

A³T: Alignment-Aware Acoustic and Text Pretraining for Speech Synthesis and Editing

He Bai^{*1} Renjie Zheng^{*2} Junkun Chen³ Xintong Li² Mingbo Ma² Liang Huang³

Abstract

Recently, speech representation learning has improved many speech-related tasks such as speech recognition, speech classification, and speech-to-text translation. However, all the above tasks are in the direction of *speech understanding*, but for the inverse direction, *speech synthesis*, the potential of representation learning is yet to be realized, due to the challenging nature of generating high-quality speech. To address this problem, we propose our framework, Alignment-Aware Acoustic-Text Pretraining (A³T), which reconstructs masked acoustic signals with text input and acoustic-text alignment during training. In this way, the pretrained model can generate high quality reconstructed spectrogram, which can be applied to the speech editing and unseen speaker TTS directly. Experiments show A³T outperforms SOTA models on speech editing, and improves multi-speaker speech synthesis without the external speaker verification model.¹

1. Introduction

Recently, speech representation learning has attracted much attention in the speech community due to its strong performance to many speech-related downstream tasks, such as speech recognition, speech classification, and speech translation (Baeovski et al., 2020; Chen et al., 2020; Liu et al., 2020; Zheng et al., 2021; Hsu et al., 2021)

However, all these efforts can only support *speech under-*

^{*}Equal contribution ¹University of Waterloo, Waterloo, ON, Canada (work done at Baidu Research USA) ²Baidu Research, Sunnyvale, CA, USA ³Oregon State University, Corvallis, OR, USA. Correspondence to: Renjie Zheng <zrenj11@gmail.com>, Liang Huang <liang.huang.sh@gmail.com>.

Proceedings of the 39th International Conference on Machine Learning, Baltimore, Maryland, USA, PMLR 162, 2022. Copyright 2022 by the author(s).

¹See demos at <https://educated-toothpaste-462.notion.site/Demo-b0edd300e6004c508744c6259369a468>. Code available at: <https://github.com/richardbaihe/a3t>

Model	Data		Reconstructed Masked Speech	Tasks	
	$\langle s, x \rangle$	$\langle s \rangle$		Speech Editing	Text-to-Speech
wav2vec 2.0		✓	discrete units		
Hubert		✓			
SLAM	✓	✓			
MAM		✓	low-quality spectrogram		
FAT-MLM	✓	✓			
A ³ T	✓	✓	high-quality spectrogram	✓	✓

Table 1. Comparisons of A³T with other existing speech pretraining models. Here s stands for speech input, while x stands for text, and $\langle s, x \rangle$ denotes parallel speech-text data.

standing tasks which take speech as input, but for the inverse direction, *speech synthesis*, which synthesis speech as output, the potential of representation learning is yet to be realized. For example, one line of work, such as wav2vec 2.0 (Baeovski et al., 2020), Hubert (Hsu et al., 2021) and SLAM (Bapna et al., 2021), learn discrete quantized speech units as latent representations. In this way, these models are good at recognizing and extracting discrete information from speech and successfully improves automatic speech recognition (ASR), but they are unable to generate continuous acoustic signals for speech synthesis. On the other hand, another line of work, such as MAM (Chen et al., 2020) and FAT-MLM (Zheng et al., 2021), show that reconstructing masked spectrogram with continuous units can improve speech-to-text translation. However, the quality of their proposed speech reconstruction is far from the requirement of speech synthesis tasks (see Fig. 5(f)).

To address this problem, we propose our framework, Alignment-Aware Acoustic-Text Pretraining (A³T), where we introduce cross-modal alignment embeddings which make the model easier to learn the alignment between the acoustic and phoneme input during multi-modal pretraining, and significantly improve the quality of the reconstructed acoustic signals. Different from the segment embeddings used in Segatron and SegBERT (Bai et al., 2021; 2022), which improve language modeling by grouping tokens according to the sentence position and paragraph position, our alignment embeddings can align a phoneme and its frames together to learn the cross-modal self-attention (Fig. 6).

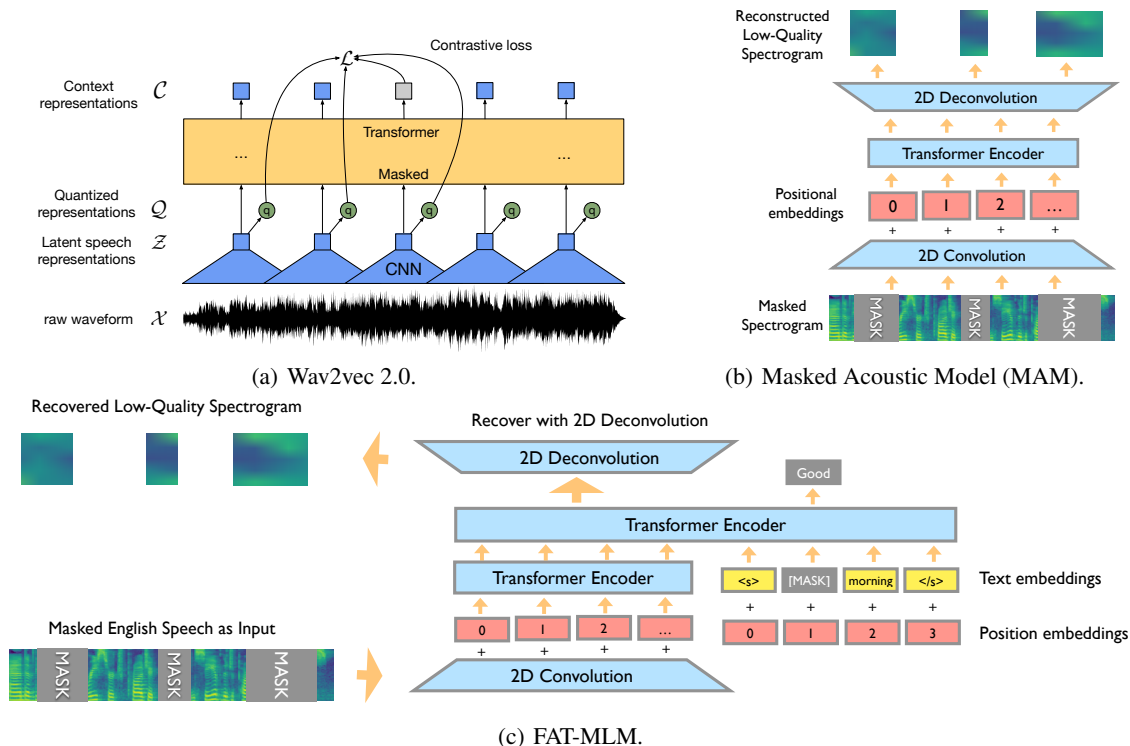


Figure 1. Previous work for speech representation learning.

Moreover, we borrow several useful ideas from recent text-to-speech (TTS) literature, including Conformer (Gulati et al., 2020; Guo et al., 2021) and Post-Net (Shen et al., 2018b), to further improve the quality of our reconstructed spectrograms.

Without any finetuning, the proposed model can be adopted as a speech-editing system, a task that modifies an existing speech, by reconstructing the desired acoustic signals given original contextual speech and modified text. Furthermore, the model can be adopted as a multi-speaker TTS system with our proposed prompt-based decoding method, to synthesis unseen speaker’s speech without the external speaker verification model (speaker embeddings). Our experiments show that our A³T with prompt-based decoding can outperform the TTS model equipped with both the speaker embedding (Jia et al., 2018) and the global style token (GST) (Wang et al., 2018).

We make the following contributions:

- We propose the Alignment-Aware Acoustic-Text Pretraining (A³T), a BERT-style pretraining model, which takes both phonemes and partially-masked spectrograms as inputs. It can reconstruct masked spectrograms with high quality without finetuning and uses the identical framework for decoding.
- We show that the proposed A³T model has the ability to do speech editing and outperforms the current SOTA.

- We propose the prompt-based decoding method. We show that our A³T model has the ability to do speech synthesis for unseen speaker and outperforms the speaker-embedding-based multi-speaker TTS system.

2. Previous Work

2.1. Speech Synthesis and Editing

Recently, neural TTS systems become capable of generating audios with high naturalness (Oord et al., 2016; Shen et al., 2018a; Ren et al., 2019; Peng et al., 2019; Ren et al., 2020). SOTA neural TTS systems generally consist of two stages: the *text-to-spectrogram* stage which generates an intermediate acoustic representation (linear- or mel-spectrogram) from the text, and the *spectrogram-to-wave* stage (vocoder) which converts the aforementioned acoustic representation into actual wave signals (Oord et al., 2018; Prenger et al., 2019).²

In the multi-speaker and unseen-speaker settings, the existing TTS models need to be trained with an additional input feature: speaker embedding (Jia et al., 2018), which is extracted from an external speaker verification model trained with tens of thousands of speakers’ audio. And during the inference for an unseen speaker, the embedding will be extracted from one of this speaker’s other audio exam-

²We focus on the *text-to-spectrogram* stage and use an off-the-shelf vocoder Parallel WaveGAN (Yamamoto et al., 2020).

ples. However, the embedding from the speaker verification model is not optimized directly to capture speaker characteristics relevant to synthesis, and cannot provide enough information for the TTS model to generate audio similar to the example.

The input of speech editing includes the original speech, the original and modified text. Jin et al. (2017) propose to insert a regenerated audio clip back into the original recording. However, due to the absence of speech contextual information, the boundaries of the modified region would be not smooth. Morrison et al. (2021) propose to retrieve the modified speech segments from other utterances of the same speaker and correct the prosody with a context-aware TD-PSOLA corrector (Moulines & Charpentier, 1990). However, the edited content may not be found in the speech data of the same speaker. Most recently, (Tan et al., 2021) use neural TTS model to generate better-modified speech. This method is only compatible with auto-regressive decoding models and highly relies on the speaker embeddings, which limits its efficiency and transferability to new speakers.

2.2. Speech Pretraining

To improve the Text-to-Speech model from larger-scale pure speech data, one idea is to do pretraining on speech data. All existing speech pretraining work learn either discrete units, which can only support speech understanding tasks, or spectrogram, but with very low quality.

2.2.1. RECONSTRUCTING DISCRETE UNITS


Wav2vec 2.0 proposed by Baevski et al. (2020) is the most popular speech pretrain model recently. It masks the speech input in the latent space and pretrains the model by predicting discrete units via a contrastive task defined over a quantization of the latent representations, as shown in Fig. 1(a). Similar to wav2vec 2.0, Hubert (Hsu et al., 2021) and SLAM (Bapna et al., 2021) also learn discrete speech units from contextualized representations to represent the latent representations. Thus these models can achieve good performance in speech recognition tasks, but they are unable to generate continuous acoustic signals for speech synthesis.

2.2.2. RECONSTRUCTING LOW-QUALITY SPECTROGRAM

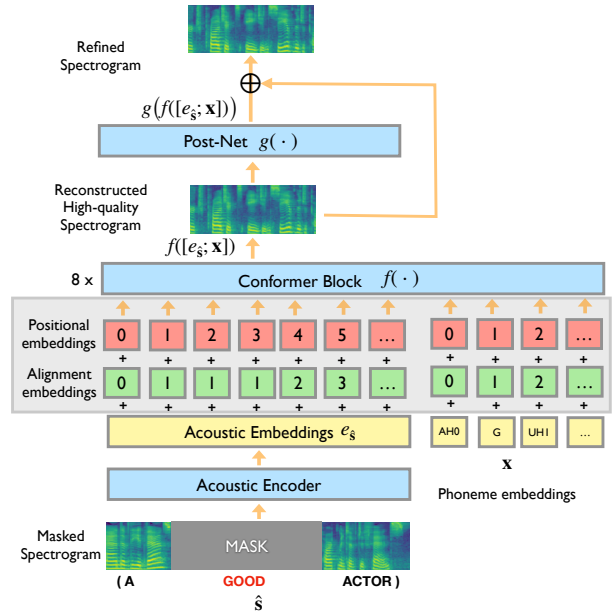
Recently, Chen et al. (2020) propose to learn a speech encoder in a self-supervised fashion on the speech side, which can utilize speech data without transcription. Fig. 1(b) demonstrate the architecture of this model, termed Masked Acoustic Modeling (MAM). MAM replaces a span of speech spectrogram with mask tokens, and learns to recover the masked spectrogram during training. On the other hand, Zheng et al. (2021) propose a Fused Acoustic and Text Masked Language Model (FAT-MLM) which jointly

learns a unified representation for both acoustic and text input from various types of corpora including parallel data for speech recognition and machine translation, and even pure speech and text data, as shown in Fig. 1(c).


Both MAM and FAT-MLM reconstruct spectrograms, however, the quality of their spectrogram output is far from the requirement of speech synthesis tasks (see Fig. 5(f)), since these pretrained models are all used in speech understanding task (speech-to-text translation), where the quality of the reconstructed spectrogram is not very important.




(a) Forced Alignment Preprocessing.



(b) A³T Model. Segment embeddings are assigned according to the Segment Index produced by forced alignment preprocessing (see subfigure (a)).



(c) Conformer Block.



(d) Post-Net.

Figure 2. Alignment-Aware Acoustic-Text Pretraining (A³T).

3. Alignment-Aware Acoustic-Text Pretraining

Although existing speech pretraining models show a strong representation learning ability and significantly improve upon many down-stream tasks in *speech understanding*, all these efforts can not support *speech synthesis* tasks. To address this problem, we propose the Alignment-Aware Acoustic-Text Pretraining (A³T) which learns to generate high-quality spectrogram given speech context and text.

3.1. A³T

A³T takes speech and transcription tuples as input, denotes as $D_{s,x} = \{(\mathbf{s}, \mathbf{x})^{(n)}\}_{n=1}^{|D|}$, where $\mathbf{s} = (s_1, \dots, s_{|\mathbf{s}|})$ is a sequence of acoustic features $s_i \in \mathbb{R}^{d_s}$ which can be the spectrogram or mel-spectrogram of the speech audio, and each s_i represents the frame-level speech feature, and $\mathbf{x} = (x_1, \dots, x_{|\mathbf{x}|})$ is the sequence of corresponding transcription.

As shown in Fig. 2, we first randomly mask several spans of \mathbf{s} by a random masking function over the input \mathbf{s} : $\hat{\mathbf{s}} \sim \text{Mask}_{\text{span}}(\mathbf{s}, \lambda)$, where $\text{Mask}_{\text{span}}(\cdot)$ replaces several random spans of \mathbf{s} by the probability of λ with the same number of a random initialized masking vector $\epsilon_s \in \mathbb{R}^{d_s}$. Then we encode $\hat{\mathbf{s}}$ with a acoustic encoder for acoustic embeddings $e_{\hat{\mathbf{s}}}$. In this work, we use a nonlinear feed-forward layer as the acoustic encoder.

3.2. Cross-modal Alignment Embedding

To strengthen the interaction between the speech and text input, we introduce cross-modal alignment embedding as one input of encoder, where we sum the i th acoustic embedding e_{s_i} or text embedding x_i with its positional embedding e_{pos_i} and alignment embedding e_{aln_i} all together: $e_{s_i} + e_{\text{pos}_i} + e_{\text{aln}_i}$, where previous work have proved the embedding sum operation is simple and effective (Devlin et al., 2018; Bai et al., 2021). After that, the phoneme embedding and its acoustic embeddings will share the same alignment embedding. We use a forced aligner (Yuan & Liberman, 2008) to pre-process the dataset to get the alignment information, which is shown in Fig. 2(a).

3.3. Conformer

Given the recent success of Convolution-augmented Transformer (Conformer) on various speech tasks (Gulati et al., 2020; Guo et al., 2021), we adopt Conformer as the backbone of our encoder and decoder. Compared with Transformer, Conformer introduces a convolution module and an additional feedforward module, which is shown in Fig. 2(c). In our experiments, we find Conformer is better than Transformer for acoustic-text pretraining.

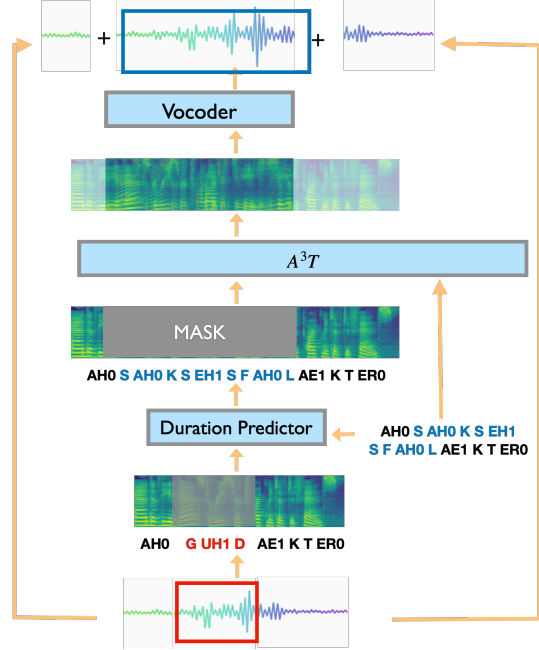


Figure 3. Speech editing pipeline.

3.4. Post-Net and Loss Function

We follow Tacotron 2 (Shen et al., 2018b) to use Post-Net to refine the generated spectrogram. The predicted spectrogram is passed through a 5-layer convolution Post-Net to be refined as shown in Fig. 2(d).

The training objective of multi-modal A³T includes a speech reconstruction loss $\ell_s(D_{s,x})$ which takes a spectrogram \mathbf{s} and a text sequence \mathbf{x} as input. We have the following training objective to reconstruct the original speech signal with the surrounding context information:³

$$\ell_s(D_{s,x}) = \sum_{(\mathbf{s}, \mathbf{x}) \in D_{s,x}} \underbrace{\|f([e_{\hat{\mathbf{s}}}; \mathbf{x}]) + g(f([e_{\hat{\mathbf{s}}}; \mathbf{x}])) - \mathbf{s}\|_1}_{\text{refined spectrogram}} + \underbrace{\|f([e_{\hat{\mathbf{s}}}; \mathbf{x}]) - \mathbf{s}\|_1}_{\text{reconstructed spectrogram}} \quad (1)$$

where g is a Post-Net which tries to recover a better original signal from encoded representation $f([e_{\hat{\mathbf{s}}}; \mathbf{x}])$. We use mean absolute error (MAE) for measuring the difference between \mathbf{s} and the reconstructed spectrogram.

3.5. A³T for Speech Editing

Once A³T finishes the pretraining process, it can be used as a speech editing system directly with an external duration

³Similar with previous work using masked language model objective, this loss only takes the masked input into consideration.




Figure 4. Illustrations for one-shot TTS. The prompt speech and text are wrapped with blue rectangles, and the target speech and text are wrapped with red.

predictor, which is shown in Fig. 3.

Given a speech s , its original phonemes \tilde{x} , and the target modified phonemes x , our system first finds the phonemes that need to be modified \hat{x} . To predict the duration \hat{d}_i of modified phonemes $x_i \in \hat{x}$, we use an external duration predictor. Since the duration predictor only takes phoneme sequence as input, we adjust model predicted duration d'_i according to the origin speech durations as: $\hat{d}_i = d'_i \sum_{j=1}^{|\tilde{x}|} \frac{\tilde{d}_j}{d'_j}$, where \tilde{d}_j/d'_j is the ratio between the duration of the original phone \tilde{x}_j in the given speech and the predicted duration of the original phone \tilde{x}_j in the original sentence \tilde{x} . We compute this ratio from all phones in the original sentence \tilde{x} and use it to adjust the predicted duration d'_i and get the final duration \hat{d}_i for phone \hat{x}_i .

With this predicted duration, we insert $\sum_{i=1}^{|\hat{x}|} \hat{d}_i \cdot sr/h^4$ number of [MASK] frames into the non-modified spectrogram context. This masked spectrogram \hat{s} is the input of A^3T and $f([e_{\hat{s}}; \mathbf{x}]) + g([e_{\hat{s}}; \mathbf{x}])$ is the prediction of the modified portion spectrogram. Then, we use a vocoder to generate the waveform of this spectrogram and output the final edited speech by replacing the modified part of the original speech.

⁴ sr stands for sample rate and h stands for hop size.

3.6. A^3T for Multi-speaker TTS

In addition to the speech editing, we find our model has the potential for unseen speaker TTS.

Existing popular unseen speaker TTS models (Jia et al., 2018) are trained with seen speaker embeddings and generalizes to unseen speaker embeddings during the inference. However, such speaker embeddings are extracted from an external speaker verification model which is trained with tens of thousands of speakers.

In this work, we find our model can achieve comparable naturalness to models with speaker embeddings for unseen speaker TTS task; What’s more, our generations are more similar to the unseen speaker’s reference speech. The illustrations of how to synthesis speech for unseen speakers with our A^3T model are shown in Fig. 4, which is named prompt-based A^3T .

The key idea is to concatenate the prompt and the target together into a new utterance input, where the target speech is consist of n [MASK] and n is predicted by a duration predictor. By inputting the concatenated speech and text, A^3T model will predict the spectrogram of these masked frames. The role of the reference text and speech in our model is similar to prompts in language model (Brown et al., 2020), and hence we call it prompt-based decoding/generation.

4. Experiments

In this section, we introduce our experiments for spectrogram reconstruction pretraining task, speech-editing task, and multi-speaker TTS. The spectrogram reconstruction is our pretraining task, where we conduct ablation study to show the contributions of different components and also the effects of different masking rates. The experiment settings of speech-editing are followed Tan et al. (2021), where we deploy two speech-editing systems with two datasets and evaluate the Mel-cepstral distortion (MCD) score and human-annotated mean opinion score (MOS) (Chu & Peng, 2006) using Amazon Mechanical Turk. The multi-speaker TTS experiments include seen speaker TTS and unseen speaker TTS evaluated with the MOS scores.

4.1. Datasets

Following Tan et al. (2021), we conduct our speech-editing experiments with a single-speaker TTS dataset LJSpeech (Ito & Johnson, 2017) and a multi-speaker TTS dataset VCTK (Yamagishi et al., 2019). The LJSpeech dataset is a single-speaker dataset with 13K examples in 24 hours. The VCTK dataset is a multi-speaker dataset with 109 speakers and 44K examples in 44 hours. It should be noted that after finishing the pretraining process with LJSpeech or VCTK, our A^3T will be used as a speech-

editing system without any further finetuning.

We test multi-speaker TTS task with VCTK dataset. For seen multi-speaker TTS, each speaker’s examples would be split into train and test sets. For unseen multi-speaker TTS, the test set contains 10 speakers’ examples, and the other 99 speaker’s examples are used for training.

4.2. Configuration Details

Raw audio files are processed with 50 ms frame size and 12.5 ms frame hop with the Hann window function to extract 80-dimensional log-Mel filterbanks. We use 24K sampling rate for VCTK and 22K for LJSpeech. The forced alignment and G2P are both carried out by HTK (Young et al., 2002) to convert English words to phones and align phones with audio segments. For speech-editing systems and prompt-based TTS, we use the publicly available duration predictor from FastSpeech 2 implemented in ESPnet (Inaguma et al., 2020). We use Parallel-WaveGAN (Yamamoto et al., 2020) vocoder for all the systems.

All A³T models pretrained in our experiments share the same architecture: 4 layers Conformer encoder, 4 layers Conformer decoder, and 5 layers Conv1d Post-Net, with 2 heads multi-head attention in 384-dim. The convolution kernel sizes of the encoder and decoder are 7 and 31, respectively. The shape of alignment embeddings is (500, 384), where we assume the number of phones will not exceed 500 for a single input. The shape of input phone embeddings is (73, 384), and we use a ReLU (Agarap, 2018) nonlinear layer to transform 80-dim log-Mel filterbanks features to 384-dim. The total number of parameters is 67.7M.

During training, we use Adam optimizer with a 1.0 initial learning rate, 4000 warmup steps, and Noam learning rate scheduler. Instead of setting a fixed batch size, we adjust the batch size according to the length of the input example and set a maximum batch-bin (the total number of input elements) for each model. Following MAM (Chen et al., 2020), 15% frames will be masked for speech-only input, For speech-text input, we randomly select several phonemes spans (80% phonemes) and mask their corresponding frames. For speech-editing experiments, we use 2.4M batch-bin, 1M steps for LJSpeech, and 3M batch-bin, 1.2M steps for VCTK.

4.3. Ablation Study with Spectrogram Reconstruction

We first conduct an ablation study with LJSpeech dataset for our pretraining task: spectrogram reconstruction. This task requires A³T to predict the masked frames. We sample 30 utterances randomly from the test set, and 1/3 phones in the middle of each sentence are masked. We adopt MCD to measure the difference between the ground-truth audio and the reconstructed audio, where we only measure the MCD

Example	Model	MCD ↓
Fig. 5(b)	A ³ T	8.09
Fig. 5(c)	- Alignment Embeddings	10.73
Fig. 5(d)	- Conformer	12.43
Fig. 5(e)	- Post-Net	12.94
Fig. 5(f)	- L1 loss	11.55

Table 2. Ablation study for A³T pretrained with LJSpeech. We remove modules one by one. For each model, we show a spectrogram example in Fig. 5.

MaskRate	Seen MCD ↓	Unseen MCD ↓
20%	10.35	12.22
50%	7.75	9.99
80%	8.72	9.68


Table 3. MCD scores of A³T pretrained in different masking rates with VCTK.

of the masked region an lower MCD means higher similarity. We incrementally discard the components of A³T: removing the cross-modal alignment embedding, replacing the Conformer with Transformer, removing the Post-Net, and using L2 (MSE) loss instead of L1 (MAE) loss.

Results are shown in Tab. 2. An example of different models’ reconstruction is shown in Fig. 5. By comparing Fig. 5(b) and Fig. 5(c), we can see that many details are lost when A³T trained without the alignment embedding, and the MCD scores rise from 8.09 to 10.73. Similar degrading can be observed after replacing Conformer with Transformer: the MCD scores rise from 10.73 to 12.43 and the spectrogram becomes blurrier (Fig. 5(d)). Compared with the alignment embedding and Conformer, Post-Net contributes only 0.49 MCD score, and L2 loss even achieves better MCD score than L1 loss. However, when looking into the spectrograms, we can see that Fig. 5(f) is blurrier than Fig. 5(e), which conforms to the previous finding (Klimkov et al., 2018) that L1 loss is better than L2 loss for speech synthesis. Hence, we choose L1 loss for A³T pretraining. Also, Fig. 5(f) indicates the quality that previous pretrained model (MAM/FAT-MAM) could achieve, and the other figures show how our A³T transforms Fig. 5(f) to Fig. 5(b).

We also conduct a study with VCTK to show the impacts of difference masking rates. Results are shown in Tab. 3. We can see that 20% masking rate leads to large MCD scores, while 50% and 80% are better. Also, 50% masking rate outperforms 80% on the seen test cases, but not on the unseen. Considering 80% masking rate has a better generalization on unseen cases, we choose 80% for all the following experiments.

Finally, we plot the attention heat maps of encoder with and



(a) Groundtruth spectrogram from LJSpeech.







 (b) Reconstructed spectrogram by A^3T .

 (c) Reconstructed spectrogram based on A^3T in (b) without segment embedding.

 (d) Reconstructed spectrogram based on A^3T in (c) with Transformer instead of Conformer.

 (e) Reconstructed spectrogram based on A^3T in (d) without Post-Net.

 (f) Reconstructed spectrogram based on A^3T in (e) with L2 loss instead of L1 loss. This model uses the similar architecture of FAT-MLM.

Figure 5. An example of ablation study in LJSpeech. Original text is “and of the **Advanced Research** Projects Agency of the Department of Defense”. The portion with red box is “**Advanced Research**” which is masked in (b,c,d,e,f) subfigures.


without our proposed cross-modal alignment embedding in Fig. 6. The attention matrices are collected from the encoder’s last layer with a mean-pooling across heads. It should be noted that the original attention matrix is 310×310 , which contains both the speech and phones, and for clarity, we plot only 11 phones and their corresponding frames in Fig. 6. We can see that our A^3T is aware of the speech segmentations and their corresponding phones, while the baseline model fails to capture such alignment information. This observation demonstrates the effectiveness of our A^3T for cross-modal pretraining. This observation also conforms previous finding that Transformer-based language model

cannot align the tokens within the same sentence/paragraph together, even pre-trained with the BERT-large setting (Bai et al., 2021).


4.4. Speech Editing

Following Tan et al. (2021), we list several baseline systems below:

- **Baseline 1:** This is a TTS system regenerating a complete waveform from the whole sentence to be edited.
- **Baseline 2:** This system generates the modified region




(a) Attention map of A^3T w/o alignment embeddings



(b) Attention map of A^3T with alignment embeddings

Figure 6. Attention map between speech and text of A^3T with and without alignment embeddings.

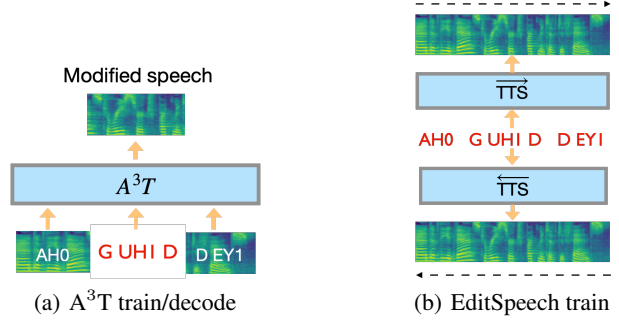


(a) Baseline 1. (b) Baseline 2. (c) Baseline 3.

Figure 7. Illustrations for speech editing baselines.


with a TTS model and insert the generation back to the original waveform with a forced aligner.

- **Baseline 3:** This system is similar to Baseline 1, but we cut the modified region from the generation and insert it back to the original waveform with a forced aligner.
- **Tan et al. (2021):** This is a speech-editing system which introduces partial inference and bidirectional fusion to sequence-to-sequence neural TTS model. EditSpeech trains two conventional autoregressive TTS models, one left-to-right and the other right-to-left (Fig. 8(b)). For decoding, the left-to-right TTS model



(a) A^3T train/decode

(b) EditSpeech train



(c) EditSpeech decode

Figure 8. Comparisons between A^3T and EditSpeech. \rightarrow : free decoding, \dashrightarrow : forced decoding. A^3T trains a non-autoregressive encoder to reconstruct masked acoustic signals, and uses the identical framework for decoding. EditSpeech trains two autoregressive TTS models: a left-to-right and a right-to-left. For decoding, these two models synthesize two speeches and are fused for the output.

force-decodes the prefix speech context and synthesizes the modified region, and the right-to-left TTS model force-decodes the suffix context and generates the modified region reversely. Finally, the two synthesized speeches are fused for final output (Fig. 8(c)).

Following Tan et al. (2021), we also evaluate our speech editing system with an identical reconstruction task, which is similar to the above ablation experiments but without the ground-truth duration length and can be evaluated with MCD metric. 30 utterances are randomly sampled for each dataset, and a part of speech, which corresponds to 1/3 phonemes in the middle of each sentence, is masked. The audio of the masked region is replaced with each system’s generation. A duration model is used to predict the length of masked speech from phonemes. Results are shown in Tab. 4. From this table, we can see that our system achieves the best MCD score. Besides, alignment embedding is the key to reducing MCD, which confirms our observation in Fig. 5(c). For TTS-based systems, we find that generating the whole audio and then extracting the modified region is better than generating the modified region only.

We then conduct the human evaluation with Amazon Mechanical Turk for the real speech insertion and replacement tasks using the VCTK dataset. To compare our results with Tan et al. (2021), we use the same 15 audio samples and modification operations from their work. For each audio sample, we use 10 English native speakers to evaluate the naturalness of synthesized audios. In Tab.5, our A³T speech editing system outperforms Tan et al. (2021)’s and gets the highest MOS scores among all these systems. Audio examples can be found at our demo link.

Model	VCTK MCD ↓	LJSpeech MCD ↓
Baseline 1/3	10.66	10.32
Baseline 2	12.06	10.91
A ³ T	7.76	9.26
w/o Alignment Emb.	11.37	10.30

Table 4. MCD evaluation on identity speech reconstruction using VCTK and LJSpeech.

Model	Insert	Replace
Baseline 1	3.02 ± 0.20	2.64 ± 0.16
Baseline 2	2.89 ± 0.17	2.70 ± 0.16
Baseline 3	2.89 ± 0.17	2.44 ± 0.16
Tan et al. (2021)	3.50 ± 0.16	3.58 ± 0.16
A ³ T	3.53 ± 0.17	3.65 ± 0.15
w/o Alignment Emb.	2.48 ± 0.21	1.98 ± 0.17

Table 5. The MOS evaluation (↑) on speech editing task on VCTK with 95% confidence intervals.

Model	Seen	Unseen
FastSpeech 2	3.33 ± 0.10	3.78 ± 0.10
+GST (Wang et al., 2018)	3.42 ± 0.10	3.81 ± 0.11
A ³ T	3.61 ± 0.09	3.90 ± 0.10
Groundtruth	3.94 ± 0.08	4.09 ± 0.10

Table 6. The MOS evaluation (↑) for **speaker similarity** on multi-speaker TTS on VCTK with 95% confidence intervals. The FastSpeech2 model is equipped with X-vectors (Snyder et al., 2018).

Model	Seen	Unseen
FastSpeech 2	3.34 ± 0.11	3.85 ± 0.11
+GST (Wang et al., 2018)	3.27 ± 0.11	3.72 ± 0.11
A ³ T	3.63 ± 0.10	3.94 ± 0.11
Groundtruth	4.04 ± 0.08	4.05 ± 0.10

Table 7. The MOS evaluation (↑) for **speech quality** on multi-speaker TTS on VCTK with 95% confidence intervals. The FastSpeech2 model is equipped with X-vectors (Snyder et al., 2018).

4.5. Prompt-based Multi-speaker TTS

We also conduct the human evaluation for multi-speaker TTS systems with seen speaker (30 test cases, 15 human

annotations for each test case) and unseen speaker (20 test cases, 15 human annotations for each test case) testing cases. The quality of the generations and the speaker similarity between the generation and the reference are evaluated, and the results are shown in Tab. 6 and Tab. 7. From this table, we can see that the style embedding GST (Wang et al., 2018) improves the similarity scores but harms the quality scores, while our A³T model is the most favorable system in both the speaker similarity and the speech quality. Strikingly, we observe that the average score of the Unseen cases is higher than the Seen, which is counterintuitive. However, when looking into the MOS of the Groundtruth, the gap is still there and we believe this is due to the difference between these two test case sets.

5. Discussion

A³T is a pretraining method on parallel data. A³T is a BERT-style pretraining method, which takes both phonemes and partially-masked spectrograms as inputs (Fig. 8(a)). Although A³T can be first trained with speech-only data (Appendix B), but a second stage of training with parallel data is necessary for speech synthesis. A³T trains a non-autoregressive encoder to reconstruct masked acoustic signals, and uses the identical framework for decoding. It is therefore akin to cross-lingual BERT like XLM (Lample & Conneau, 2019) which also trains on parallel data.

Finetuning. In this paper, our major finding is that A³T can be directly used without finetuning (Tab 2- 7), like GPT-3, for downstream tasks such as speech editing (Tab. 5) and unseen-speaker TTS (Tab. 7). We also find A³T can be pretrained with more data and be finetuned, like BERT, to improve downstream tasks. Finetuning results for multi-speaker TTS are reported in Appendix B.

A³T is not a TTS model. The input of A³T must be both the text and the speech context, while the traditional TTS models’ input is only the text. We show a synthesized speech example in our demo, whose input is the text and a piece of silent audio. We find the generated speech sounds like multiple speakers are speaking the text simultaneously. This observation shows that A³T generates speech based on the given context and follows its properties. On the other hand, A³T can become a TTS model after finetuned with TTS task and data, which is introduced in the Appendix B.

6. Conclusions

In this paper, we propose Alignment-Aware Acoustic-Text Pretraining (A³T) which can reconstruct masked acoustic signals with high quality. We show that our proposed A³T model has the ability to do speech editing and outperforms the current SOTA models, and also improves unseen-speaker speech synthesis with our proposed prompt-based decoding.

References

- Agarap, A. F. Deep learning using rectified linear units (relu). *arXiv preprint arXiv:1803.08375*, 2018.
- Baevski, A., Zhou, H., Mohamed, A., and Auli, M. wav2vec 2.0: A framework for self-supervised learning of speech representations. *NeurIPS 2020*, 2020.
- Bai, H., Shi, P., Lin, J., Xie, Y., Tan, L., Xiong, K., Gao, W., and Li, M. Segatron: Segment-aware transformer for language modeling and understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 12526–12534, 2021.
- Bai, H., Wang, T., Sordoni, A., and Shi, P. Better language model with hypernym class prediction. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1352–1362, May 2022.
- Bapna, A., Chung, Y.-a., Wu, N., Gulati, A., Jia, Y., Clark, J. H., Johnson, M., Riesa, J., Conneau, A., and Zhang, Y. Slam: A unified encoder for speech and language modeling via speech-text joint pre-training. *arXiv preprint arXiv:2110.10329*, 2021.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901, 2020.
- Chen, J., Ma, M., Zheng, R., and Huang, L. Mam: Masked acoustic modeling for end-to-end speech-to-text translation. *arXiv preprint arXiv:2010.11445*, 2020.
- Chu, M. and Peng, H. Objective measure for estimating mean opinion score of synthesized speech, April 4 2006. US Patent 7,024,362.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Gulati, A., Qin, J., Chiu, C.-C., Parmar, N., Zhang, Y., Yu, J., Han, W., Wang, S., Zhang, Z., Wu, Y., et al. Conformer: Convolution-augmented transformer for speech recognition. *arXiv preprint arXiv:2005.08100*, 2020.
- Guo, P., Boyer, F., Chang, X., Hayashi, T., Higuchi, Y., Inaguma, H., Kamo, N., Li, C., Garcia-Romero, D., Shi, J., et al. Recent developments on espnet toolkit boosted by conformer. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5874–5878. IEEE, 2021.
- Hsu, W.-N., Bolte, B., Tsai, Y.-H. H., Lakhotia, K., Salakhutdinov, R., and Mohamed, A. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *arXiv preprint arXiv:2106.07447*, 2021.
- Inaguma, H., Kiyono, S., Duh, K., Karita, S., Soplin, N. E. Y., Hayashi, T., and Watanabe, S. Espnet-st: All-in-one speech translation toolkit. *arXiv preprint arXiv:2004.10234*, 2020.
- Ito, K. and Johnson, L. The LJ Speech Dataset. <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- Jia, Y., Zhang, Y., Weiss, R., Wang, Q., Shen, J., Ren, F., Nguyen, P., Pang, R., Lopez Moreno, I., Wu, Y., et al. Transfer learning from speaker verification to multispeaker text-to-speech synthesis. *Advances in neural information processing systems*, 31, 2018.
- Jin, Z., Mysore, G. J., Diverdi, S., Lu, J., and Finkelstein, A. Voco: Text-based insertion and replacement in audio narration. *ACM Transactions on Graphics (TOG)*, 36(4): 1–13, 2017.
- Kahn, J., Rivière, M., Zheng, W., Kharitonov, E., Xu, Q., Mazaré, P. E., Karadayi, J., Liptchinsky, V., Collobert, R., Fuegen, C., Likhomanenko, T., Synnaeve, G., Joulin, A., Mohamed, A., and Dupoux, E. Libri-light: A benchmark for asr with limited or no supervision. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7669–7673, 2020. <https://github.com/facebookresearch/libri-light>.
- Klimkov, V., Moinet, A., Nadolski, A., and Drugman, T. Parameter generation algorithms for text-to-speech synthesis with recurrent neural networks. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pp. 626–631. IEEE, 2018.
- Lample, G. and Conneau, A. Cross-lingual language model pretraining. *arXiv preprint arXiv:1901.07291*, 2019.
- Liu, A. T., Li, S.-W., and Lee, H.-y. Tera: Self-supervised learning of transformer encoder representation for speech. *arXiv preprint arXiv:2007.06028*, 2020.
- Morrison, M., Rencker, L., Jin, Z., Bryan, N. J., Caceres, J.-P., and Pardo, B. Context-aware prosody correction for text-based speech editing. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7038–7042. IEEE, 2021.
- Moulines, E. and Charpentier, F. Pitch-synchronous waveform processing techniques for text-to-speech synthesis using diphones. *Speech communication*, 9(5-6):453–467, 1990.


- Oord, A., Li, Y., Babuschkin, I., Simonyan, K., Vinyals, O., Kavukcuoglu, K., Driessche, G., Lockhart, E., Cobo, L., Stimberg, F., et al. Parallel WaveNet: Fast high-fidelity speech synthesis. In *International Conference on Machine Learning*, pp. 3915–3923, 2018.
- Oord, A. v. d., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., and Kavukcuoglu, K. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016.
- Panayotov, V., Chen, G., Povey, D., and Khudanpur, S. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5206–5210. IEEE, 2015.
- Peng, K., Ping, W., Song, Z., and Zhao, K. Parallel neural text-to-speech. *arXiv preprint arXiv:1905.08459*, 2019.
- Prenger, R., Valle, R., and Catanzaro, B. Waveglow: A flow-based generative network for speech synthesis. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3617–3621. IEEE, 2019.
- Ren, Y., Ruan, Y., Tan, X., Qin, T., Zhao, S., Zhao, Z., and Liu, T.-Y. FastSpeech: Fast, robust and controllable text to speech. *arXiv preprint arXiv:1905.09263*, 2019.
- Ren, Y., Hu, C., Tan, X., Qin, T., Zhao, S., Zhao, Z., and Liu, T.-Y. FastSpeech 2: Fast and high-quality end-to-end text to speech. *arXiv preprint arXiv:2006.04558*, 2020.
- Shen, J., Pang, R., Weiss, R., Schuster, M., Jaitly, N., Yang, Z., Chen, Z., Zhang, Y., Wang, Y., Skerrv-Ryan, R., Saurous, R., Agiomvrgiannakis, Y., and Wu, Y. Natural TTS synthesis by conditioning WaveNet on MEL spectrogram predictions. In *Interspeech*, 2018a.
- Shen, J., Pang, R., Weiss, R. J., Schuster, M., Jaitly, N., Yang, Z., Chen, Z., Zhang, Y., Wang, Y., Skerrv-Ryan, R., et al. Natural tts synthesis by conditioning wavenet on mel spectrogram predictions. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4779–4783. IEEE, 2018b.
- Snyder, D., Garcia-Romero, D., Sell, G., Povey, D., and Khudanpur, S. X-vectors: Robust dnn embeddings for speaker recognition. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5329–5333. IEEE, 2018.
- Tan, D., Deng, L., Yeung, Y. T., Jiang, X., Chen, X., and Lee, T. Editspeech: A text based speech editing system using partial inference and bidirectional fusion. *arXiv preprint arXiv:2107.01554*, 2021.
- Wang, Y., Stanton, D., Zhang, Y., Ryan, R.-S., Battenberg, E., Shor, J., Xiao, Y., Jia, Y., Ren, F., and Saurous, R. A. Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis. In *International Conference on Machine Learning*, pp. 5180–5189. PMLR, 2018.
- Yamagishi, J., Veaux, C., and MacDonald, K. CSTR VCTK Corpus: English multi-speaker corpus for CSTR voice cloning toolkit (version 0.92), 2019.
- Yamamoto, R., Song, E., and Kim, J.-M. Parallel wavegan: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6199–6203. IEEE, 2020.
- Young, S., Evermann, G., Gales, M., Hain, T., Kershaw, D., Liu, X., Moore, G., Odell, J., Ollason, D., Povey, D., et al. The htk book. *Cambridge university engineering department*, 3(175):12, 2002.
- Yuan, J. and Liberman, M. Speaker identification on the scotus corpus. *Journal of the Acoustical Society of America*, 123(5):3878, 2008.
- Zheng, R., Chen, J., Ma, M., and Huang, L. Fused acoustic and text encoding for multimodal bilingual pretraining and speech translation. *Proceedings of the 38th International Conference on Machine Learning*, 2021.

Appendix


In this appendix, we first show the generated spectrogram comparisons between our system and Tan et al. (2021)'s system in Sec. A.

Then, we show our A³T can also be used for TTS finetuning in Sec. B. As experiments in Sec. B are using more training steps or more data, we decide to put these experiments in the appendix instead of the main text for readers interested in the finetuning.


A. Spectrogram Comparison of Speech Editing



(a) Groundtruth spectrogram from VCTK.




(b) Generated modified spectrogram by A³T.




(c) Generated modified spectrogram by Tan et al. (2021).


Figure 9. An speech editing example from VCTK. Original text: For that reason cover **should not** be given. Modified text: For that reason cover **is impossible to** be given.



(a) Groundtruth spectrogram from VCTK.



(b) Reconstructed spectrogram by A³T.



(c) Reconstructed spectrogram by Tan et al. (2021).

Figure 10. Original text: This would give **Scotland** around eight members. Modified text: This would give **China and Japan** around eight members.

B. Pretraining for Multi-speaker TTS

Method	Pretrain Data	MOS \uparrow
Groundtruth		4.07 \pm 0.07
FastSpeech 2	none	3.63 \pm 0.07
A ³ T	LibriTTS	3.72 \pm 0.07
A ³ T	ASR + Speech	3.77 \pm 0.07

Table 8. MOS evaluation for multi-speaker speech synthesis.




Figure 11. Training and validation loss using LibriTTS between TTS models initialized with and without A³T.

We conduct finetuning experiments with a large multi-speaker TTS dataset LibriTTS but split the validation and test set with only the new speakers. We test on 50 test cases with 15 human annotators for each case. In this setting, we find the FastSpeech 2 fails to generate high-quality audio for these new speakers, even equipped with X-Vector (Snyder et al., 2018) to generate speaker embeddings for new speakers. After initializing the FastSpeech 2 model with our LibriTTS pretrained A³T, the generated audio can be improved significantly. Results are shown in Tab. 8. We also plot the validation loss and training loss during the training of TTS models with and without A³T in Fig. 11. We can see that both the training and validation loss is improved with the initialization from the A³T model, which demonstrates the effectiveness of our method. Finally, we also observe the improvement from the external data (LibriSpeech (Panayotov et al., 2015) and LibriLight (Kahn et al., 2020)) pretraining for the A³T model, which achieves 3.77 MOS scores in Tab. 8. It should be noted that when training with speech only data LibriLight, our model is similar to MAM (Chen et al., 2020) and the alignment embedding are discarded.