/1 stop token  $s^t$ . Finally, to incorporate residuals in predicted Mel-spectrogram, these predictions are passed through 5-layer convolutional post-net, each layer having 512 kernels of shape  $5 \times 1$  followed by batch normalization and tanh activation. The post-net predicts the residual that is added to the prediction, which has been shown to improve the Mel-spectrogram reconstruction [28]:

$$\hat{\mathbf{z}}^t = \text{PostNet}(\hat{\mathbf{z}}_{\text{pre}}^t) + \hat{\mathbf{z}}_{\text{pre}}^t \quad (7)$$

## 2.5. Multi-source attention mechanism

We use a multi-source attention mechanism to summarize the encoded sequences from both audio and text encoders. The multi-source attention consists of two individual attentions for the two encoders, respectively, without sharing the weights between the two encoders. During each decoding step  $t$ , the two attentions first produce two fixed-length attention contexts:

$$\mathbf{c}_x^t = \text{Attention}_{\text{audio}}(\mathbf{q}^t, \mathbf{c}_x^{t-1}, \mathbf{h}_x) \quad (8)$$

$$\mathbf{c}_y^t = \text{Attention}_{\text{text}}(\mathbf{q}^t, \mathbf{c}_y^{t-1}, \mathbf{h}_y) \quad (9)$$

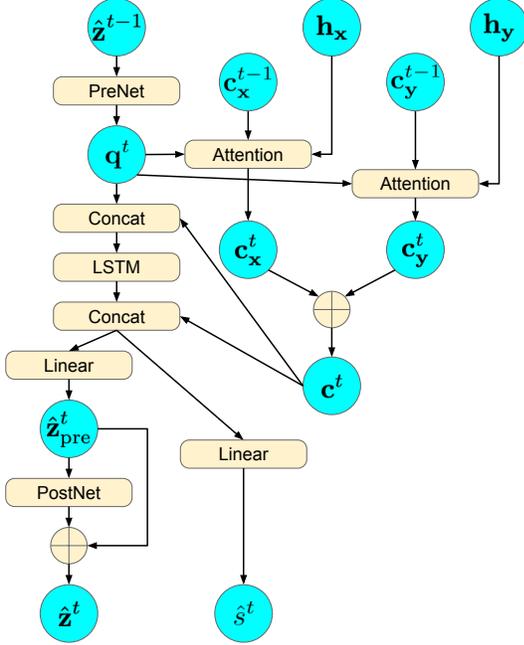


Fig. 3. Diagram of the decoder with multi-source attention.

Then we obtain the final context vector by summing up the two context vectors:

$$\mathbf{c}^t = \mathbf{c}_x^t + \mathbf{c}_y^t \quad (10)$$

There are other strategies to combine the two context vectors, such as averaging, concatenation, and hierarchical attention combination [29]. However, our preliminary results show that the difference between different combination strategies are minimal, so we use the simplest summation operation here. In addition, we use Gaussian mixture attention mechanism [30] for both audio and text, which has been shown to achieve superior performance than conventional additive attention mechanism [31] on speech synthesis [32, 33, 34].

## 2.6. Model training and inference

Following [35], during training, the model is optimized by minimizing the sum of the L1 and L2 distances computed from the output before and after the post-net. We apply the teacher-forcing training procedure (feeding in the correct output instead of the predicted output on the decoder side). As a result, we need to jointly minimize an extra cross-entropy loss to learn the stop token for model inference. The overall loss function of the proposed model is:

$$L = \|\hat{\mathbf{z}}_{\text{pre}} - \mathbf{z}\|_2^2 + \|\hat{\mathbf{z}} - \mathbf{z}\|_2^2 + \|\hat{\mathbf{z}}_{\text{pre}} - \mathbf{z}\|_1 + \|\hat{\mathbf{z}} - \mathbf{z}\|_1 + \text{CrossEntropy}(\hat{\mathbf{s}}, \mathbf{s}) \quad (11)$$

where  $\hat{\mathbf{s}}$  is the sequence of the predicted stop token and  $\mathbf{s}$  is the sequence of the target stop token.

Once we have a trained model, we can pass the Mel-spectrogram of the microphone mixture signal along with the source text of the TTS playback to the model to acquire the Mel-spectrogram of the enhanced signal. The Mel-spectrogram can be directly consumed by downstream ASR. Additionally, we can also use a vocoder (e.g.

WaveNet [36] or WaveRNN [37]) to synthesize the waveform of the enhanced audio if it is needed for other downstream modules.

## 3. EXPERIMENTS

We conduct experiments under two different conditions to evaluate the proposed approach. In the first experiment, we consider the TTS voice being generated from a canonical speaker (denoted as *single interfering voice condition*). Following this, we extend the TTS voice to have multiple different speakers' identities in the second experiment (denoted as *multiple interfering voices condition*), which is closer to real-world scenarios (e.g. personalized playback voice in smart speaker devices) but more challenging.

### 3.1. Datasets

Following previous echo cancellation studies [3, 5, 4, 38], we use synthetic data for the evaluations. In both conditions, we use the LibriTTS dataset [39] for the user query. To produce microphone mixture signal, we mix utterances from LibriTTS with the utterances from the LJ Speech dataset [40] and the CSTR VCTK dataset [41] in the single/multiple interfering voice(s) conditions, respectively (see Section 3.2). Comparing against TIMIT dataset that is commonly used in previous studies [3, 5, 4, 38], these datasets contain continuous sentences instead of just the recording of ten digits, which is more appropriate in simulating the practical use cases.

The LibriTTS dataset consists of 585 hours of audio book speech data from 2,456 speakers. The dataset is divided into three parts: 555 hours of training sets, 15 hours of development sets, and 15 hours of testing set. Each of them contains both clean and noisy speech.

The LJ Speech dataset has 24 hours clean audio book speech data from a single speaker. The original LJ Speech dataset does not have training and testing subsets. For evaluation purpose, we randomly selected 90% of the utterances as the training set and the remaining 10% as the testing set.

The CSTR VCTK dataset contains 44 hours of clean speech from 109 speakers. Similarly, we randomly selected 90% of the utterances from each speaker as the training set and the remaining 10% as the testing set, since there are no official training and testing subsets.

### 3.2. Generating synthetic microphone mixture signal

We use the acoustic signal model described in [38, 4] to generate synthetic microphone mixture signals. In their model, the microphone mixture signal  $x(n)$  is generated as:

$$x(n) = z(n) + y(n) * h(n) \quad (12)$$

where  $z(n)$  is user's speech signal,  $y(n)$  is TTS playback,  $h(n)$  is the room impulse response (RIR), and  $*$  is the convolutional operation. A total number of 3 million RIR were generated using a room simulator [42, 43, 44] to cover different reverberation conditions. With Eq. 12, we generated three subsets of synthetic data: (1) training, (2) test-clean, and (3) test-other, as shown in Table 2.

For each user's speech utterance, we randomly chose an interfering utterance and followed the above model to generate a microphone mixture signal with a 0-dB Signal-Noise-Ratio (SNR). As a result, the number of mixtures in the synthetic datasets is the same as the number of utterances in the LibriTTS dataset. Additionally, we padded the two utterances to have the same length to handle the duration difference between two utterances.

**Table 2.** Data configuration for our experiments. This table shows how the microphone signal mixtures were generated. For example, in single interfering voice condition, the synthetic training set was mixed using LibriTTS training set and LJ Speech training set.

Condition	Synthetic subset	User’s speech	Interfering speech
Single interfering voice	training	LibriTTS training	LJ Speech training
	test-clean	LibriTTS test-clean	LJ Speech test
	test-other	LibriTTS test-other	LJ Speech test
Multiple interfering voices	training	LibriTTS training	VCTK training
	test-clean	LibriTTS test-clean	VCTK test
	test-other	LibriTTS test-other	VCTK test

For interfering speech, we use the same TTS speakers during model training and evaluation. This is based on the fact that for smart home speaker devices, there is typically a fixed set of TTS voice options. However, for user’s speech, we use different sets of speakers during training and evaluation (*e.g.*, LibriTTS train vs. dev/test sets), since the user of the device is unknown at runtime.

### 3.3. Metrics

To evaluate the proposed approach, we consider three metrics: (1) Word Error Rate (WER), (2) Mel-Cepstral Distortion (MCD) [45], and (3) Mean Opinion Score (MOS) of the naturalness of the enhanced speech audio. Additionally, we also estimate the side input size and computational complexity for different approaches.

#### 3.3.1. Word Error Rate

As described in Section 1, the main purpose of our proposed approach is to improve the speech recognition performance of smart speaker devices when the user’s query and the TTS playback echo have overlaps. As a result, we use WER as the major evaluation metric for our experiments. The speech recognizer we used for WER evaluation is a state-of-the-art model proposed in [46], which is trained on the LibriSpeech [47] training set. We did not re-train the recognizer on the generated features.

#### 3.3.2. Mel-Cepstral Distortion

MCD (dB) is a commonly used objective metric to evaluate the quality of the synthesized speech, which is defined as

$$\text{MCD} = \frac{10}{\ln 10} \sum_{t=1}^{T_z} \sqrt{2 \sum_{d=1}^{13} (\hat{z}_{t,d} - z_{t,d})^2} \quad (13)$$

where  $\hat{z}_{t,d}$  and  $z_{t,d}$  are the  $d$ -th Mel-Frequency Cepstral Coefficient (MFCC) of the enhanced speech and the time-aligned<sup>3</sup> target speech at the  $t$ -th time step, respectively. In this paper, we used 13 MFCCs (skipping MFCC<sub>0</sub>, which is energy) to compute MCD. Lower MCD indicates that the enhanced speech is more similar to the target clean speech.

We did not include the metrics that are used in conventional echo cancellation approaches (*e.g.* echo return loss enhancement metric [7], perceptual evaluation of speech quality [49]) for two reasons. First, the downstream ASR model can directly take the Mel-spectrogram as the input, and therefore, generating waveform becomes redundant and may cause extra distortions. Second, the waveform of the enhanced signal is generated using generative neural vocoder models, and the generated waveform can be very different from the target waveform, even if the linguistic content of the waveforms are exactly the same. As a result, these waveform based metrics become ill-defined in our case.

<sup>3</sup>We use dynamic time warping [48] for time alignment.

#### 3.3.3. Speech naturalness Mean Opinion Score

We measured the speech naturalness of the enhanced signal with a 5-point Mean Opinion Score (1-bad; 5-excellent). For each system, we randomly chose 1,000 utterances from test-clean/test-other subsets for MOS evaluation. All the utterances were synthesized using a separately trained WaveRNN model [37]. Each sample was rated by six raters, and each evaluation was conducted independently: the outputs of different models were not compared directly.

#### 3.3.4. Side input size and computational complexity

We also computed the side input size and the computational complexity to evaluate the resources that are required for practical applications. To measure the computational complexity, we estimate the number of floating-point operations (FLOPS) required following [21]:

$$\text{FLOPS} = M_{\text{audio}} \cdot T_x + M_{\text{text}} \cdot T_y + M_{\text{dec}} \cdot T_z + \text{FLOPS}_{\text{atten}} \quad (14)$$

where  $M_{\text{audio}}$ ,  $M_{\text{text}}$ , and  $M_{\text{dec}}$  are number of model parameters of the audio encoder, text encoder, and decoder, respectively.  $T_x$ ,  $T_y$ , and  $T_z$  are the sequence lengths of the Mel-spectrogram of the microphone mixture signal, text embedding, and user’s actual clean speech signal, respectively.  $\text{FLOPS}_{\text{atten}}$  is the FLOPS required for the multi-source attention layer. For each attention source, we first multiply the size of the attention source matrix with  $T_x/T_y$ , and multiply the size of query matrix with  $T_z$ , and then the FLOPS for multi-source attention is computed as the sum of the two.

### 3.4. Implementation details

We implemented the model using the Lingvo [50] framework in TensorFlow [51]. Our model was trained on  $2 \times 2$  Tensor Processing Units (TPU) slices with a global batch size of 32. During training, we use Adam optimizer [52] with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , and  $\epsilon = 10^{-6}$ . We set the initial learning rate to  $10^{-4}$  and exponentially decays to  $10^{-5}$  after 50,000 iterations.

### 3.5. Results

In each condition, we compared the proposed approach against three baselines that we implemented: (1) AEC-NLMS: A conventional signal-processing AEC algorithm based on normalized least mean square (NLMS) algorithm [7] that is widely used in prior studies [3, 38, 4]. (2) Vanilla-Seq2seq: A sequence-to-sequence network that directly transforms the microphone mixture signal to enhanced signal without any side input using single attention. This network is similar to [24, 53] except for not using auxiliary decoders, and we use a Gaussian mixture attention for it. (3) AEC-Seq2seq: An end-to-end model with similar architecture as the proposed TEC model. We replaced the text encoder in the proposed approach with an audio encoder that takes the TTS playback as the input, which operates similarly to other model-based AEC models.

**Table 3.** Word Error Rate (WER), Mel-Cepstral Distortion (MCD), and speech naturalness Mean Opinion Score (MOS) evaluation results of the single speaker interfering voice and multiple interfering voices conditions. The MOS is presented with 95% confidence intervals. We also include the size of the side input in kilobytes (the TTS playback echo or the TTS source text) and the floating point operations per second in Giga (GFLOPS).

Condition	Method	WER (%)		MCD		MOS		Side input (KB)	GFLOPS
		test-clean	test-other	test-clean	test-other	test-clean	test-other		
Ground-truth LibriTTS	-	2.30	4.50	0.00	0.00	4.43 ± 0.04	3.82 ± 0.06	-	-
Single interfering voice	Microphone signal	89.9	120.5	18.83	21.44	-	-	-	-
	AEC-NLMS	48.6	60.1	12.26	12.57	1.95 ± 0.10	1.28 ± 0.09	310	0
	Vanilla-Seq2seq	25.4	54.0	7.85	8.84	1.99 ± 0.06	1.47 ± 0.05	0	6.32
	AEC-Seq2seq	8.30	24.3	6.38	7.07	2.77 ± 0.07	1.90 ± 0.06	310	9.51
	TEC (proposed)	15.5	39.8	7.51	8.54	2.20 ± 0.07	1.65 ± 0.06	0.10	7.27
Multiple interfering voices	Microphone signal	29.7	44.6	10.75	12.88	-	-	-	-
	AEC-NLMS	15.5	35.5	6.57	8.13	2.06 ± 0.11	1.60 ± 0.08	230	0
	Vanilla-Seq2seq	19.7	38.7	7.53	8.87	2.16 ± 0.07	1.50 ± 0.05	0	6.32
	AEC-Seq2seq	6.90	19.8	5.04	5.72	2.90 ± 0.07	2.03 ± 0.07	230	8.62
	TEC (proposed)	14.8	32.5	6.46	7.71	2.39 ± 0.07	1.70 ± 0.06	0.06	6.90

★ The side input size and GFLOPS in the two conditions are different since the average lengths of the echo signal are different in the two conditions.

The WER, MCD, and MOS evaluation results of the two conditions are shown in Table 3. We include the three measurements of the ground-truth LibriTTS test-set, acting as a performance upper bound. Under single interfering voice condition, we observed that TEC achieves 15.5% WER, 7.51 MCD, and 2.20 MOS on test-clean subset as well as 39.8% WER, 8.54 MCD, and 1.65 MOS on test-other subset. Under multiple interfering voices condition, TEC achieves 14.8% WER, 6.46 MCD, and 2.39 MOS on test-clean subset as well as 32.5% WER, 7.71 MCD, and 1.70 MOS on test-other subset. Comparing against the baseline systems, the WER, MCD, and MOS results in both conditions consistently suggest that our proposed approach significantly outperforms AEC-NLMS, corresponding to the observations obtained from prior studies that deep learning models usually achieve superior performance than signal-processing based AEC algorithms. In addition, TEC also achieves essential improvement than Vanilla-Seq2seq, indicating that the information of TTS playback is key in achieving reasonable echo cancellation performance. However, we found that TEC is not as good as AEC-Seq2seq. This is expected since TEC only uses the source text of the TTS playback instead of the TTS playback echo, and the source text contains less information about how the TTS playback actually sounds like. Although there is a performance gap between TEC and AEC-Seq2seq, the size of the side inputs and the GFLOPS of TEC are much lower than AEC-Seq2seq, which supports our argument that TEC is more efficient in terms of Internet communications and latency.

The performances of all the systems on test-clean subset are better than those on test-other subset, since the utterances in test-other subset have considerable background noises, which degrades the quality of the output enhanced audio. Additionally, it is interesting to observe that the WER, MCD and MOS of all the systems under multiple interfering voices condition are better than those under single interfering voice condition. A possible explanation of this observation is that the utterances in VCTK are much shorter than those in LJ Speech (~2 seconds vs. ~7 seconds, in terms of average duration per utterance), and therefore, the interfered intervals under multiple interfering voices condition are much shorter than those under single interfering voice condition. It makes the test set under multiple interfering voices condition an easier case, which is also supported by our results that the WER and MCD of the microphone signal under this condition are lower than those under single interfering voice condition. Besides, most of the utterances in VCTK

have a British English accent, while LibriTTS and LJ Speech are dominated by American English accent. Both factors make it easier to separate the user’s speech from the interfering speech under multiple interfering voices condition than that under single interfering voice condition.

#### 4. CONCLUSIONS AND FUTURE WORK

In this paper, we proposed textual echo cancellation, a framework to cancel the TTS playback echo from overlapped speech, which is useful when a user talks to an intelligent device while the device is still playing synthesized response to a previous query. Our proposed approach uses the source text as the side input instead of the TTS playback, which can be efficiently transmitted between servers and from server to device, thus largely reducing Internet communications and latency compared with conventional AEC-Seq2seq approaches. We conducted experiments under a single interfering voice condition and a multiple interfering voices condition. Our experimental results show that TEC significantly outperforms the baseline of not using any side input, indicating that the textual information of the TTS playback is critical to the enhancement performance. In addition, the side input size and GFLOPS of TEC are much lower than model based AEC-Seq2seq methods.

In our experiments, the performance of TEC is still not as good as that of AEC-Seq2seq. Several directions can be explored in the future. First, a second decoder for phonetic recognition can be added during training, which has shown to be effective in [24, 53] for speech-to-speech conversion models. Additionally, alternative architectures such as frame-to-frame TEC models can be implemented and compared with our current multi-source attention sequence-to-sequence TEC model. Furthermore, in applications where there is no restrictions on computations and data transmissions (e.g., offline echo cancellation), we can consider both TTS playback and source text as the side inputs to the model, which may provide extra performance gains. Last but not least, we can train and evaluate the proposed approach on data in the wild instead of synthetic data. As mentioned in Section 1, AEC-Seq2seq model is subjected to the mismatch problem between the TTS playback echo and the TTS playback. By contrast, TEC does not have such a problem, and therefore, we believe the gap between AEC-Seq2seq and TEC will become smaller on wild data.

## 5. REFERENCES

- [1] Eberhard Hänsler and Gerhard Schmidt, *Acoustic echo and noise control: A practical approach*, vol. 40, John Wiley & Sons, 2005.
- [2] Jacob Benesty, Tomas Gänsler, Dennis R Morgan, M Mohan Sondhi, and Steven L Gay, *Advances in network and acoustic echo cancellation*, Springer, 2001.
- [3] Hao Zhang and DeLiang Wang, “Deep learning for acoustic echo cancellation in noisy and double-talk scenarios,” in *Proceedings of Interspeech*, 2018, pp. 3239–3243.
- [4] Amin Fazel, Mostafa El-Khamy, and Jungwon Lee, “Deep multitask acoustic echo cancellation,” in *Proceedings of Interspeech*, 2019, pp. 4250–4254.
- [5] Qinhui Lei, Hang Chen, Junfeng Hou, Liang Chen, and Lirong Dai, “Deep neural network based regression approach for acoustic echo cancellation,” in *Proceedings of the 2019 4th International Conference on Multimedia Systems and Signal Processing*, 2019, pp. 94–98.
- [6] Jacob Benesty, Constantin Paleologu, Tomas Gänsler, and Silviu Ciochină, *A perspective on stereophonic acoustic echo cancellation*, vol. 4, Springer Science & Business Media, 2011.
- [7] Gerald Enzner, Herbert Buchner, Alexis Favrot, and Fabian Kuech, “Acoustic echo control,” in *Academic Press Library in Signal Processing*, vol. 4, pp. 807–877. Elsevier, 2014.
- [8] Ke Hu and DeLiang Wang, “An unsupervised approach to cochannel speech separation,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 21, no. 1, pp. 122–131, 2012.
- [9] Po-Sen Huang, Minje Kim, Mark Hasegawa-Johnson, and Paris Smaragdis, “Deep learning for monaural speech separation,” in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 1562–1566.
- [10] Jun Du, Yanhui Tu, Yong Xu, Lirong Dai, and Chin-Hui Lee, “Speech separation of a target speaker based on deep neural networks,” in *Proceedings of International Conference on Signal Processing (ICSP)*. IEEE, 2014, pp. 473–477.
- [11] John R Hershey, Zhuo Chen, Jonathan Le Roux, and Shinji Watanabe, “Deep clustering: Discriminative embeddings for segmentation and separation,” in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 31–35.
- [12] Yi Luo and Nima Mesgarani, “TasNet: Time-domain audio separation network for real-time, single-channel speech separation,” in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 696–700.
- [13] Jun Wang, Jie Chen, Dan Su, Lianwu Chen, Meng Yu, Yanmin Qian, and Dong Yu, “Deep extractor network for target speaker recovery from single channel speech mixtures,” *arXiv preprint arXiv:1807.08974*, 2018.
- [14] Marc Delcroix, Katerina Zmolikova, Keisuke Kinoshita, Atsunori Ogawa, and Tomohiro Nakatani, “Single channel target speaker extraction and recognition with speaker beam,” in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5554–5558.
- [15] Quan Wang, Hannah Muckenhirn, Kevin Wilson, Prashant Sridhar, Zelin Wu, John R Hershey, Rif A Saurous, Ron J Weiss, Ye Jia, and Ignacio Lopez Moreno, “VoiceFilter: Targeted voice separation by speaker-conditioned spectrogram masking,” in *Proceedings of Interspeech*, 2019, pp. 2728–2732.
- [16] Keisuke Kinoshita, Marc Delcroix, Atsunori Ogawa, and Tomohiro Nakatani, “Text-informed speech enhancement with deep neural networks,” in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [17] Luc Le Magoarou, Alexey Ozerov, and Ngoc QK Duong, “Text-informed audio source separation using nonnegative matrix partial co-factorization,” in *2013 IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, 2013, pp. 1–6.
- [18] Hakan Erdogan, John R Hershey, Shinji Watanabe, and Jonathan Le Roux, “Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 708–712.
- [19] Brian King, Paris Smaragdis, and Gautham J Mysore, “Noise-robust dynamic time warping using plca features,” in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2012, pp. 1973–1976.
- [20] Xianglilan Zhang, Jiping Sun, and Zhigang Luo, “One-against-all weighted dynamic time warping for language-independent and speaker-dependent speech recognition in adverse conditions,” *PloS one*, vol. 9, no. 2, pp. e85458, 2014.
- [21] Ke Hu, Tara N Sainath, Ruoming Pang, and Rohit Prabhavalkar, “Deliberation model based two-pass end-to-end speech recognition,” in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7799–7803.
- [22] Yingce Xia, Fei Tian, Lijun Wu, Jianxin Lin, Tao Qin, Nenghai Yu, and Tie-Yan Liu, “Deliberation networks: Sequence generation beyond one-pass decoding,” in *Proceedings of Advances in Neural Information Processing Systems*, 2017, pp. 1784–1794.
- [23] Anna Currey and Kenneth Heafield, “Multi-source syntactic neural machine translation,” in *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2018, pp. 2961–2966.
- [24] Fadi Biadsy, Ron J. Weiss, Pedro J. Moreno, Dimitri Kanvesky, and Ye Jia, “Parrottron: An end-to-end speech-to-speech conversion model and its applications to hearing-impaired speech and speech separation,” in *Proceedings of Interspeech*, 2019, pp. 4115–4119.
- [25] SHI Xingjian, Zhouong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo, “Convolutional LSTM network: A machine learning approach for precipitation nowcasting,” in *Proceedings of Advances in neural information processing systems*, 2015, pp. 802–810.
- [26] Mike Schuster and Kuldip K Paliwal, “Bidirectional recurrent neural networks,” *IEEE transactions on Signal Processing*, vol. 45, no. 11, pp. 2673–2681, 1997.
- [27] Sergey Ioffe and Christian Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” *arXiv preprint arXiv:1502.03167*, 2015.

- [28] Jonathan Shen, Ruoming Pang, Ron J Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, Rj Skerrv-Ryan, et al., “Natural TTS synthesis by conditioning WaveNet on Mel spectrogram predictions,” in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4779–4783.
- [29] Jindřich Libovický and Jindřich Helcl, “Attention strategies for multi-source sequence-to-sequence learning,” in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 2017, pp. 196–202.
- [30] Alex Graves, “Generating sequences with recurrent neural networks,” *arXiv preprint arXiv:1308.0850*, 2013.
- [31] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio, “Neural machine translation by jointly learning to align and translate,” *arXiv preprint arXiv:1409.0473*, 2014.
- [32] Mutian He, Yan Deng, and Lei He, “Robust sequence-to-sequence acoustic modeling with stepwise monotonic attention for neural TTS,” in *Proceedings of Interspeech*, 2019, pp. 1293–1297.
- [33] RJ Skerry-Ryan, Eric Battenberg, Ying Xiao, Yuxuan Wang, Daisy Stanton, Joel Shor, Ron Weiss, Rob Clark, and Rif A Saurous, “Towards end-to-end prosody transfer for expressive speech synthesis with Tacotron,” in *Proceedings of International Conference on Machine Learning*, 2018, pp. 4693–4702.
- [34] Adam Polyak and Lior Wolf, “Attention-based WaveNet autoencoder for universal voice conversion,” in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6800–6804.
- [35] Ye Jia, Yu Zhang, Ron Weiss, Quan Wang, Jonathan Shen, Fei Ren, Patrick Nguyen, Ruoming Pang, Ignacio Lopez Moreno, Yonghui Wu, et al., “Transfer learning from speaker verification to multispeaker text-to-speech synthesis,” in *Advances in neural information processing systems*, 2018, pp. 4480–4490.
- [36] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu, “WaveNet: A generative model for raw audio,” in *Proceedings of ISCA Speech Synthesis Workshop*, 2016, pp. 125–125.
- [37] Nal Kalchbrenner, Erich Elsen, Karen Simonyan, Seb Noury, Norman Casagrande, Edward Lockhart, Florian Stimberg, Aäron van den Oord, Sander Dieleman, and Koray Kavukcuoglu, “Efficient neural audio synthesis,” in *Proceedings of International Conference on Machine Learning*, 2018.
- [38] Amin Fazel, Mostafa El-Khamy, and Jungwon Lee, “CAD-AEC: Context-aware deep acoustic echo cancellation,” in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6919–6923.
- [39] Heiga Zen, Viet Dang, Rob Clark, Yu Zhang, Ron J Weiss, Ye Jia, Zhifeng Chen, and Yonghui Wu, “LibriTTS: A corpus derived from LibriSpeech for text-to-speech,” in *Proceedings of Interspeech*, 2019, pp. 1526–1530.
- [40] Keith Ito and Linda Johnson, “The LJ speech dataset,” <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [41] Christophe Veaux, Junichi Yamagishi, Kirsten MacDonald, et al., “Superseded-CSTR VCTK corpus: English multi-speaker corpus for CSTR voice cloning toolkit,” 2016.
- [42] Richard Lippmann, Edward Martin, and D Paul, “Multi-style training for robust isolated-word speech recognition,” in *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. IEEE, 1987, vol. 12, pp. 705–708.
- [43] Tom Ko, Vijayaditya Peddinti, Daniel Povey, Michael L Seltzer, and Sanjeev Khudanpur, “A study on data augmentation of reverberant speech for robust speech recognition,” in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 5220–5224.
- [44] Chanwoo Kim, Ananya Misra, Kean Chin, Thad Hughes, Arun Narayanan, Tara N Sainath, and Michiel Bacchiani, “Generation of large-scale simulated utterances in virtual rooms to train deep-neural networks for far-field speech recognition in Google Home,” in *Proceedings of Interspeech*, 2017, pp. 379–383.
- [45] R Kubichek, “Mel-cepstral distance measure for objective speech quality assessment,” in *Proceedings of IEEE Pacific Rim Conference on Communications Computers and Signal Processing*. IEEE, 1993, vol. 1, pp. 125–128.
- [46] Daniel S Park, Yu Zhang, Ye Jia, Wei Han, Chung-Cheng Chiu, Bo Li, Yonghui Wu, and Quoc V Le, “Improved noisy student training for automatic speech recognition,” *arXiv preprint arXiv:2005.09629*, 2020.
- [47] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “LibriSpeech: An ASR corpus based on public domain audio books,” in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [48] Stan Salvador and Philip Chan, “Toward accurate dynamic time warping in linear time and space,” *Intelligent Data Analysis*, vol. 11, no. 5, pp. 561–580, 2007.
- [49] Antony W Rix, John G Beerends, Michael P Hollier, and Andries P Hekstra, “Perceptual evaluation of speech quality (PESQ) - a new method for speech quality assessment of telephone networks and codecs,” in *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (ICASSP)*. IEEE, 2001, vol. 2, pp. 749–752.
- [50] Jonathan Shen, Patrick Nguyen, Yonghui Wu, Zhifeng Chen, Mia X Chen, Ye Jia, Anjali Kannan, Tara Sainath, Yuan Cao, Chung-Cheng Chiu, et al., “Lingvo: A modular and scalable framework for sequence-to-sequence modeling,” *arXiv preprint arXiv:1902.08295*, 2019.
- [51] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al., “Tensorflow: A system for large-scale machine learning,” in *Proceedings of USENIX symposium on operating systems design and implementation (OSDI 16)*, 2016, pp. 265–283.
- [52] Diederik P Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [53] Ye Jia, Ron J Weiss, Fadi Biadsy, Wolfgang Macherey, Melvin Johnson, Zhifeng Chen, and Yonghui Wu, “Direct speech-to-speech translation with a sequence-to-sequence model,” in *Proceedings of Interspeech*, 2019.