

SpeechX: Neural Codec Language Model as a Versatile Speech Transformer

Xiaofei Wang, Manthan Thakker, Zhuo Chen, Naoyuki Kanda, Sefik Emre Eskimez, Sanyuan Chen, Min Tang, Shujie Liu, Jinyu Li, Takuya Yoshioka

Abstract—Recent advancements in generative speech models based on audio-text prompts have enabled remarkable innovations like high-quality zero-shot text-to-speech. However, existing models still face limitations in handling diverse audio-text speech generation tasks involving transforming input speech and processing audio captured in adverse acoustic conditions. This paper introduces SpeechX, a versatile speech generation model capable of zero-shot TTS and various speech transformation tasks, dealing with both clean and noisy signals. SpeechX combines neural codec language modeling with multi-task learning using task-dependent prompting, enabling unified and extensible modeling and providing a consistent way for leveraging textual input in speech enhancement and transformation tasks. Experimental results show SpeechX’s efficacy in various tasks, including zero-shot TTS, noise suppression, target speaker extraction, speech removal, and speech editing with or without background noise, achieving comparable or superior performance to specialized models across tasks. See <https://aka.ms/speechx> for demo samples.

Index Terms—Speech generation, audio-text input, multi-task learning, zero-shot text-to-speech, noise suppression, target speaker extraction, speech editing, speech removal

I. INTRODUCTION

The technology of generative models has undergone rapid and transformative advancements in various machine learning applications, encompassing text [1], [2], vision [3], and audio [4]. These advancements have had significant implications for both the industry and society at large. Notably, generative models using multi-modal input have emerged as a remarkable innovation [5], [6], [7], [8], [9], [10].

In the speech domain, one prominent speech generation task that leverages audio-text input is zero-shot text-to-speech (TTS). Zero-shot TTS involves converting a given text into speech with the voice characteristics and speaking style of a desired talker by using only a brief audio sample of that person. Early studies in zero-shot TTS employed fixed-dimensional speaker embeddings [11], [12], [13], [14]. This approach limited their usage to TTS alone and did not adequately support speaker cloning capabilities.

In contrast, recent approaches have embraced more generic formulations, such as masked speech prediction [15] or neural codec language modeling [16], [17], [18], [19]. These novel approaches directly utilize the target speaker’s audio without compressing it into a fixed-dimensional representation. Consequently, these models have not only achieved remarkable zero-shot TTS performance but also demonstrated additional capabilities, including voice conversion [15], [18] and speech edit-

ing [15]. This enhanced flexibility holds tremendous promise for unlocking new possibilities in speech generation models.

However, despite their impressive achievements, these recent generative models still have certain limitations, particularly when it comes to addressing various audio-text-based speech generation tasks involving transforming input speech. For instance, existing speech editing models [20], [21] are restricted to handling clean signals only, lacking the ability to modify spoken content while preserving background sounds. Additionally, to perform denoising, the model discussed in [15] necessitates the noisy signal to be surrounded by clean speech segments, imposing significant constraints on its practical applications. In the context of transforming non-clean speech, another particularly useful task is target speaker extraction [22], [23], [24]. Target speaker extraction involves extracting the voice of a desired speaker from a speech mixture containing multiple talkers. The desired speaker can be specified using a short voice recording of that individual. Despite its potential significance as discussed in [25], this task remains unaddressed by existing generative speech models.

It is noteworthy that traditional approaches to speech enhancement tasks, such as denoising and target speaker extraction, have relied on regression models for faithful signal recovery. However, these prior methods typically required distinct expert models for each task, which is not ideal, given the potential diversity of acoustic disturbances [26]. Furthermore, there has been a lack of comprehensive audio-text-based speech enhancement models that leverage reference transcriptions to generate intelligible speech, except for limited studies focusing only on particular speech enhancement tasks [27], [28].

Given the aforementioned considerations and the successful precedents in other domains, the creation of audio-text-based generative speech models unifying generation and transformation capabilities assumes crucial research importance. These models should possess an overarching capability to tackle a diverse array of speech generation tasks. We propose that such models should be equipped with the following key properties:

- **Versatility:** Similar to unified or foundation models developed in other machine learning domains, the unified audio-text-based generative speech models must handle a wide range of tasks involving speech generation from audio and text inputs. These tasks should encompass not only zero-shot TTS but also various forms of speech transformation, including speech enhancement and speech editing, to name a few.

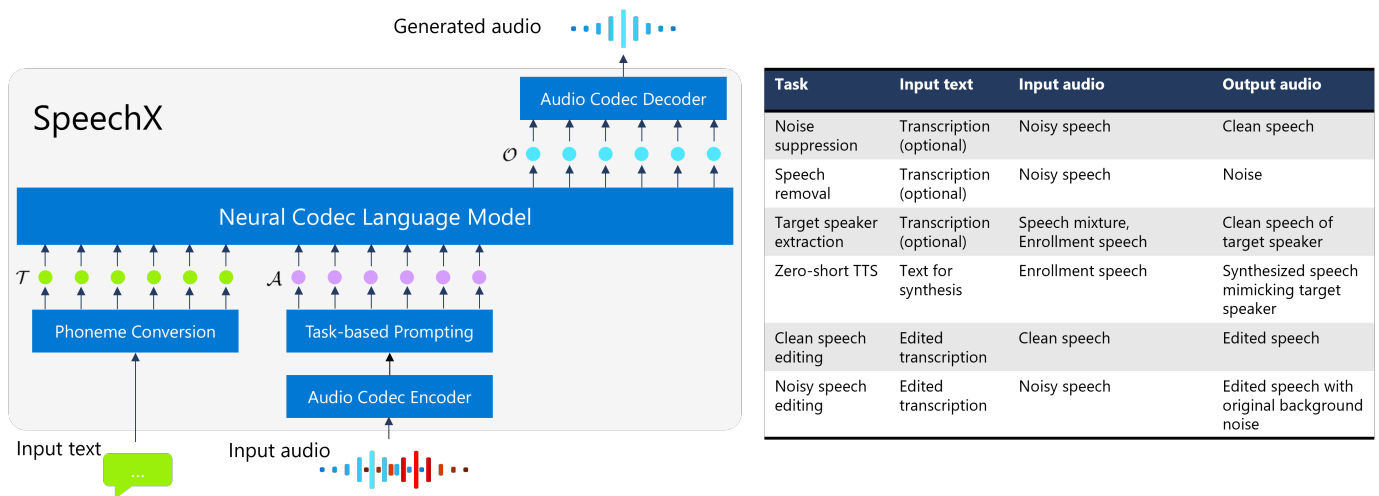


Fig. 1. Overview of SpeechX. SpeechX handles multiple audio-text-based speech generation tasks, including noise suppression, speech removal, target speaker extraction, zero-shot TTS, clean speech editing, and noisy speech editing, by using neural codec language model conditioned on the text and acoustic token stream. Text input is optional for some tasks.

- **Robustness:** It is essential for the unified models to exhibit robustness to various acoustic distortions since they are likely to be applied in acoustically challenging environments. By ensuring reliable performance, these models can be deemed highly usable in real-world scenarios where background sounds are prevalent.
- **Extensibility:** The unified models must employ flexible architectures, allowing for seamless extensions of task support. One approach to achieving this involves accommodating additional elements, such as input tokens or extra modules. Such flexibility will empower the models to adapt to future speech generation tasks efficiently.

In pursuit of this objective, this paper introduces a versatile speech generation model capable of performing multiple tasks, including zero-shot TTS, noise suppression using an optional transcript input, speech removal, target speaker extraction using an optional transcript input, and speech editing for both quiet and noisy acoustic environments (Fig. 1). We refer to our proposed model as **SpeechX**¹. As with VALL-E, SpeechX adopts a language modeling approach that generates codes of a neural codec model, or acoustic tokens, based on textual and acoustic inputs. To enable the handling of diverse tasks, we incorporate additional tokens in a multi-task learning setup, where the tokens collectively specify the task to be executed. Experimental results, using 60K hours of speech data from LibriLight [29] as a training set, demonstrate the efficacy of SpeechX, showcasing comparable or superior performance compared to expert models in all the aforementioned tasks. Notably, SpeechX also exhibits novel or expanded capabilities, such as preserving background sounds during speech editing and leveraging reference transcriptions for noise suppression and target speaker extraction. Audio samples showcasing the capabilities of our proposed SpeechX model are available at <https://aka.ms/speechx>.

¹X stands for transformation to highlight that our model performs various speech transformation tasks in addition to zero-shot TTS.

II. RELATED WORK

A. Autoregressive generative models

Generative models based on a language modeling approach using autoregressive Transformers, also known as decoder-only Transformers, have garnered significant success in various application domains. Notable examples of such models include the GPT series [1], [2] and DALL-E [30]. The autoregressive approach has also been extended to the audio and speech domains. AudioLM [4] and MusicLM [10] are pioneering efforts that exploit multiple types of tokens, each with a distinct time scale and degree of semantic granularity, allowing for hierarchical token generation. This hierarchical structure, comprising both coarse and fine-grained tokens, enables the synthesis of sounds with both nuanced details and long-term regularities.

For zero-shot TTS, VALL-E [16] and SPEAR-TTS [17] employ the autoregressive Transformers by representing textual (semantic) and acoustic tokens as a single data stream. This approach enables the models to perform zero-shot speaker adaptation, facilitating the generation of TTS voices that mimic a specific person’s voice. It was demonstrated that these models could perform zero-shot TTS from speech clips as short as three seconds. A notable advantage of these autoregressive speech generation models is their ability to perform TTS without requiring a separate duration model. This streamlined architecture simplifies the training process and potentially offers increased flexibility needed to subsume various speech generation tasks. For this reason, we opt to build our SpeechX models by using autoregressive Transformers.

B. Multi-task generative speech models

Several papers have recently reported efforts in developing audio-text-based speech generation models that support zero-shot TTS and several related tasks. These tasks include voice or style conversion (Make-A-Voice [18], NaturalSpeech2 [31]),

and Voicebox [15]), speech editing (Mega-TTS [21] and Voicebox), and denoising (NaturalSpeech2 and Voicebox). Voicebox has showcased noteworthy advancements by facilitating a multitude of tasks through its masked speech prediction principle. Nevertheless, its capabilities are still limited to clean speech generation alone, falling short of effectively dealing with noisy speech or encompassing conventional audio enhancement tasks such as noise suppression and target speaker extraction.

In this study, we deal with both clean and noisy speech and unify the generation and transformation tasks. To accomplish this, we extend VALL-E by performing multi-task learning with task-dependent prompts. The resulting model, which we call SpeechX, exhibits versatility in various speech processing tasks. The model excels not only in speech generation tasks like zero-shot TTS and speech editing but also performs effectively in enhancement tasks such as noise suppression and target speaker extraction. It also realizes novel capabilities, such as editing spoken content while retaining the background noise or effectively leveraging transcriptions for enhancement tasks.

III. METHOD

A. Overview

Fig. 1 illustrates an overview of the SpeechX architecture. Building upon the principles introduced in VALL-E, SpeechX employs a neural codec language model based on Transformers. The model learns to perform conditional generation of a neural code sequence, denoted as \mathcal{O} , based on two input prompts: textual prompt \mathcal{T} and acoustic prompt \mathcal{A} . The neural codes may also be referred to as acoustic tokens.

The textual prompt \mathcal{T} is a sequence of phonemes obtained by applying grapheme-to-phoneme conversion² to an input text. The textual prompt conveys the semantic information, and thus it is called semantic tokens. Conversely, the acoustic prompt \mathcal{A} encapsulates the acoustic information of an input speech signal. It is obtained by converting the input audio into a sequence of acoustic tokens with an encoder of the neural codec model. Furthermore, to specify the task to be executed, or equivalently the desired output, we incorporate additional tokens in the acoustic prompt. The details will be explained in Section III-C. The output \mathcal{O} is a sequence of neural codes of the desired signal, which is then translated into a waveform signal with the codec decoder.

We use EnCodec [32] as the neural codec model, following the prior work. Encodec is based on an encoder-decoder architecture with L quantization layers. In our experiments, we use $L = 8$ to be consistent with the configuration of [16]. Each layer of the EnCodec model produces discrete codes consisting of 1024 entries at a sampling rate of 75 Hz.

We emphasize that the proposed simple architecture capitalizes on the end-to-end modeling capability of the neural language modeling approach. In contrast to other zero-shot TTS or speech generation methods, this approach eliminates the need for a separate model, such as a speaker embedding

model or a duration model, apart from the neural codec model. This key property allows SpeechX to acquire knowledge of diverse tasks with varying requirements and input-output relationships, thereby facilitating a versatile and highly extensible speech generation process.

B. Neural codec language model

As with VALL-E [16], SpeechX makes use of auto-regressive (AR) and non-auto-regressive (NAR) Transformer models. Specifically, the AR model is used to output the neural codes corresponding to the first quantization layer of EnCodec. On the other hand, the NAR model generates the neural codes of all the layers above the first layer, namely the second through eighth layers. Combining the AR and NAR models provides a reasonable trade-off between generation flexibility and inference speed, as discussed in [16].

Let output \mathcal{O} be specifically represented as matrix $\mathbf{O} = [o_{t,l}] \in \mathbb{N}^{T \times L}$, where $o_{t,l}$ represents the code for the l -th codec layer at time frame t and it can take one of the 1024 values. The output sequence length is denoted by T . The AR model comprises a stack of Transformer decoder layers [33] and is optimized by minimizing the negative log-likelihood of the first layer code of the desired output, which is defined as follows:

$$\mathcal{L}_{AR} = - \sum_{t=1}^T \log P(o_{t,1} | \mathcal{T}, \mathcal{A}, \mathbf{o}_{<t,1}; \theta_{AR}), \quad (1)$$

where $\mathbf{o}_{<t,1} = [o_{1,1}, \dots, o_{t-1,1}]$, while θ_{AR} represents the AR Transformer model parameters. Different embedding projections are applied to the textual and acoustic tokens, and they are superimposed by sinusoidal positional embeddings. Note that the AR model in SpeechX is conditioned on the concatenated embeddings of both the acoustic and textual prompts. This formulation differs from that of VALL-E, where the AR model is only conditioned on the textual prompt and the past acoustic history.

After obtaining the first layer codes with the AR model, the NAR model is used to generate the l th layer codes based on the text and acoustic prompts as well as the output codes for the first $l-1$ layers, which have already been produced. The model is used repeatedly for $l = 2, \dots, 8$. Since we use the same NAR model for the remaining seven layers, the NAR model is trained to minimize the following negative log-likelihood function:

$$\mathcal{L}_{NAR} = - \sum_{l=2}^8 \log P(\mathbf{o}_{:,l} | \mathcal{T}, \mathcal{A}, \mathbf{o}_{:,<l}; \theta_{NAR}), \quad (2)$$

where θ_{NAR} represents the NAR model parameters, while $\mathbf{o}_{:,l}$ denotes the entire sequence of $o_{t,l}$ for the l th layer, and $\mathbf{o}_{:,<l} = [o_{:,1}, \dots, o_{:,l-1}]$. In this formulation, in order for the single NAR model to process each of the seven layers, the acoustic tokens from the first to $(l-1)$ th layers, $\mathbf{o}_{:,<l}$, are embedded and summed up.

C. Task-based prompting

SpeechX aims to handle multiple tasks with one model. To this end, we adopt task-based prompting, as illustrated in Table I and explained in detail below.

²<https://github.com/Kyubyong/g2p>

TABLE I
TASK-BASED PROMPTING: PROMPTS AND DESIRED OUTPUT FOR INDIVIDUAL TASKS. G2P(\cdot) DENOTES GRAPHEME-TO-PHONEME CONVERSION.

Task	Textual prompt \mathcal{T}	Acoustic prompt \mathcal{A}	Desired output \mathcal{O}
Noise suppression	G2P(text) / null	$\langle ns \rangle, C(s+n)$	$C(s)$
Speech removal	G2P(text) / null	$\langle sr \rangle, C(s+n)$	$C(n)$
Target speaker extraction	G2P(text) / null	$C(s'_1), \langle tse \rangle, C(s_1+s_2)$	$C(s_1)$
Zero-shot TTS	G2P(text)	$C(s)$	$C(s')$
Clean speech editing	G2P(text)	$C(s_{pre}), \langle soe \rangle, \langle mask \rangle, \langle eoe \rangle, C(s_{post})$	$C(s_{pre}), C(s_{edit}), C(s_{post})$
Noisy speech editing	G2P(text)	$C(s_{pre}+n_{pre}), \langle soe \rangle, C(s_{mid}+n_{mid}), \langle eoe \rangle, C(s_{post}+n_{post})$	$C(s_{pre}+n_{pre}), C(s_{edit}+n_{mid}), C(s_{post}+n_{post})$

Noise suppression is a task of extracting clean speech signal s from its noise-corrupted observation $s+n$, where n denotes the noise. For the noise suppression task, we incorporate a special token, denoted as $\langle ns \rangle$, to form the acoustic prompt, resulting in $\mathcal{A} = [\langle ns \rangle, C(s+n)]$. Here, $C(\cdot)$ denotes the function used to convert an audio signal into a neural codec token sequence. While the textual prompt \mathcal{T} is supposed to be provided by a user as a reference transcription, we let the use of the textual prompt be optional to accommodate the scenario where the human transcription is unavailable. The desired output is the acoustic token sequence of the clean audio, $C(s)$.

Speech removal involves removing speech from a noisy speech signal while preserving the background noise. It is useful for removing only unwanted speech from recordings. To address this task, we employ a special token, $\langle sr \rangle$, to construct the acoustic prompt as $\mathcal{A} = [\langle sr \rangle, C(s+n)]$. The desired output is the acoustic token sequence of the noise signal, $C(n)$. As in the case of noise suppression, the textual prompt can be omitted.

Target speaker extraction aims at isolating clean speech s_1 of a target speaker from a mixture of s_1 and interfering speech s_2 from a secondary speaker. The target speaker is identified through a short enrollment audio s'_1 of that individual, where we assumed three seconds for the enrollment. For this task, we form the acoustic prompt by concatenating the acoustic tokens extracted from the enrollment audio, $C(s'_1)$, and those of the mixed speech, $C(s_1+s_2)$, with a task-specifying token, denoted as $\langle tse \rangle$. That is, we have $\mathcal{A} = [C(s'_1), \langle tse \rangle, C(s_1+s_2)]$. The desired output is $C(s_1)$. As with the previous tasks, the inclusion of the textual prompt is optional.

Zero-shot TTS aims to generate a speech signal s' by leveraging both the provided input text and an enrollment speech s . The goal is to ensure that the speech characteristics of s' closely resemble those of s , while also accurately reflecting the input text. For this task, we employ the acoustic tokens extracted from the enrollment audio, denoted as $C(s)$, as the acoustic prompt. The model generates acoustic tokens for the synthesized speech, $C(s')$, based on the input text. These acoustic tokens are then converted into the corresponding waveform.

Clean speech editing is defined as modifying a segment of input speech to align with an input text. Let s denote the input speech signal to be edited. We divide s into three distinct portions, s_{pre} , s_{mid} , and s_{post} , with s_{mid} being the target segment for editing, without loss of generality (s_{pre} and s_{post} can be empty). We construct the acous-

tic prompt as $[C(s_{pre}), \langle soe \rangle, \langle mask \rangle, \langle eoe \rangle, C(s_{post})]$, where new tokens $\langle soe \rangle$, $\langle mask \rangle$, $\langle eoe \rangle$ are introduced to specify the task and the speech segment designated for editing. The desired output is a sequence of neural codes, $[C(s_{pre}), C(s_{edit}), C(s_{post})]$, where the spoken content of $[s_{pre}, s_{edit}, s_{post}]$ matches the input text. The speaker characteristics of s_{edit} must be consistent with those of s_{pre} and s_{post} .

Noisy speech editing, in contrast, operates on noisy speech as input, aiming to modify the speech content within a segment while keeping the underlying background noise intact. Therefore, this task would be more challenging than the clean speech editing task because the model needs to distinguish between speech and noise during the editing process. To accomplish this objective, it is crucial to provide the model with the complete input speech signal instead of masking out the segment for editing with $\langle mask \rangle$ token. Therefore, we construct the acoustic prompt as $[C(s_{pre}+n_{pre}), \langle soe \rangle, C(s_{mid}+n_{mid}), \langle eoe \rangle, C(s_{post}+n_{post})]$, with the subscripts corresponding to pre, mid, or post as previously defined. The desired output comprises a sequence of neural codes, $[C(s_{pre}+n_{pre}), C(s_{edit}+n_{mid}), C(s_{post}+n_{post})]$. This formulation makes it clear that the model must transform s_{mid} into s_{edit} based on the text input while retaining n_{mid} .

In practical speech editing scenarios, the input text is often obtained by first applying automatic speech recognition (ASR) to the input speech and then having a user edit the transcription. In such situations, it is simple to identify the positions at which $\langle soe \rangle$ and $\langle eoe \rangle$ must be inserted. Also, it is noteworthy that, in clean speech editing, the use of $\langle mask \rangle$ allows the model to adaptively change the output speech length in such a way that the output speech sounds natural in terms of speaking speed.

The outlined task-based prompting strategy equips the SpeechX model with the ability to uniquely decide the desired output during inference. This approach enables flexibility for incorporating additional tasks. Adding new tasks entails integrating corresponding prompting schemes and continuing model training from an existing checkpoint, where only embeddings for newly introduced task-specific tokens are randomly initialized. This can be performed without changing the underlying model architecture.

D. Model training

During training, we randomly sample the task for each model update at an equal probability. This is intended to ensure the model does not unduly favor any particular tasks. For noise suppression, speech removal, and target speaker extraction

tasks, we include the textual prompt at a 50% probability so that the model equally experiences both text and text-less scenarios.

To help model to acquire basic generation capabilities, we first train the model only for zero-shot TTS and then continue the training process using all the tasks to perform multi-task learning. In other words, we initialize the model with an existing VALL-E model checkpoint. Precisely speaking, the SpeechX model trained solely for zero-shot TTS exhibits slight divergence from VALL-E. This difference arises from the fact that the former explicitly incorporates a distinct enrollment audio, originating from the same speaker, for each training sample, while the latter does not. Nevertheless, for the sake of simplicity, we refer to this initialization approach as VALL-E initialization. When starting the multi-task training stage, randomly initialized embeddings are appended for the special tokens related to the task-dependent prompts. This two-stage training strategy substantially enhances performance across all tasks, as evidenced by our experimental results.

IV. EVALUATION SETUPS

Evaluating versatile speech generation models like SpeechX requires performing an array of tests, each focusing on individual tasks. To keep the experiments manageable as well as ensure consistency across the tasks, we used evaluation datasets that were derived from the test-clean split of LibriSpeech for all evaluations. In this section, we provide the details of our evaluation setups. Following previously established practices [15], [16], we selected the test samples with durations between 4 and 10 seconds.

A. Evaluation data

Zero-shot TTS: For each test sample, we used the reference transcription to create the textual prompt. The acoustic prompt was generated by randomly choosing another utterance of the same speaker and extracting a 3-second-long clip.

Noise suppression: We mixed each test sample with a noise sample randomly picked from the MUSAN dataset [34] at a signal-to-noise ratio (SNR) which was randomly determined from the range between 0 dB and 20 dB. The task was to recover the uncorrupted speech from the noisy speech. The acoustic prompt was obtained by applying EnCodec to the noisy signal. As regards the textual prompt, we considered both text-less (i.e., using no semantic prompt) and text-guided noise suppression, where we used the reference transcription for the text-guided setting.

Target speaker extraction: We mixed each test sample with an utterance of a different speaker at a signal-to-interference ratio (SIR) which was randomly determined from the range between 0 dB and 20 dB. Also, we randomly chose one or more other utterances of the same speaker to create a 3-second-long enrollment clip to help models identify who the desired speaker is. Both the mixed and enrollment signals were used to derive the acoustic prompt as described in Section III-C. The task was to recover the original uncorrupted speech

of the target speaker. As with the noise suppression task, we considered both text-less and text-guided settings.

Clean speech editing: For each test sample, we randomly selected a period of length between 10% and 50% of the whole utterance. We replaced the speech of the selected period with another randomly chosen speech sample of the same speaker. Given the partially replaced, speaker homogeneous speech and the reference transcription, the task was to generate a speech signal that follows the transcription without changing the speaker characteristics and the unreplaced portion of the input signal. In our experiments, we used the correct `<soe>` and `<eoe>` locations based on the knowledge of the replaced segment.

Noisy speech editing: We added a randomly picked MUSAN noise sample to each test sample of the clean speech editing task. The SNR was chosen from the range of 0 dB to 20 dB. Given the noise-corrupted partially replaced speech and the reference transcription, the task was to generate a noisy speech signal that follows the transcription without changing the background noise, the speaker characteristics, and the unreplaced portion of the input speech.

Speech removal: The same dataset was used as the one used for noise suppression. Given a noisy speech signal, the task was to extract the noise signal by removing the speech. We considered only the textless case. Consequently, the input exclusively comprised the acoustic prompt corresponding to the noisy speech.

B. Metrics

For consistency and reproducibility, we opted to use objective metrics for individual tasks as described below.

Word error rate (WER): We employed the WER as a metric to evaluate the fidelity of the generated audio in adhering to the provided transcription. The ASR system utilized for our experiments was NeMo’s `stt_en_conformer_transducer_large` model³, which is based on the Conformer Transducer architecture [35]. We selected this particular ASR model based on its superior stability and robustness against noise and processing artifacts in comparison to other publicly available ASR models, as was observed during our preliminary experiments. Robustness in ASR is particularly crucial for tasks such as noise suppression and noisy speech editing. The WER metric was employed across all tasks, with the exception of speech removal.

Speaker similarity score (SIM): The speaker similarity score served as a metric to assess the coherence of the generated speech in relation to the speaker’s characteristics. This score was calculated as the cosine similarity between the speaker embeddings of the generated speech and the desired speech signals. The computation of speaker embeddings was performed using NeMo’s `TitaNet-Large`⁴. We employed the original audio data instead of utilizing an EnCodec-processed signal for the speaker similarity measurement to capture and

³https://huggingface.co/nvidia/stt_en_conformer_transducer_xlarge

⁴https://huggingface.co/nvidia/speakerverification_en_titanet_large

reflect any potential speech deformation effects that may arise due to the use of the EnCodec model. SIM was used in zero-shot TTS, clean speech editing, and noisy speech editing.

DNSMOS: For evaluation in the noise suppression and target speaker extraction tasks, we utilized DNSMOS [36], a well-established model-based metric for predicting the perceived quality of acoustically corrupted speech⁵. Specifically, we employed the OVRL score from the DNSMOS P.835 model. To evaluate the performance of target speaker extraction, we employed a personalized DNSMOS model, which was tailored for this particular task and is available on the same webpage.

Perceptual Evaluation of Speech Quality (PESQ): For the noise suppression and target speaker extraction tasks, we also utilized PESQ [37]. Unlike DNSMOS, PESQ is an intrusive metric that necessitates the clean reference signals. Consequently, PESQ is expected to assess the fidelity of the generated audio with respect to the original clean data.

Mel-cepstral distortion (MCD): MCD⁶ is a metric used to quantify the dissimilarity between two sequences of mel cepstra. We employed this metric to objectively measure the speech removal accuracy by comparing the estimated noise with the ground truth noise audio.

V. EXPERIMENTS

A. Training data

We sourced clean speech data from LibriLight, comprising 60 thousand hours of untranscribed English reading speech from over 7,000 speakers [29], as was performed in the zero-shot TTS experiment using VALL-E [16]. To meet the specific training requirements for each task, data simulation was performed by following the methods employed for creating the evaluation data, as elaborated below. Note that, as discussed in Section III-D, we formed individual training mini-batches based on randomly selected tasks for each iteration.

For the noise suppression and speech removal tasks, we mixed the clean speech with noise samples from the DNS challenge corpus [38] at SNRs between -5 dB and 20 dB. Our models were trained to recover the acoustic tokens of the clean speech and noise for noise suppression and speech removal, respectively. For the target speaker extraction task, we mixed the individual clean speech samples with those of other randomly chosen speakers with SNRs ranging from -5 dB to 20 dB. As regards clean speech editing, for each clean utterance, we randomly selected a subsegment of length ranging from 10% to 70%, and then substituted it with another audio segment from the same speaker with different content. We saved the start and end times of the replaced segment, which were used to insert the <soe> and <eoe> tokens to the correct positions in the acoustic prompt during training. Furthermore, to create training samples for noisy speech editing, we added noise samples used in the noise suppression task to the partially replaced clean audio. As a

result, we obtained pairs of noisy partially replaced speech and the corresponding original noisy speech, which served as the training data for the noisy speech editing task. The SNR range used for noisy speech editing training was also [-5, 20] dB.

Since LibriLight does not provide reference transcriptions, we adopted a pseudo-labeling approach to derive the semantic prompts, i.e., the phoneme sequences of the individual training samples, by following [15], [16]. Specifically, we transcribed the LibriLight training data with an off-the-shelf Kaldi model that was trained on the 960-hour Librispeech data with 3x speed perturbation⁷.

B. Model and training configurations

Both the SpeechX AR and NAR models share the same Transformer architecture, featuring 12 layers, 16 attention heads, an embedding dimension of 1024, a feed-forward layer dimension of 4096, and a dropout rate of 0.1.

We conducted experiments employing two initialization methods: random initialization and VALL-E initialization (refer to Section III-D for details). In the random initialization scenario, we trained the SpeechX model for 800K iterations. The model optimization utilized the AdamW optimizer, with the learning rate undergoing a warm-up phase for the initial 32K updates, peaking at 5×10^{-4} , before transitioning into a linear decay phase. Conversely, with VALL-E initialization, we opted for 400K iterations, as the initial model already underwent zero-shot TTS training over 400K iterations. In this instance, the learning rate scheduler was retained, but the warm-up period was shortened to the first 20K updates.

C. Baseline expert models

We employed expert models for different tasks to establish comparison baselines. For zero-shot TTS, we utilized VALL-E by following the model configuration outlined in the original paper [16]. For the noise suppression task, we employed a non-causal Deep Complex Convolutional Recurrent Network (DCCRN) [39], which is a widely recognized model for noise suppression. Our training data for DCCRN came from Microsoft’s internal dataset, and we further fine-tuned the model using the ASR objective based on the training recipe of [40]. For target speaker extraction, we leveraged VoiceFilter [22], employing a bidirectional LSTM configuration. We relied on a publicly available implementation of VoiceFilter⁸. Finally, for speech editing, we employed A3T [20] as the baseline. The implementation of A3T that we used is also publicly accessible⁹.

D. Results

1) *Result overview:* Table. II shows the performance analysis of SpeechX in various tasks compared to the individual expert models. We can see that initializing the model parameters using an exiting VALL-E model checkpoint was beneficial across all tasks, especially in terms of WER.

⁵<https://github.com/microsoft/DNS-Challenge/tree/master/DNSMOS>

⁶<https://pyip.org/project/pymcd>

⁷<https://kaldi-asr.org/models/m13>

⁸<https://github.com/Edresson/VoiceSplit>

⁹<https://github.com/richardbaihe/a3t>

TABLE II

RESULTS FOR VARIOUS SPEECH GENERATION/TRANSFORMATION TASKS BY SPEECHX COMPARED TO EXPERT MODELS FOR INDIVIDUAL TASKS. TEXTUAL PROMPTS WERE USED FOR NOISE SUPPRESSION AND TARGET SPEAKER EXTRACTION. IN ZERO-SHOT TTS, "NO PROCESSING" ROW SHOWS THE RESULTS OF DESIRED SPEECH SIGNALS.

Model	Noise suppression			Target speaker extraction			Zero-shot TTS		Clean speech editing		Noisy speech editing		Speech removal
	WER↓	DNSMOS↑	PESQ↑	WER↓	DNSMOS↑	PESQ↑	WER↓	SIM↑	WER↓	SIM↑	WER↓	SIM↑	MCD↓
No processing	3.29	2.42	1.93	12.55	3.04	2.27	1.71	1.00	38.29	0.96	42.48	0.87	12.57
Expert model	DCCRN [39], [40]			VoiceFilter [22]			VALL-E [16]		A3T [20]		A3T [20]		N/A
	6.39	3.25	3.52	5.09	3.39	2.90	5.90	0.57	17.17	0.29	32.17	0.18	
SpeechX (random init.)	2.56	3.05	2.24	3.12	3.46	2.27	5.40	0.57	8.10	0.75	15.33	0.64	3.04
SpeechX (VALL-E init.)	2.48	3.05	2.24	2.53	3.46	2.28	4.66	0.58	5.63	0.76	13.95	0.65	3.05

TABLE III

RESULTS OF NOISE SUPPRESSION AND TARGET SPEAKER EXTRACTION WITH OR WITHOUT TEXTUAL PROMPT.

Prompt	Noise suppression			Target speaker extraction		
	WER↓	DNSMOS↑	PESQ↑	WER↓	DNSMOS↑	PESQ↑
w/ text	2.48	3.05	2.24	2.53	3.46	2.28
w/o text	6.76	3.05	2.20	5.00	3.01	2.23

In noise suppression and target speaker extraction, SpeechX exhibited superior performance in terms of WER compared to the respective expert models. Conventional regression-based noise suppression and target speaker extraction models are known to suffer from processing artifacts, which our WER results confirmed. SpeechX was able to avoid this detrimental effect thanks to the audio-text-based generation capability. On the other hand, in terms of DNSMOS and PESQ scores, it lagged behind the expert models. This can largely be attributed to the impact of the codec model used, as discussed in detail in Section V-D5. The investigation into the speech removal task revealed that SpeechX demonstrated substantial improvement in MCD, showcasing its efficacy in removing speech. These results underscore the versatility of the SpeechX model in handling enhancement-related tasks, while also highlighting the usefulness of the audio-text-based speech generation capability that SpeechX provides.

In the zero-shot TTS task, SpeechX demonstrated a slight advantage over the baseline VALL-E model in terms of WER while concurrently achieving a comparable speaker similarity score¹⁰. Furthermore, for the clean speech editing task, SpeechX exhibited significant improvement over the baseline A3T model. The WER observed in the speech editing task was slightly higher than the WER obtained in the zero-shot TTS task, even though one might anticipate that they should fall within the same range. This discrepancy could be attributed to certain test samples where the length of non-edited speech was shorter than three seconds. These results highlight that SpeechX is equally effective in tasks primarily focusing on speech generation capability, rather than transformation ability.

2) *Speech editing for clean and noisy speech:* Table II also compares the speech editing results between clean and noisy speech in terms of WER and SIM. Editing noisy speech poses

greater challenges than clean speech, as it requires modifying the spoken content while preserving background noise. This difficulty is evident from a WER gap of 38.29% vs. 42.48% observed between the clean and noisy audio signals to be edited as well as A3T’s limited WER improvement from 42.48% to 32.17%.

Nonetheless, the SpeechX model successfully edited the noisy speech, reducing the WER to 13.95% after processing. This demonstrates the model’s robustness to acoustic noise in the input signal. The high SIM score of 0.65 shows the model largely preserved speaker characteristics, even with noise present. Our observation revealed the model retained background noise, as confirmed by our provided demo samples. Fig. 2 compares mel spectrograms for two exemplary pairs of input and generated speech signals. In the first example, the input speech contained periodic noise in the middle frequency range. SpeechX preserved this background noise over the full input duration while selectively modifying only the foreground speech during the period beginning at two seconds. A similar observation can be made for the second example, wherein the alteration was applied to the first half of the speech content. In summary, the results demonstrate SpeechX model’s effectiveness at noisy speech editing while maintaining speaker identity and background noise. Future work should develop a metric to quantitatively evaluate noise cloning capability.

3) *Effectiveness of text input in noise suppression and target speaker extraction:* With SpeechX, it is feasible to perform noise suppression and target speaker extraction using solely the acoustic prompt as input. To assess the efficacy of incorporating additional text input in the SpeechX model, we conducted noise suppression and target speaker extraction experiments where we employed only the acoustic prompt as the model input. Specifically, the input for noise suppression comprised the noisy speech, while for target speaker extraction, it consisted of the mixed speech and the target speaker’s enrollment audio.

The experimental results are presented in Table III. For both tasks, omitting the text input resulted in a noticeable increase in WER, whereas the degradation in DNSMOS and PESQ scores was modest. These findings suggest that leveraging the text input was particularly beneficial for enhancing the intelligibility of the output speech. In target speaker extraction, a significant impact on the DNSMOS score was observed, indicating that the text input aids in disentangling the target

¹⁰To avoid potential confusion, it should be noted that our experimental setup corresponds to the non-continual evaluation configuration utilized in the original VALL-E work.

TABLE IV
EFFECTS OF ADDING TASKS DURING TRAINING. ZS: ZERO-SHOT, SE: SPEECH EDITING, NS: NOISE SUPPRESSION, SR: SPEECH REMOVAL, TSE: TARGET SPEAKER EXTRACTION.

Training tasks	Zero-shot TTS		Speech editing (clean/noisy)		Noise suppression		Speech removal	Target speaker extraction	
	WER↓	SIM↑	WER↓	SIM↑	WER↓	DNSMOS↑	MCD↓	WER↓	DNSMOS↑
ZS-TTS	5.90	0.57	-	-	-	-	-	-	-
ZS-TTS + SE	4.55	0.58	5.79 / 13.80	0.76 / 0.65	-	-	-	-	-
ZS-TTS + SE + NS/SR	5.11	0.57	6.91 / 13.23	0.77 / 0.66	2.59	3.03	3.04	-	-
ZS-TTS + SE + NS/SR + TSE	4.66	0.58	5.63 / 13.95	0.76 / 0.65	2.48	3.05	3.05	2.53	3.46

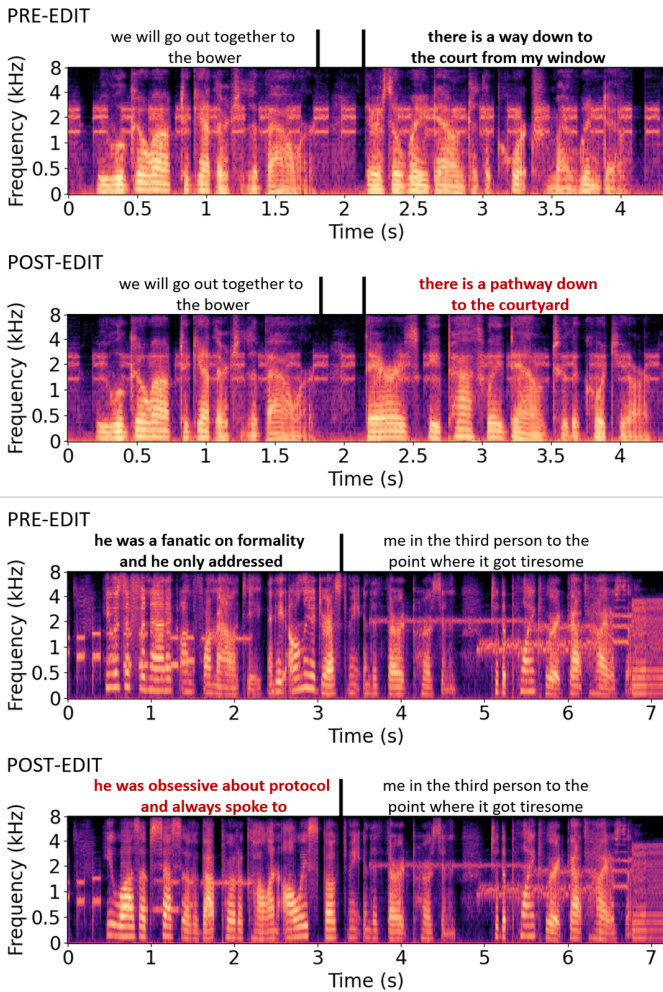


Fig. 2. Mel spectrograms of pre-edit and post-edit noisy signals. The pre-edit signal was obtained by applying EnCodec compression and decompression without intermediate processing to highlight the change made by SpeechX’s neural codec language model. See Section V-D5 for discussion on codec’s impact.

speaker’s voice from the interfering talker. Notably, while relying solely on the acoustic prompt led to WER degradation, the achieved WERs were still comparable to those of the baseline expert models.

4) *Effect of multi-task training:* We also conducted experiments where we used subsets of the tasks during training to explore potential interactions between different tasks. Specifically, in addition to the VALL-E model and the fully-trained SpeechX models that used the complete set of the tasks, we trained two additional SpeechX models: one trained exclusively for zero-shot TTS and speech editing tasks, and

TABLE V
IMPACT OF NEURAL CODEC ON PERFORMANCE METRICS FOR CLEAN AND NOISY SPEECH.

Audio type	WER↓	DNSMOS↑	PESQ↑	SIM↑
Raw clean speech	1.71	3.22	4.64	1.00
↔ EnCodec	1.81	2.97	2.69	0.81
Raw noisy speech	3.29	2.42	1.93	0.95
↔ EnCodec	5.08	2.19	1.63	0.75

the other trained on the zero-shot TTS, speech editing, noise suppression, and speech removal data.

Table IV shows the experimental results. The inclusion of speech editing during training led to an enhancement in WER for zero-shot TTS while allowing the model to learn about the speech editing task. Considering the strong parallels between zero-shot TTS and speech editing, this improvement can be attributed to the speech editing training task introducing additional variations to the distribution of the training data. Further inclusion of the noise suppression and speech removal tasks during training resulted in degradation in clean speech editing performance, while concurrently enhancing the performance for noisy speech editing. This suggests that exposing the model to noisy speech samples from these additional tasks improved the model’s robustness to acoustic noise at the expense of clean speech generation. Also, it is noteworthy that introduction of the target speaker extraction tasks to the training data did not compromise the model’s proficiency in noise suppression and speech removal.

5) *Limitation of current neural codec model:* The performance of SpeechX is inherently constrained by the accuracy of the neural codec model employed for acoustic tokenization. It should be noted that, in all previous experiments, we compared SpeechX’s results with the reference (i.e., no-processing and expert model) results obtained without any neural codec processing. To gain a more precise interpretation of SpeechX’s results, we conducted an additional experiment where we applied compression and decompression to the LibriSpeech test-clean data without any intermediate processing, measuring EnCodec’s impact on performance metrics.

Table V shows the experimental results. It is evident that processing the signals with the codec model resulted in varying degrees of performance regression across all metrics. Notably, the PESQ score dropped from 4.64 to 2.69 for the clean speech input. Our assessment indicates that while EnCodec produced slightly noticeable speech quality degradation, the significant PESQ degradation may be partly attributed to the mismatch between the PESQ algorithm and EnCodec’s training objective.

While we utilized EnCodec due to its accessibility and prior usage, future work should address this issue by developing an acoustic tokenization model more suitable for handling speech under various acoustic conditions.

VI. CONCLUSION

In this paper, we described SpeechX, a novel versatile speech generation model capable of handling diverse audio-text-based speech generation tasks, including zero-shot TTS, noise suppression, speech removal, target speaker extraction, and speech editing. For noise suppression and target speaker extraction, the proposed model provides a unified way for incorporating the knowledge of transcriptions. Also, regarding speech editing, SpeechX enables modifying the spoken content of a speech signal that contains a fair amount of background noise. SpeechX adopts a language modeling approach to generate acoustic tokens conditioned on textual and acoustic prompts, where additional task-dependent tokens are incorporated in a multi-task learning framework to support various speech transformation capabilities beyond zero-shot TTS. We demonstrated SpeechX’s efficacy through comprehensive experiments. The proposed model represents an important step toward unified generative speech models. Further research can build on this work by expanding the tasks supported, enhancing robustness, and developing more advanced conditioning mechanisms.

REFERENCES

- [1] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, “Language models are few-shot learners,” in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 1877–1901.
- [2] OpenAI, “GPT-4 technical report,” *arXiv preprint arXiv:2204.06125*, 2023.
- [3] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, “High-resolution image synthesis with latent diffusion models,” in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 10674–10685.
- [4] Z. Borsos, R. Marinier, D. Vincent, E. Kharitonov, O. Pietquin, M. Sharifi, D. Roblek, O. Teboul, D. Grangier, M. Tagliasacchi *et al.*, “AudioLM: a language modeling approach to audio generation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 2523–2533, 2023.
- [5] J.-B. Alayrac, J. Donahue, P. Luc, A. Miech, I. Barr, Y. Hasson, K. Lenc, A. Mensch, K. Millican, M. Reynolds, R. Ring, E. Rutherford, S. Cabi, T. Han, Z. Gong, S. Samangooei, M. Monteiro, J. L. Menick, S. Borgeaud, A. Brock, A. Nematzadeh, S. Sharifzadeh, M. a. Bińkowski, R. Barreira, O. Vinyals, A. Zisserman, and K. Simonyan, “Flamingo: a visual language model for few-shot learning,” in *Advances in Neural Information Processing Systems*, vol. 35, 2022, pp. 23716–23736.
- [6] J. Wang, Z. Yang, X. Hu, L. Li, K. Lin, Z. Gan, Z. Liu, C. Liu, and L. Wang, “GIT: A generative image-to-text transformer for vision and language,” *Transactions on Machine Learning Research*, 2022.
- [7] N. Ruiz, Y. Li, V. Jampani, Y. Pritch, M. Rubinstein, and K. Aberman, “DreamBooth: Fine tuning text-to-image diffusion models for subject-driven generation,” *arXiv preprint arXiv:2208.12242*, 2023.
- [8] Z. Yang, M. Khademi, Y. Xu, R. Pryzant, Y. Fang, C. Zhu, D. Chen, Y. Qian, M. Gao, Y.-L. Chen, R. Gmyr, N. Kanda, N. Codella, B. Xiao, Y. Shi, L. Yuan, T. Yoshioka, M. Zeng, and X. Huang, “i-Code V2: An autoregressive generation framework over vision, language, and speech data,” *arXiv preprint arXiv:2305.12311*, 2023.
- [9] P. K. Rubenstein, C. Asawaroengchai, D. D. Nguyen, A. Bapna, Z. Borsos, F. d. Quiry, P. Chen, D. E. Badawy, W. Han, E. Kharitonov, H. Muckenhirn, D. Padfield, J. Qin, D. Rozenberg, T. Sainath, J. Schalkwyk, M. Sharifi, M. T. Ramanovich, M. Tagliasacchi, A. Tudor, M. Velićmirović, D. Vincent, J. Yu, Y. Wang, V. Zayats, N. Zeghidour, Y. Zhang, Z. Zhang, L. Zilka, and C. Frank, “AudioPaLM: A large language model that can speak and listen,” *arXiv preprint arXiv:2306.12925*, 2023.
- [10] A. Agostinelli, T. I. Denk, Z. Borsos, J. Engel, M. Verzett, A. Caillon, Q. Huang, A. Jansen, A. Roberts, M. Tagliasacchi, M. Sharifi, N. Zeghidour, and C. Frank, “MusicLM: Generating music from text,” *arXiv preprint arXiv:2301.11325*, 2023.
- [11] E. Casanova, J. Weber, C. D. Shulby, A. C. Junior, E. Gölge, and M. A. Ponti, “YourTTS: Towards zero-shot multi-speaker TTS and zero-shot voice conversion for everyone,” in *Proceedings of the 39th International Conference on Machine Learning*, 2022, pp. 2709–2720.
- [12] Y. Jia, Y. Zhang, R. Weiss, Q. Wang, J. Shen, F. Ren, Z. Chen, P. Nguyen, R. Pang, I. Lopez Moreno, and Y. Wu, “Transfer learning from speaker verification to multispeaker text-to-speech synthesis,” in *Advances in Neural Information Processing Systems*, 2018.
- [13] E. Cooper, C.-I. Lai, Y. Yasuda, F. Fang, X. Wang, N. Chen, and J. Yamagishi, “Zero-shot multi-speaker text-to-speech with state-of-the-art neural speaker embeddings,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6184–6188.
- [14] E. Casanova, C. Shulby, E. Gölge, N. M. Müller, F. S. de Oliveira, A. Candido Jr., A. da Silva Soares, S. M. Aluisio, and M. A. Ponti, “SC-GlowTTS: An efficient zero-shot multi-speaker text-to-speech model,” in *Proc. Interspeech 2021*, 2021, pp. 3645–3649.
- [15] M. Le, A. Vyas, B. Shi, B. Karrer, L. Sari, R. Moritz, M. Williamson, V. Manohar, Y. Adi, J. Mahadeokar, and W.-N. Hsu, “Voicebox: Text-guided multilingual universal speech generation at scale,” *arXiv preprint arXiv:2306.15687*, 2023.
- [16] C. Wang, S. Chen, Y. Wu, Z. Zhang, L. Zhou, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li, L. He, S. Zhao, and F. Wei, “Neural codec language models are zero-shot text to speech synthesizers,” *arXiv preprint arXiv:2301.02111*, 2023.
- [17] E. Kharitonov, D. Vincent, Z. Borsos, R. Marinier, S. Girgin, O. Pietquin, M. Sharifi, M. Tagliasacchi, and N. Zeghidour, “Speak, read and prompt: High-fidelity text-to-speech with minimal supervision,” *arXiv preprint arXiv:2302.03540*, 2023.
- [18] R. Huang, C. Zhang, Y. Wang, D. Yang, L. Liu, Z. Ye, Z. Jiang, C. Weng, Z. Zhao, and D. Yu, “Make-A-Voice: Unified voice synthesis with discrete representation,” *arXiv preprint arXiv:2305.19269*, 2023.
- [19] Z. Zhang, L. Zhou, C. Wang, S. Chen, Y. Wu, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li, L. He, S. Zhao, and F. Wei, “Speak foreign languages with your own voice: Cross-lingual neural codec language modeling,” *arXiv preprint arXiv:2303.03926*, 2023.
- [20] H. Bai, R. Zheng, J. Chen, M. Ma, X. Li, and L. Huang, “A³T: Alignment-aware acoustic and text pretraining for speech synthesis and editing,” in *Proceedings of the 39th International Conference on Machine Learning*, 2022, pp. 1399–1411.
- [21] Z. Jiang, Y. Ren, Z. Ye, J. Liu, C. Zhang, Q. Yang, S. Ji, R. Huang, C. Wang, X. Yin *et al.*, “Mega-TTS: Zero-shot text-to-speech at scale with intrinsic inductive bias,” *arXiv preprint arXiv:2306.03509*, 2023.
- [22] Q. Wang, H. Muckenhirn, K. Wilson, P. Sridhar, Z. Wu, J. Hershey, R. A. Saurous, R. J. Weiss, Y. Jia, and I. L. Moreno, “VoiceFilter: Targeted voice separation by speaker-conditioned spectrogram masking,” *arXiv preprint arXiv:1810.04826*, 2018.
- [23] K. Žmolíková, M. Delcroix, K. Kinoshita, T. Ochiai, T. Nakatani, L. Burget, and J. Černocký, “SpeakerBeam: Speaker aware neural network for target speaker extraction in speech mixtures,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 4, pp. 800–814, 2019.
- [24] S. E. Eskimez, T. Yoshioka, H. Wang, X. Wang, Z. Chen, and X. Huang, “Personalized speech enhancement: new models and comprehensive evaluation,” in *2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 356–360.
- [25] K. Zmolikova, M. Delcroix, T. Ochiai, K. Kinoshita, J. Černocký, and D. Yu, “Neural target speech extraction: An overview,” *IEEE Signal Processing Magazine*, vol. 40, no. 3, pp. 8–29, 2023.
- [26] J. Serra, S. Pascual, J. Pons, R. O. Araz, and D. Scaini, “Universal speech enhancement with score-based diffusion,” *arXiv preprint arXiv:2206.03065*, 2022.
- [27] K. Kinoshita, M. Delcroix, A. Ogawa, and T. Nakatani, “Text-informed speech enhancement with deep neural networks,” in *Proc. Interspeech 2015*, 2015, pp. 1760–1764.

- [28] K. Schulze-Forster, C. S. J. Doire, G. Richard, and R. Badeau, "Joint phoneme alignment and text-informed speech separation on highly corrupted speech," in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7274–7278.
- [29] J. Kahn, M. Rivière, W. Zheng, E. Kharitonov, Q. Xu, P.-E. Mazaré, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen, T. Likhomanenko, G. Synnaeve, A. Joulin, A. Mohamed, and E. Dupoux, "Libri-Light: A benchmark for ASR with limited or no supervision," in *ICASSP 2020*, 2020, pp. 7669–7673.
- [30] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever, "Zero-shot text-to-image generation," in *Proceedings of the 38th International Conference on Machine Learning*, 2021, pp. 8821–8831.
- [31] K. Shen, Z. Ju, X. Tan, Y. Liu, Y. Leng, L. He, T. Qin, S. Zhao, and J. Bian, "NaturalSpeech 2: Latent diffusion models are natural and zero-shot speech and singing synthesizers," *arXiv preprint arXiv:2304.09116*, 2023.
- [32] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, "High fidelity neural audio compression," *arXiv preprint arXiv:2210.13438*, 2022.
- [33] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [34] D. Snyder, G. Chen, and D. Povey, "MUSAN: A music, speech, and noise corpus," *arXiv preprint arXiv:1510.08484*, 2015.
- [35] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented transformer for speech recognition," in *Proc. Interspeech 2020*, 2020, pp. 5036–5040.
- [36] C. K. Reddy, V. Gopal, and R. Cutler, "Dnsmos: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors," in *2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6493–6497.
- [37] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, "Perceptual evaluation of speech quality (PESQ) – a new method for speech quality assessment of telephone networks and codecs," in *2001 IEEE international conference on acoustics, speech, and signal processing. Proceedings*, vol. 2, 2001, pp. 749–752.
- [38] H. Dubey, A. Aazami, V. Gopal, B. Naderi, S. Braun, R. Cutler, A. Ju, M. Zohourian, M. Tang, H. Gamper, M. Golestaneh, and R. Aichner, "ICASSP 2023 deep noise suppression challenge," *arXiv preprint arXiv:2303.11510*, 2023.
- [39] Y. Hu, Y. Liu, S. Lv, M. Xing, S. Zhang, Y. Fu, J. Wu, B. Zhang, and L. Xie, "DCCRN: Deep complex convolution recurrent network for phase-aware speech enhancement," in *Proc. Interspeech 2020*, 2020, pp. 2472–2476.
- [40] S. E. Eskimez, X. Wang, M. Tang, H. Yang, Z. Zhu, Z. Chen, H. Wang, and T. Yoshioka, "Human listening and live captioning: Multi-task training for speech enhancement," in *Proc. Interspeech 2021*, 2021, pp. 2686–2690.