

FluentSpeech: Stutter-Oriented Automatic Speech Editing with Context-Aware Diffusion Models

Ziyue Jiang*
Zhejiang University
ziyuejiang@zju.edu.cn

Qian Yang*
Zhejiang University
qyang1021@zju.edu.cn

Jialong Zuo
Zhejiang University
jialongzuo@gmail.com

Zhenhui Ye
Zhejiang University
zhenhuiye@zju.edu.cn

Rongjie Huang
Zhejiang University
rongjiehuang@zju.edu.cn

Yi Ren
Bytedance AI Lab
ren.yi@bytedance.com

Zhou Zhao[†]
Zhejiang University
zhaozhou@zju.edu.cn

Abstract

Stutter removal is an essential scenario in the field of speech editing. However, when the speech recording contains stutters, the existing text-based speech editing approaches still suffer from: 1) the over-smoothing problem in the edited speech; 2) lack of robustness due to the noise introduced by stutter; 3) to remove the stutters, users are required to determine the edited region manually. To tackle the challenges in stutter removal, we propose FluentSpeech, a stutter-oriented automatic speech editing model. Specifically, 1) we propose a context-aware diffusion model that iteratively refines the modified mel-spectrogram with the guidance of context features; 2) we introduce a stutter predictor module to inject the stutter information into the hidden sequence; 3) we also propose a stutter-oriented automatic speech editing (SASE) dataset that contains spontaneous speech recordings with time-aligned stutter labels to train the automatic stutter localization model. Experimental results on VCTK and LibriTTS datasets demonstrate that our model achieves state-of-the-art performance on speech editing. Further experiments on our SASE dataset show that FluentSpeech can effectively improve the fluency of stuttering speech in terms of objective and subjective metrics. Code and audio samples can be found at <https://github.com/Zain-Jiang/Speech-Editing-Toolkit>.

1 Introduction

Recently, text-based speech editing (Jin et al., 2017, 2018; Morrison et al., 2021; Tan et al., 2021; Tae et al., 2021; Wang et al., 2022; Bai et al., 2022) has made rapid progress, and stutter removal is

a critical sub-task in speech editing. There are various application scenarios for stutter removal, like short-form videos, movies, podcasts, YouTube videos, and online lectures, since it provides great convenience for media producers.

Previous speech editing systems (Jin et al., 2017, 2018) successfully enable the user to edit the speech recording through operations in the text transcript. Some neural text-to-speech (TTS) based methods (Tan et al., 2021; Tae et al., 2021) achieve smooth transition at the boundaries of the edited region. And most recently, the mask prediction based methods (Wang et al., 2022; Bai et al., 2022) learn better contextual information from the input mel-spectrogram and outperform previous approaches at speech quality and prosody modeling. However, the existing approaches only aim at modifying reading-style speeches, while removing stutters from spontaneous speeches remains a considerable challenge.

When applied to the stutter removal task, previous efforts are still subject to the following limitations: 1) the generated mel-spectrogram is usually blurry and lacks frequency bin-wise details, resulting in unnatural sounds in the boundaries of the modified region; 2) when the speech recording is full of stutters, the edited speech is usually not robust due to the noise introduced by the discrepancy between text and stuttering speech content; 3) the stutter region should be manually determined one by one, which is costly and laborious for media producers.

To tackle these challenges, we propose FluentSpeech, the first generative model to solve the stutter removal task, which automatically detects the stutter regions, removes them, and generates fluent speech with natural details. Specifically,

*Equal contribution.

[†]Corresponding author.

- Non-probabilistic models tend to generate over-smooth mel-spectrograms (Huang et al., 2022; Popov et al., 2021), while probabilistic models (e.g., GAN and diffusion) generate mel-spectrograms with richer frequency details and natural sounds. Based on this observation, we adopt a context-aware diffusion model that utilizes rich contextual information to guide the diffusion and reverse processes, which helps FluentSpeech to generate high-quality and expressive results.
- To improve the robustness against stuttering speeches, we introduce a conditional stutter predictor that localizes the stutter region and injects the stutter information into the frame-level hidden sequence to reduce the discrepancy between text and stuttering speech. Moreover, the predicted stutter region can be utilized as the mask for automatic stutter removal.
- We propose a novel dataset called the stutter-oriented automatic speech editing (SASE) dataset, which contains spontaneous speech recordings with time-aligned stutter labels for automatic stutter removal.

Experiments on the VCTK (Yamagishi et al., 2019) and LibriTTS (Zen et al., 2019) dataset show that FluentSpeech outperforms state-of-the-art models on speech editing towards reading-style speech with fewer model parameters. And in the experiments on our newly collected SASE dataset, FluentSpeech enjoys much robustness against stuttering speech and demonstrates the ability to improve the fluency of stuttering speech significantly. The main contributions of this work can be summarized as follows:

- We analyze the characteristics of different speech editing approaches (e.g., algorithm, architecture, alignment learning approaches, etc.) and propose a context-aware diffusion probabilistic model that achieves state-of-the-art performance on speech editing.
- We propose a stutter predictor module to improve the robustness against the stuttering speech and localize the stutter region. The stutter predictor can also control the stutter representations by removing the stutters from the spontaneous speech to improve its fluency,

which solves the automatic stutter removal task for the first time.

- We contribute a novel SASE dataset which contains 40 hours of spontaneous speech crawled from online lectures or open courses given by 46 speakers. We will publish our model and dataset as the benchmark for the evaluation of future SASE algorithms.

2 Background

In this section, we describe the background of speech editing and the basic knowledge of diffusion model. We also review the existing applications of diffusion model in speech tasks and analyze their advantages and disadvantages.

2.1 Speech Editing

Conventional speech editing methods (Derry, 2012; Whittaker and Amento, 2004) provide users with interfaces for cut, copy, paste, volume adjustment, time-stretching, pitch bending, de-noising, etc. Then text-based speech editing systems (Jin et al., 2017, 2018) allow the editor to perform select, cut, and paste operations in the text transcript of the speech and apply the changes to the waveform accordingly. However, they mainly face two problems. One is that the edited speech often sounds unnatural because the edited region does not match the prosody of the speech context. (e.g., mismatches in intonation, stress, or rhythm) (Jin et al., 2017). Another is that the interfaces do not support the ability to synthesize new words not appearing in the transcript (Morrison et al., 2021). There are a series of studies on these problems. Jin et al. (2017) propose to insert a synthesized audio clip using a combination of the text-to-speech model and voice conversion model (Sun et al., 2016), which leads to unnatural prosody near the boundaries of the edited regions. Tan et al. (2021) use neural TTS model with auto-regressive partial inference to maintain a coherent prosody and speaking style. Most recently, the mask prediction based methods (Wang et al., 2022; Bai et al., 2022) can capture more contextual information from the input mel-spectrogram. Wang et al. (2022) propose to learn the relation between text and audio through cross-attention but suffer from the extremely slow convergence rate. Bai et al. (2022) introduce the alignment embedding into the Conformer-based (Gulati et al., 2020; Guo et al., 2021) backbone to improve the speech quality. However, previous methods only focus on

the modification of reading-style speeches, which is not stutter-oriented.

2.2 Diffusion Model

Basic knowledge of diffusion model Denoising diffusion probabilistic models (DDPMs) have achieved state-of-the-art performances in both image and audio synthesis (Dhariwal and Nichol, 2021; Kong et al., 2020b; Huang et al., 2022). DDPMs (Ho et al., 2020; Dhariwal and Nichol, 2021) are designed to learn a data distribution $p(x)$ by gradually denoising a normally distributed variable through the reverse process of a fixed Markov Chain of length T . Denote x_t as a noisy version of the clean input x_0 . DDPMs choose to parameterize the denoising model θ through directly predicting ϵ with a neural network ϵ_θ . The corresponding objective can be simplified to:

$$\mathcal{L}_\theta^{\text{Grad}} = \left\| \epsilon_\theta \left(\alpha_t x_0 + \sqrt{1 - \alpha_t^2} \epsilon \right) - \epsilon \right\|_2^2, \epsilon \sim \mathcal{N}(0, \mathbf{I}), \quad (1)$$

with t uniformly sampled from $\{1, \dots, T\}$.

Applications of diffusion model in speech tasks

Applications of diffusion model in speech tasks mainly lie in speech synthesis. Diff-TTS (Jeong et al., 2021), Grad-TTS (Popov et al., 2021), and DiffSpeech (Liu et al., 2021) are gradient-based models with score-matching objectives to generate high-quality speeches, which require hundreds of iterations with small β_t to guarantee high sample quality. Most recently, ProDiff (Huang et al., 2022) parameterize the denoising model by directly predicting clean data and avoids significant perceptual quality degradation when reducing reverse iterations. In the field of speech editing, Tae et al. (2021) propose a diffusion model that requires a pre-trained TTS model to synthesize the target audio and eliminate the artifacts of concatenation by a score-based manipulation algorithm, which is not text-based speech editing.

3 FluentSpeech

This section presents our proposed FluentSpeech, a stutter-oriented automatic speech editing model that solves the stutter removal task. We firstly overview the motivation and the architecture of FluentSpeech. Secondly, we describe the detailed designs of alignment modeling, context-aware spectrogram denoiser, and stutter predictor. Finally, we describe the training objectives of FluentSpeech,

following with the illustration of training and inference procedures.

3.1 Model Overview

The overall model architecture of FluentSpeech is shown in Figure 1. FluentSpeech consists of a linguistic encoder and a context-aware spectrogram denoiser. Denote the phoneme sequence of the transcription as $\mathbf{p} = (p_1, \dots, p_{|\mathbf{p}|})$ and the acoustic feature sequence as $\mathbf{x} = (x_1, \dots, x_{|\mathbf{x}|})$. \mathbf{x} can be the spectrogram or mel-spectrogram of the speech audio, and each x_i represents the speech feature of frame i . The Transformer-based (Vaswani et al., 2017) linguistic encoder converts \mathbf{p} into the text hidden sequence e_p . Denote $\hat{\mathbf{x}} = \text{Mask}(\mathbf{x}, \boldsymbol{\lambda})$ as the masked acoustic feature sequence, where $\text{Mask}(\cdot)$ replaces several random spans of x by the probability of λ with the same number of a random initialized masking vector. Then, the context-aware spectrogram denoiser θ aggregates phoneme embedding e_p and other features like acoustic embedding e_x , pitch embedding e_{pitch} as the condition c to guide the reverse process of the diffusion model $f_\theta(x_t | t, c)$.

3.2 Alignment Modeling

Due to the modality gap between text and speech, alignment modeling is essential in text-based speech editing. There are three types of approaches to model the monotonous alignment between text and speech: 1) cross-attention, Wang et al. (2022) propose to learn the alignment information with the cross-attention module in the transformer decoder, which suffers from the slow convergence rate and is usually not robust; 2) alignment embedding, Bai et al. (2022) introduce the alignment embedding from external alignment tools into the self-attention based architecture to guide the alignment modeling; 3) length regulator (Ren et al., 2019; Tan et al., 2021), the length regulator expand text embedding into frame-level embedding according to the phoneme duration predicted by the duration predictor (Ren et al., 2019; Tan et al., 2021), which ensures hard alignments and is more robust than the above two methods. However, the duration predictor in Tan et al. (2021) does not consider the existing context duration. It only predicts the duration of the entire sentence from text representations and applies the duration of the edited words to the masked region, which results in unnatural prosody. Therefore, in FluentSpeech, we train the duration predictor with the mask prediction proce-

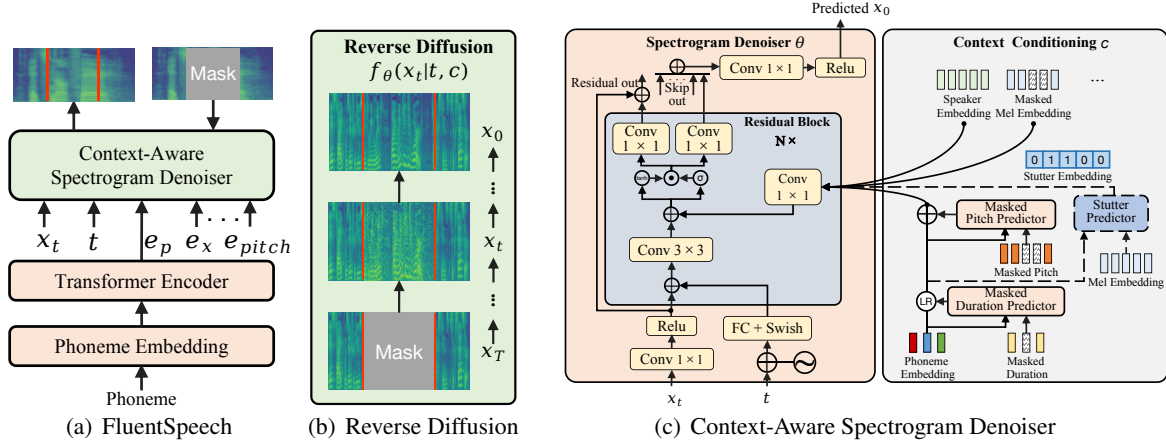


Figure 1: The overall architecture for FluentSpeech. In subfigure (c), the spectrogram denoiser θ takes noisy spectrogram x_t as input and computes $f_\theta(x_t|t, c)$ conditioned on diffusion time index t and context information c . The sinusoidal-like symbol, FC, Swish, and \bullet denote the positional encoding, fully-connected layer, swish activation function (Ramachandran et al., 2017), and element-wise multiple operation. LR denotes the Length Regulator proposed in FastSpeech (Ren et al., 2019). N is the number of residual blocks. The dashed black line denotes that the operation is only executed when the dataset contains spontaneous speeches.

ture to achieve the fluent duration transition at the edited region, which is called the masked duration predictor.

3.3 Context-Aware Spectrogram Denoiser

Context Conditioning As shown in Figure 1(c), in the context conditioning module, we adopt frame-level text embedding e_t , acoustic feature sequence x , masked acoustic feature sequence \hat{x} , speaker embedding e_{spk} , pitch embedding e_{pitch} , and stutter embedding $e_{stutter}$ as the condition for our spectrogram denoiser. The phoneme embedding e_p is first expanded into frame-level text embedding e_t by the length regulator with the duration information from the masked duration predictor. We add e_t to the context condition c . We also extract the speaker embeddings e_{spk} from audio samples using open-source voice encoder¹ and feed them into the context condition c following the common practice (Min et al., 2021; Huang et al., 2022; Tan et al., 2021). Then we adopt a nonlinear feed-forward acoustic encoder to transform the speech feature x and \hat{x} into the acoustic embeddings e_x and $e_{\hat{x}}$ following Bai et al. (2022). The masked acoustic embedding $e_{\hat{x}}$ is also added to the condition to provide more contextual information for mel-spectrogram reconstruction. Moreover, the masked pitch predictor utilizes e_t and the masked pitch embedding \hat{e}_{pitch} to predict the pitch F_0 of each frame in the edited region. We further con-

vert it into the pitch embedding vector and add it to the context condition c . To promote the natural transition at the edited boundaries, we train the duration predictor and pitch predictor with the mask prediction procedure:

$$\mathcal{L}_p = \|p - g_p(e_t, \hat{e}_{pitch})\|_2^2, \quad (2)$$

$$\mathcal{L}_d = \|d - g_d(e_t, \hat{e}_{dur})\|_2^2 \quad (3)$$

where we use d and p to denote the target duration and pitch respectively, and use g_d and g_p to denote the corresponding duration predictor and pitch predictor, which share the same architecture of 1D convolution with ReLU activation and layer normalization. The loss weights are all set to 0.1 and the reconstruction losses are also added to train the linguistic encoder.

Spectrogram Denoiser Following Liu et al. (2021); Huang et al. (2022), we adopt a non-causal WaveNet (Oord et al., 2016) architecture to be our spectrogram denoiser. The decoder comprises a 1×1 convolution layer and N convolution blocks with residual connections to project the input hidden sequence with 256 channels. For any step t , we use the cosine schedule $\beta_t = \cos(0.5\pi t)$. Different from the aforementioned diffusion models that require hundreds of steps with small β_t to estimate the gradient for data density, we choose to parameterize the denoising model by directly predicting the clean data x_0 following recent researches in image generation and TTS literature (Salimans and

¹<https://github.com/resemble-ai/Resemblyzer>

Ho, 2021; Liu et al., 2022; Huang et al., 2022) to significantly accelerate sampling from a complex distribution. Specifically, in the generator-based diffusion models, $p_\theta(x_0|x_t)$ is the implicit distribution imposed by the neural network $f_\theta(x_t, t)$ that outputs x_0 given x_t . And then x_{t-1} is sampled using the posterior distribution $q(x_{t-1}|x_t, x_0)$ given x_t and the predicted x_0 . The training loss is defined as the mean absolute error (MAE) in the data x space:

$$\mathcal{L}_\theta^{MAE} = \left\| \mathbf{x}_\theta \left(\alpha_t \mathbf{x}_0 + \sqrt{1 - \alpha_t^2} \epsilon \right) - \mathbf{x}_0 \right\|, \epsilon \sim \mathcal{N}(0, \mathbf{I}), \quad (4)$$

and efficient training is optimizing a random t term with stochastic gradient descent. Inspired by (Ren et al., 2022), we also adopt structural similarity index (SSIM) loss $\mathcal{L}_\theta^{SSIM}$ in training to capture structural information in mel-spectrogram and improve the perceptual quality:

$$\mathcal{L}_\theta^{SSIM} = 1 - \text{SSIM} \left(\mathbf{x}_\theta \left(\alpha_t \mathbf{x}_0 + \sqrt{1 - \alpha_t^2} \epsilon \right), \hat{\mathbf{x}}_0 \right). \quad (5)$$

The loss weights are all set to 0.5. Since the capability of our spectrogram denoiser is powerful enough, we do not adopt the convolutional Post-Net to refine the predicted spectrogram like previous works (Wang et al., 2022; Bai et al., 2022).

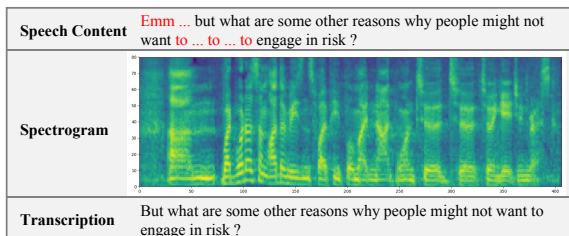


Figure 2: The illustration of the discrepancy between the given transcription and stuttering speech content.

3.4 Stutter Predictor

The stutter predictor is introduced only when the speech corpus contains stuttering recordings. The stutters in the speech content will introduce noise to the training pipeline due to the noise introduced by the information gap between text and stuttering speech content. As shown in Figure 2, the stuttering word “to” in the speech content makes the speech editing model learn unintentional sounds in the pronunciation of the word “to”. Therefore, we introduce the stutter embedding into the text hidden sequence to disentangle the stutter-related

gradients from the speech content, which significantly improves the pronunciation robustness of our FluentSpeech.

Let $s = (s_1, \dots, s_s)$ be a time-aligned stutter label that defines the stutter regions in the corresponding spontaneous speech, where $s_i \in \{0, 1\}$ (0 for normal and 1 for stutter) for each frame (See Appendix C for further details about the stutter label in our SASE dataset). In training, we take the ground-truth value of the stutter label as input into the hidden sequence to predict the target speech. At the same time, we use the ground-truth labels as targets to train the stutter predictor, which is used in inference to localize the stutter region in target speech.

The stutter predictor consists of 1) a 4-layer 1D conditional convolutional network with ReLU activation, each followed by the layer normalization and the dropout layer; 2) an extra linear layer and a softmax layer to predict the probability of stutter tag. As shown in Figure 1(c), we propose a text-guided stutter predictor module, which takes frame-level text embedding e_t and mel-spectrogram embedding e_x as input and seeks to locate the text-irrelevant stutter regions. The main objective function for stutter prediction is the binary cross-entropy loss \mathcal{L}_{BCE} . The Focal loss (Lin et al., 2017) \mathcal{L}_{Focal} is also introduced since the misclassification of fluent regions is tolerable and we want the stuttering regions to be accurately classified. The α_0, α_1 is set to $5e^{-3}, 1$ and γ is set to 3.

3.5 Training and Inference Procedures

Training The final training loss terms consist of the following parts: 1) sample reconstruction loss \mathcal{L}_θ^{MAE} ; 2) structural similarity index (SSIM) loss $\mathcal{L}_\theta^{SSIM}$; 3) reconstruction loss for pitch and duration predictor $\mathcal{L}_p, \mathcal{L}_d$; 4) classification loss for stutter predictor $\mathcal{L}_{BCE}, \mathcal{L}_{Focal}$. In the training stage, we randomly select 80% phonemes spans and mask their corresponding frames since 80% masking rate shows good performances on both seen and unseen cases. Then we add the stutter embedding to the context condition. The objective functions only take the masked region into consideration.

Inference for reading-style speech editing

Given a speech spectrogram x , its original phonemes \tilde{p} and the target phonemes p . Denote the spectrogram region that needs to be modified as μ . When the speech recording is reading-style, we do not utilize the stutter predictor. We first use an ex-

ternal alignment tool² to extract the spectrogram-to-phoneme alignments. \hat{x} is the spectrogram masked according to the region μ . FluentSpeech takes p , \hat{x} , x , e_{spk} , \hat{e}_{dur} , and \hat{e}_{pitch} as inputs and generates the spectrogram of the masked region μ . Finally, we use a pre-trained vocoder to transform this spectrogram into the waveform.

Inference for stutter removal When the speech recording is spontaneous, the stutter predictor first predicts the stutter region μ' . Since the stutter region μ' also influences the prosody (e.g., duration and pitch) of the neighboring words, we find all of the phoneme spans that overlap with or are adjacent³ to μ' and denote them as $\hat{\mu}$. Then the spectrogram region that needs to be modified can be defined as $\mu = \mu' \cup \hat{\mu}$. To make the spontaneous speech fluent, the stutter embedding is not added to the hidden sequence. Following the masked spectrogram reconstruction process in the inference for reading-style speech editing, FluentSpeech is able to perform automatic stutter removal.

4 Experiments

4.1 Datasets

Reading-Style We evaluate FluentSpeech on two reading-style datasets, including: 1) VCTK (Yamagishi et al., 2019), an English speech corpus uttered by 110 English speakers with various accents; 2) LibriTTS (Zen et al., 2019), a large-scale multi-speaker English corpus of approximately 585 hours of speech. We evaluate the text-based speech editing performance of FluentSpeech and various baselines on these datasets.

Spontaneous We also evaluate FluentSpeech on the stutter-oriented automatic speech editing (SASE) dataset collected and annotated by us (See Appendix C for further details). The SASE dataset consists of approximately 40 hours of spontaneous speech recordings from 46 speakers with various accents. All the audio files are collected from online lectures and courses with accurate official transcripts. Each recording is sampled at 22050 Hz with 16-bit quantization. We evaluate the SASE performance of FluentSpeech and various baselines on this dataset.

²<https://github.com/MontrealCorpusTools/Montreal-Forced-Aligner>

³The adjacent region is the top-1 adjacent word boundary given by the alignment tool.

For each of the three datasets, we randomly sample 400 samples for testing. We randomly choose 50 samples in the test set for subjective evaluations and use all testing samples for objective evaluations. The ground truth mel-spectrograms are generated from the raw waveform with the frame size 1024 and the hop size 256.

4.2 Experimental Setup

Model Configuration FluentSpeech consists of a linguistic encoder, an acoustic encoder, a masked variance adaptor, a spectrogram denoiser, and a stutter predictor. The linguistic and acoustic encoders consist of multiple feed-forward Transformer blocks (Ren et al., 2019) with relative position encoding (Shaw et al., 2018) following GlowTTS (Kim et al., 2020). The hidden channel is set to 256. In the spectrogram denoiser, we set $N = 20$ to stack 20 layers of convolution with the kernel size 3, and we set the dilated factor to 1 (without dilation) at each layer following (Huang et al., 2022). The number of diffusion steps T is set to 8. The stutter predictor is based on the non-causal WaveNet (Oord et al., 2016) architecture. We have attached more detailed information on the model configuration in Appendix A.1.

Training and Evaluation We train the FluentSpeech with $T = 8$ diffusion steps. The FluentSpeech model has been trained for 300,000 steps using 1 NVIDIA 3080 GPU with a batch size of 30 sentences. The adam optimizer is used with $\beta_1 = 0.9$, $\beta_2 = 0.98$, $\epsilon = 10^{-9}$. We utilize HiFi-GAN (Kong et al., 2020a) (V1) as the vocoder to synthesize waveform from the generated mel-spectrogram in all our experiments. To measure the perceptual quality, we conduct human evaluations with MOS (mean opinion score), CMOS (comparative mean opinion score), and average preference score on the testing set via Amazon Mechanical Turk (See Appendix A.3 for more details). We keep the text content and text modifications consistent among different models to exclude other interference factors, only examining the audio quality. We further measure the objective evaluation metrics, such as MCD (Kubichek, 1993), STOI (Taal et al., 2010), and PESQ (Rix et al., 2001). More information on evaluation has been attached in Appendix A.4.

Method	VCTK			LibriTTS			#Params.
	MCD (\downarrow)	STOI (\uparrow)	PESQ (\uparrow)	MCD (\downarrow)	STOI (\uparrow)	PESQ (\uparrow)	
EditSpeech	6.92	0.69	1.43	5.33	0.68	1.35	48.1M
CampNet	7.83	0.54	1.38	6.51	0.40	1.28	14.7M
A ³ T	6.25	0.41	1.18	5.69	0.70	1.39	67.7M
FluentSpeech	5.86	0.81	1.91	4.74	0.78	1.82	23.9M

Table 1: The objective audio quality comparisons. We only measure the MCD, STOI, and PESQ of the masked region. MCD and PESQ indicate speech quality, and STOI reflects speech intelligibility

Method	Seen	Unseen
EditSpeech	4.00 \pm 0.10	3.89 \pm 0.09
CampNet	3.59 \pm 0.11	3.04 \pm 0.18
A ³ T	4.09 \pm 0.10	3.90 \pm 0.10
FluentSpeech	4.27 \pm 0.11	4.18 \pm 0.09

Table 2: The MOS evaluation (\uparrow) for speech quality on speech editing task on the VCTK dataset with 95% confidence intervals.

Method	Seen	Unseen
EditSpeech	4.26 \pm 0.10	3.90 \pm 0.13
CampNet	3.93 \pm 0.12	3.58 \pm 0.20
A ³ T	4.27 \pm 0.09	3.53 \pm 0.14
FluentSpeech	4.42 \pm 0.06	4.21 \pm 0.11

Table 3: The MOS evaluation (\uparrow) for speaker similarity on speech editing task on the VCTK dataset with 95% confidence intervals.

4.3 Results of Reading-Style Speech Editing

We compare the quality of generated audio samples of our FluentSpeech with other baseline systems, including 1) EditSpeech (Tan et al., 2021); 2) CampNet (Wang et al., 2022); 3) A³T (Bai et al., 2022) (detailed descriptions can be found in Appendix A.2). For objective evaluation, we conduct the spectrogram reconstruction experiment to evaluate these systems. As shown in Table 1, FluentSpeech demonstrates superior performance in MCD, PESQ, and STOI metrics.

For subjective evaluation, we manually define modification operations (i.e., insertion, replacement, and deletion) of 50 audio samples. We then conduct the experiments on the VCTK dataset. For each audio sample, we ask at least 10 English speakers to evaluate the generated audios’ speech quality and speaker similarity. The results are presented Table 2 and Table 3. For the seen case, each speaker’s examples would be split into train and test sets. And for the unseen case, the test set contains 10 speakers’ examples, and the other 99 speakers’ examples are used for training following (Bai et al., 2022). It can be seen that FluentSpeech achieves the highest perceptual quality and speaker similarity on both seen and unseen settings compared to all baselines, which demonstrates the effectiveness of our proposed context-aware spectrogram denoiser.




Figure 3: Average preference score (%) evaluation on naturalness and fluency on the SASE dataset, where “Neural” stands for “no preference”.

4.4 Results of Stutter-Oriented Automatic Speech Editing

We evaluate the accuracy of FluentSpeech on the stutter localization task, and the results are shown in Table 4. It can be seen that our FluentSpeech achieves 80.5% accuracy and 94.4% precision on the stutter localization task. We then compare the naturalness and fluency of generated audio samples of our FluentSpeech with the original spontaneous recordings. We conduct a subjective average preference score evaluation, where 50 sentences are randomly selected from the test set of our SASE dataset. The listeners are asked to judge which utterance in each pair has better naturalness (or fluency) or no preference in the edited area. As shown in Figure 3, FluentSpeech achieves similar naturalness compared to the original audio. Moreover, the fluency of the speeches generated by our FluentSpeech is significantly improved, which further shows the effectiveness of our stutter-oriented

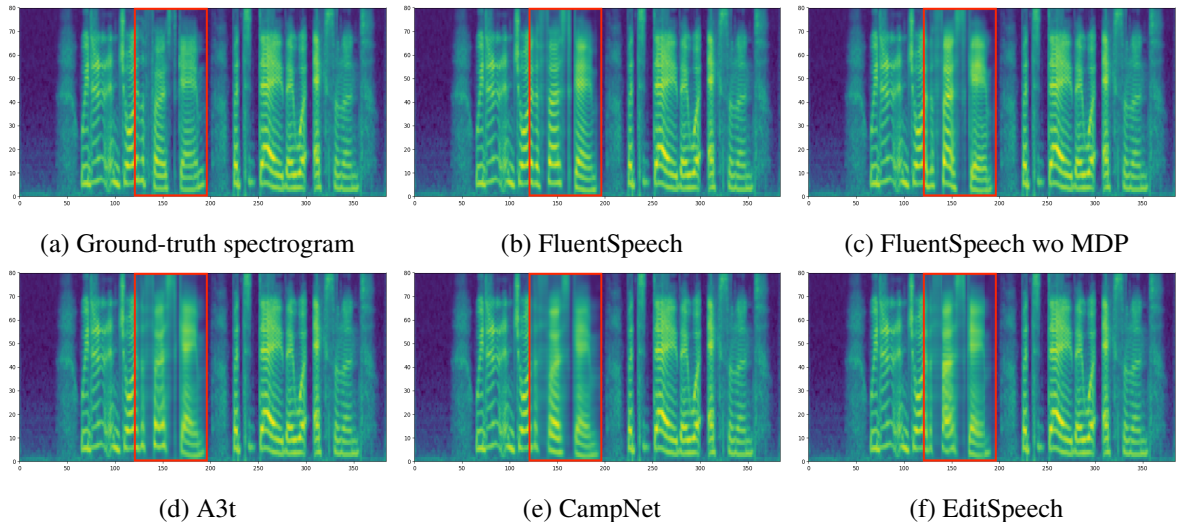


Figure 4: Visualizations of the ground-truth and generated mel-spectrograms by different speech editing models. Original text is “We didnt enjoy **the first game**, but today they were excellent”. In (b,c,d,e,f) subfigures, the portion with red box is “**the first game**” which is masked and reconstructed. MDP denotes the masked duration predictor.

Method	Accuracy (%)	Precision (%)
FluentSpeech	80.5%	94.4%

Table 4: The stutter localization evaluation (\uparrow) on the SASE dataset. Accuracy (%) denotes the overall accuracy; Precision (%) indicates the proportion of the correctly classified stutter regions.

automatic speech editing strategy.

4.5 Visualizations

As illustrated in Figure 4, we visualize the mel-spectrograms generated by FluentSpeech and baseline systems. We can see that FluentSpeech can generate mel-spectrograms with richer frequency details compared with other baselines, resulting in natural and expressive sounds. Moreover, when we substitute the masked duration predictor with the duration predictor utilized in Tan et al. (2021); Wang et al. (2022); Bai et al. (2022), an unnatural transition has occurred in the left boundary of the edited region of FluentSpeech, which demonstrates the effectiveness of our proposed masked duration predictor.

4.6 Ablation Studies

We conduct ablation studies to demonstrate the effectiveness of several designs in FluentSpeech, including the stutter embedding and the masked predictors. We perform CMOS and MCD evaluations for these ablation studies. The results are shown in Table 5. We can see that CMOS drops

Method	C-MOS	MCD (\downarrow)
FluentSpeech	0.00	4.54
- Stutter Embedding	-0.52	4.63
- MDP - MPP + DP + PP	-0.35	5.75
- MDP - MPP + DP	-0.24	5.15

Table 5: Audio quality comparisons on the SASE dataset for ablation study. MDP denotes the masked duration predictor; MPP denotes the masked pitch predictor; DP denotes the duration predictor used in Tan et al. (2021); Bai et al. (2022) and PP denotes the pitch predictor proposed in (Ren et al., 2020).

rapidly when we remove the stutter embedding, indicating that the noise introduced by the text-speech pair’s discrepancy greatly reduces the naturalness of the generated audio. Thus, the stutter embedding successfully improves the robustness of our FluentSpeech; Moreover, when we remove the MDP, MPP and use the DP following recent speech editing algorithms (Tan et al., 2021; Wang et al., 2022; Bai et al., 2022), the speech quality also drops significantly, demonstrating the effectiveness of our proposed masked predictors. It is worth mentioning that the pitch predictor without masked training also results in a performance drop in terms of voice quality.

5 Conclusion

In this work, we proposed FluentSpeech, a stutter-oriented automatic speech editing model for stutter removal. FluentSpeech adopts a context-aware spectrogram denoiser to generate high-quality and

expressive speeches with rich frequency details. To improve the robustness against stuttering speeches and perform automatic stutter removal, we propose a conditional stutter predictor that localizes the stutter region and injects the stutter embedding into the text hidden sequence to reduce the discrepancy between text and stuttering speech recording. We also contribute a novel stutter-oriented automatic speech editing dataset named SASE, which contains spontaneous speech recordings with time-aligned stutter labels. Experimental results demonstrate that FluentSpeech achieves state-of-the-art performance on speech editing for reading-style speeches. Moreover, FluentSpeech is robust against stuttering speech and demonstrates the ability to improve the fluency of stuttering speech significantly. To the best of our knowledge, FluentSpeech is the first stutter-oriented automatic speech editing model that solves the automatic stutter removal task. Our extensive ablation studies demonstrated that each design in FluentSpeech is effective. We hope that our work will serve as a basis for future stutter-oriented speech editing studies.

6 Limitations

We list the limitations of our work as follows. Firstly, the model architecture we use to localize the stuttering speech is simple. Future works could explore a more effective model to perform automatic stutter removal with the help of our SASE dataset. Secondly, we only test the English datasets. And other languages except for English and multi-language stutter-oriented speech editing remain for future works. Finally, after being pre-trained on our SASE dataset, the stutter embedding in FluentSpeech could also be used to inject stutters into the reading-style speech to change its speaking style, and we leave it for future works.

7 Ethics Statement

FluentSpeech improves the naturalness of edited speech and promotes the automatic stutter removal of stuttered speech, which may cause unemployment for people with related occupations. Besides, the free manipulation of speeches may bring potential social damage. Further efforts in automatic speaker verification should be made to lower the aforementioned risks.

8 Acknowledgments

This work was supported in part by the National Key R&D Program of China under Grant No.2022ZD0162000, National Natural Science Foundation of China under Grant No. 62222211, Grant No.61836002 and Grant No.62072397.

References

- He Bai, Renjie Zheng, Junkun Chen, Mingbo Ma, Xintong Li, and Liang Huang. 2022. A3t: Alignment-aware acoustic and text pretraining for speech synthesis and editing. In *International Conference on Machine Learning*, pages 1399–1411. PMLR.
- Roger Derry. 2012. *PC audio editing with Adobe Audition 2.0: Broadcast, desktop and CD audio production*. Routledge.
- Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794.
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, et al. 2020. Conformer: Convolution-augmented transformer for speech recognition. *arXiv preprint arXiv:2005.08100*.
- Pengcheng Guo, Florian Boyer, Xuankai Chang, Tomoki Hayashi, Yosuke Higuchi, Hirofumi Inaguma, Naoyuki Kamo, Chenda Li, Daniel Garcia-Romero, Jiatong Shi, et al. 2021. Recent developments on espnet toolkit boosted by conformer. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5874–5878. IEEE.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851.
- Yi Hu and Philipos C Loizou. 2007. Evaluation of objective quality measures for speech enhancement. *IEEE Transactions on audio, speech, and language processing*, 16(1):229–238.
- Rongjie Huang, Zhou Zhao, Huadai Liu, Jinglin Liu, Chenye Cui, and Yi Ren. 2022. Prodiff: Progressive fast diffusion model for high-quality text-to-speech. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 2595–2605.
- Myeonghun Jeong, Hyeongju Kim, Sung Jun Cheon, Byoung Jin Choi, and Nam Soo Kim. 2021. Diff-ts: A denoising diffusion model for text-to-speech. *arXiv preprint arXiv:2104.01409*.

- Zeyu Jin, Gautham J Mysore, Stephen Diverdi, Jingwan Lu, and Adam Finkelstein. 2017. Voco: Text-based insertion and replacement in audio narration. *ACM Transactions on Graphics (TOG)*, 36(4):1–13.
- Zeyu Jin et al. 2018. Speech synthesis for text-based editing of audio narration.
- Jaehyeon Kim, Sungwon Kim, Jungil Kong, and Sun-groh Yoon. 2020. Glow-tts: A generative flow for text-to-speech via monotonic alignment search. *Advances in Neural Information Processing Systems*, 33:8067–8077.
- Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. 2020a. Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis. *Advances in Neural Information Processing Systems*, 33:17022–17033.
- Zhifeng Kong, Wei Ping, Jiayi Huang, Kexin Zhao, and Bryan Catanzaro. 2020b. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761*.
- Robert Kubichek. 1993. Mel-cepstral distance measure for objective speech quality assessment. In *Proceedings of IEEE pacific rim conference on communications computers and signal processing*, volume 1, pages 125–128. IEEE.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988.
- Jinglin Liu, Chengxi Li, Yi Ren, Feiyang Chen, Peng Liu, and Zhou Zhao. 2021. DiffSinger: Diffusion acoustic model for singing voice synthesis.
- Songxiang Liu, Dan Su, and Dong Yu. 2022. Diffgan-tts: High-fidelity and efficient text-to-speech with denoising diffusion gans. *arXiv preprint arXiv:2201.11972*.
- Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger. 2017. Montreal forced aligner: Trainable text-speech alignment using kaldi. In *Interspeech*, volume 2017, pages 498–502.
- Dongchan Min, Dong Bok Lee, Eunho Yang, and Sung Ju Hwang. 2021. Meta-stylespeech: Multi-speaker adaptive text-to-speech generation. In *International Conference on Machine Learning*, pages 7748–7759. PMLR.
- Max Morrison, Lucas Rencker, Zeyu Jin, Nicholas J Bryan, Juan-Pablo Caceres, and Bryan Pardo. 2021. Context-aware prosody correction for text-based speech editing. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7038–7042. IEEE.
- Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. 2016. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*.
- Vadim Popov, Ivan Vovk, Vladimir Gogoryan, Tasnima Sadekova, and Mikhail Kudinov. 2021. Grad-tts: A diffusion probabilistic model for text-to-speech. In *International Conference on Machine Learning*, pages 8599–8608. PMLR.
- Prajit Ramachandran, Barret Zoph, and Quoc V Le. 2017. Searching for activation functions. *arXiv preprint arXiv:1710.05941*.
- Yi Ren, Chenxu Hu, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. 2020. Fastspeech 2: Fast and high-quality end-to-end text to speech. *arXiv preprint arXiv:2006.04558*.
- Yi Ren, Yangjun Ruan, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. 2019. Fastspeech: Fast, robust and controllable text to speech. *Advances in Neural Information Processing Systems*, 32.
- Yi Ren, Xu Tan, Tao Qin, Zhou Zhao, and Tie-Yan Liu. 2022. Revisiting over-smoothness in text to speech. *arXiv preprint arXiv:2202.13066*.
- Antony W Rix, John G Beerends, Michael P Hollier, and Andries P Hekstra. 2001. 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 (Cat. No. 01CH37221)*, volume 2, pages 749–752. IEEE.
- Tim Salimans and Jonathan Ho. 2021. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. Self-attention with relative position representations. *arXiv preprint arXiv:1803.02155*.
- Lifa Sun, Kun Li, Hao Wang, Shiyin Kang, and Helen Meng. 2016. Phonetic posteriorgrams for many-to-one voice conversion without parallel data training. In *2016 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE.
- Cees H Taal, Richard C Hendriks, Richard Heusdens, and Jesper Jensen. 2010. A short-time objective intelligibility measure for time-frequency weighted noisy speech. In *2010 IEEE international conference on acoustics, speech and signal processing*, pages 4214–4217. IEEE.
- Cees H Taal, Richard C Hendriks, Richard Heusdens, and Jesper Jensen. 2011. An algorithm for intelligibility prediction of time–frequency weighted noisy speech. *IEEE Transactions on Audio, Speech, and Language Processing*, 19(7):2125–2136.

- Jaesung Tae, Hyeongju Kim, and Taesu Kim. 2021. Editts: Score-based editing for controllable text-to-speech. *arXiv preprint arXiv:2110.02584*.
- Daxin Tan, Liqun Deng, Yu Ting Yeung, Xin Jiang, Xiao Chen, and Tan Lee. 2021. Editspeech: A text based speech editing system using partial inference and bidirectional fusion. In *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 626–633. IEEE.
- Tomoki Toda, Alan W Black, and Keiichi Tokuda. 2007. Voice conversion based on maximum-likelihood estimation of spectral parameter trajectory. *IEEE Transactions on Audio, Speech, and Language Processing*, 15(8):2222–2235.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Tao Wang, Jiangyan Yi, Liqun Deng, Ruibo Fu, Jianhua Tao, and Zhengqi Wen. 2022. Context-aware mask prediction network for end-to-end text-based speech editing. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6082–6086. IEEE.
- Steve Whittaker and Brian Amento. 2004. Semantic speech editing. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 527–534.
- Junichi Yamagishi, Christophe Veaux, and Kirsten Macdonald. 2019. *CSTR VCTK Corpus: English multi-speaker corpus for CSTR voice cloning toolkit (version 0.92)*.
- H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu. 2019. *LibriTTS: A corpus derived from librispeech for text-to-speech*. In *Proc. Interspeech*.
- Hui Zhang, Xueliang Zhang, and Guanglai Gao. 2018. Training supervised speech separation system to improve stoi and pesq directly. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5374–5378. IEEE.

A Detailed Experimental Settings

A.1 Model Configurations

We list the model hyper-parameters of FluentSpeech in Table 6.

A.2 Details of Baseline Systems

EditSpeech (Tan et al., 2021) is a speech-editing system that 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. For decoding, the left-to-right TTS model and the right-to-left TTS model generate the modified region simultaneously. Finally, the two synthesized speeches are fused for the final output. CampNet (Wang et al., 2022) propose a context-aware mask prediction network (CampNet) to simulate the process of text-based speech editing. Three text-based speech editing operations based on CampNet are designed: deletion, replacement, and insertion. And a word-level autoregressive generation method is proposed to improve the editing length. A³T (Bai et al., 2022) propose the alignment-aware acoustic-text pre-training, a BERT-style pre-training model, which takes both phonemes and partially-masked spectrograms as inputs. The alignment embedding from external alignment tools is introduced into the Conformer-based (Gulati et al., 2020; Guo et al., 2021) backbone to improve the speech quality.

A.3 Details in Subjective Evaluation

We perform the subjective evaluation on Amazon Mechanical Turk (MTurk). For speech editing evaluations, we randomly select 50 samples from the test set and manually define modification operations (i.e., insertion, replacement, and deletion) for these audio samples. We use FluentSpeech and the baseline speech editing systems to edit the audio samples. Each generated audio has been listened to by at least 10 native listeners. We paid \$8 to participants hourly and spent about \$400 on participant compensation. We tell the participants that the data will be used in scientific research.

- For audio quality evaluations (MOS), each tester is asked to evaluate the subjective naturalness of a sentence on a 1-5 Likert scale, and we tell listeners to “*assess the quality of the audio based on how close it is to natural speech*”.
- For speaker similarity evaluations (MOS), listeners are asked to compare pairs of audio generated by systems A and ground-truth B and indicate the speaker similarity of the two audio and choose the scores on a 1-5 similar scale. We tell listeners to answer “*How similar is this recording to the reference audio? Please focus only on the similarity of the speaker to the reference, and ignore the differences of content, grammar, or audio quality*”. The screenshots of instructions for speech edit-

Hyperparameter		FluentSpeech	Number of parameters
Text Encoder	Phoneme Embedding	192	3.7M
	Encoder Layers	4	
	Encoder Hidden	192	
	Encoder Conv1d Kernel	5	
	Encoder Conv1D Filter Size	384	
Context Condition	Predictor Conv1D Kernel	3	5.8M
	Predictor Conv1D Filter Size	256	
	Predictor Dropout	0.4	
Spectrogram Denoiser	Diffusion Embedding	256	14.4M
	Residual Layers	20	
	Residual Channels	256	
	WaveNet Conv1d Kernel	3	
	WaveNet Conv1d Filter	512	
Total Number of Parameters		23.9M	

Table 6: Hyperparameters of FluentSpeech models.

ing tests are shown in Figure 5(a) and Figure 5(b).

- For stutter removal evaluations, we perform average preference score tests for speech quality and fluency. For the speech quality AB test, each listener is asked to select their preferred audio according to audio quality. We tell listeners to answer “Which of the audio has better quality? Please focus on the audio quality and ignore other factors”. For the speech fluency AB test, each listener is asked to select the audio they prefer according to audio fluency, and we tell listeners to answer “Which of the audio sounds more fluent? Please focus on speech fluency and ignore other factors. The stutter in the audio typically sounds like “emm”, “uhhh”, “hmmm”, or words repetition”. The screenshots of instructions for stutter removal evaluations are shown in Figure 5(c) and Figure 5(d).

A.4 Details in Objective Evaluation

The effectiveness of our FluentSpeech is measured by MCD (Toda et al., 2007), STOI (Taal et al., 2011), PESQ (Hu and Loizou, 2007) metrics. MCD measures the Euclidean distance between two mel cepstral sequences, which describes the global spectral characteristics of audio signals. PESQ indicates speech quality, and STOI reflects speech intelligibility (Zhang et al., 2018). The lower MCD and higher PESQ, STOI represent better performance in the generated speech. Denote $m^t = [m_1^t, \dots, m_L^t]$ and $m^c = [m_1^c, \dots, m_L^c]$ as two mel cepstral sequences. The traditional MCD measure is given

Method	Duration Error (ms) (↓)
DP	152.9
MDP	99.9

Table 7: Average duration error comparisons on the VCTK dataset. DP denotes the duration predictor used in Tan et al. (2021); Bai et al. (2022) and MDP denotes the masked duration predictor in our FluentSpeech.

by:

$$MCD[dB] = \frac{10}{\ln 10} \sqrt{2 \sum_{i=1}^L (m_i^t - m_i^c)^2}, \quad (6)$$

where L is the order of mel cepstrum and L is 34 in our implementation.

The traditional PESQ measure is given by:

$$PESQ = a_0 + a_1 D_{ind} + a_2 A_{ind}, \quad (7)$$

where a_0, a_1, a_2 are the parameters, D_{ind} represents the average disturbance value and A_{ind} represents the average asymmetrical disturbance values.

STOI is a function of a TF-dependent intermediate intelligibility measure, which compares the temporal envelopes of the clean and degraded speech in short-time regions by means of a correlation coefficient. The following vector notation is used to denote the short-time temporal envelope of the clean speech:

$$x_{j,m} = [X_j(m-N+1), X_j(m-N+2), \dots, X_j(m)]^T, \quad (8)$$

where $N = 30$ which equals an analysis length of 384 ms.

Instructions Shortcuts How natural (i.e. human-sounding) is this recording? Please focus on examining the naturalness (noise, timbre, sound clarity and high-frequency details) of testing audio, testers could...

Instructions

Listen to the sample of computer generated speech and assess the quality of the audio based on how close it is to natural speech. The words in the audio are shown in the Original utterance below.

For better results, wear headphones and work in a quiet environment.

[More Instructions](#)

Transcripts: \${original_text}

0:00 / 0:00

Select an option

Excellent - Completely natural speech - 5	1
4.5	2
Good - Mostly natural speech - 4	3
3.5	4
Fair - Equally natural and unnatural speech - 3	5
2.5	6
Poor - Mostly unnatural speech - 2	7

(a) Screenshot of MOS testing for speech quality in the speech editing evaluation.

Instructions Shortcuts How similar is this recording to the reference audio? Please focus only on the similarity of the speaker to the reference, and ignore the differences of content, grammar, or audio qual...

Instructions

Listen to the sample of computer generated speech and assess: How similar is this recording to the reference audio? Please focus only on the similarity of the speaker to the reference, and ignore the differences of content, grammar, or audio quality. The words in the audio are shown in the Original utterance below.

For better results, wear headphones and work in a quiet environment.

[More Instructions](#)

Reference audio:

0:00 / 0:00

Corresponding transcripts: \${ref_text}

Testing audio:

0:00 / 0:00

Corresponding transcripts: \${ref_text}

Select an option

Excellent - Completely similar speaker - 5	1
4.5	2
Good - Mostly similar speaker - 4	3
3.5	4
Fair - Equally similar and dissimilar speaker - 3	5
2.5	6
Poor - Mostly dissimilar speaker - 2	7

(b) Screenshot of MOS testing for speaker similarity in the speech editing evaluation.

Instructions Shortcuts which of the audio has better quality? Please focus on the audio quality and ignore other factors.

Instructions

Listen to the sample of computer generated speech and assess: which of the audio has better quality? Please focus on the audio quality and ignore other factors. The words in the audio are shown in the Original utterance below.

For better results, wear headphones and work in a quiet environment.

[More Instructions](#)

Source audio:

0:00 / 0:00

Target audio:

0:00 / 0:00

Corresponding transcripts: \${ref_text}

Select an option

Source Better 1	1
Neural 0	2
Target Better - 1	3

(c) Screenshot of average preference score testing for speech quality in the stutter removal evaluation.

Instructions Shortcuts which of the audio sounds more fluent? Please focus on the fluency and ignore other factors.

Instructions

Listen to the sample of computer generated speech and assess: which of the audio sounds more fluent? Please focus on the fluency and ignore other factors. The stutter in the audio typically sounds like "smm, uhhh,hmmm" or other words repetition. The words in the audio are shown in the Original utterance below.

For better results, wear headphones and work in a quiet environment.

[More Instructions](#)

Source audio:

0:00 / 0:00

Target audio:

0:00 / 0:00

Corresponding transcripts: \${ref_text}

Select an option

Source Better 1	1
Neural 0	2
Target Better - 1	3

(d) Screenshot of average preference score testing for speech fluency in the stutter removal evaluation.

Figure 5: Screenshots of subjective evaluations.

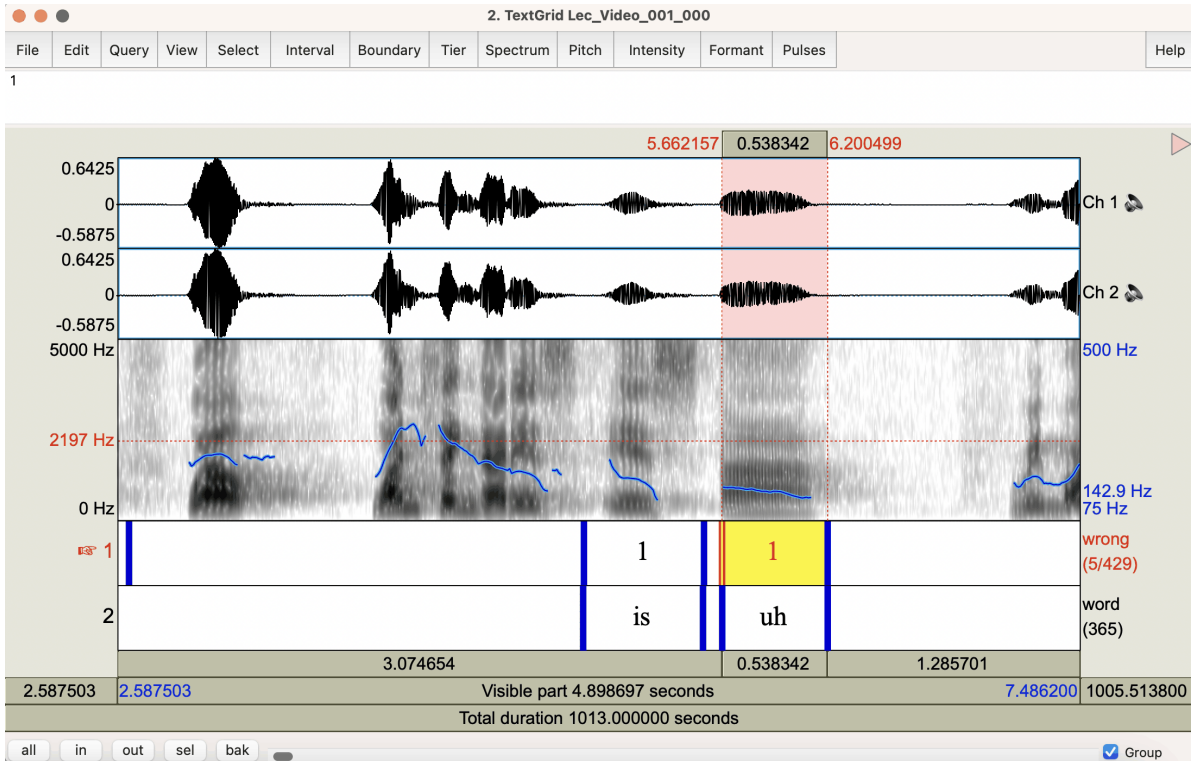


Figure 6: Screenshot of our SASE dataset.

Method	Average Pitch Error (\downarrow)
EditSpeech	5571
CampNet	6758
A ³ T	4595
FluentSpeech w/ PP	4134
FluentSpeech w/ MPP	2276

Table 8: Average pitch error comparisons on the VCTK dataset. PP denotes the pitch predictor proposed in Ren et al. (2020) and MPP denotes the masked pitch predictor in our FluentSpeech.

B Detailed analysis of duration and pitch

To further dive into the detailed performance of our model, we evaluate the duration and pitch errors between our FluentSpeech and the baseline models. For duration errors, the ground truth duration is obtained from the Montreal Forced Aligner (MFA) (McAuliffe et al., 2017). We calculate MSE of word-level durations for the duration predictor (DP) used in Tan et al. (2021); Bai et al. (2022) and the masked duration predictor (MDP) in FluentSpeech. The results on the VCTK dataset are shown in Table 7. It can be seen that the masked duration predictor predicts more accurate duration, demonstrating the effectiveness of the masked prediction training. For pitch errors, we compare our FluentSpeech with all other baseline models.

We firstly extract frame-level pitch information using parselmouth⁴, then calculate the MSE of the mean pitch distance between the model-generated speeches and the ground-truth speeches. The results on the VCTK dataset are shown in table 8. It can be seen that FluentSpeech achieves the lowest average pitch error. Moreover, the average pitch error of FluentSpeech with the masked pitch predictor (MPP) is significantly lower than the FluentSpeech with the pitch predictor proposed in Ren et al. (2020), demonstrating the effectiveness of our masked pitch predictor.

C More details of SASE dataset

The SASE dataset consists of approximately 40 hours of spontaneous speech recordings from 46 speakers with various accents. The speech recordings are crawled from online lectures and courses with accurate official transcripts. Each recording is sampled at 22050 Hz with 16-bit quantization. We substitute the speakers' names with speaker IDs to protect their personal information, and the dataset can only be accessed for research purposes.

To obtain the time-aligned stutter labels, we recruit annotators from a crowdsourcing platform, Zhengshu Technology, to label the stuttering re-

⁴<https://github.com/YannickJadoul/ParseImouth>

gion according to the audio and transcription. Specifically, the stuttering region may be 1) stammers and repetitive words, for instance, “I am go...go...going... out for a...a...a... trip”; 2) filled pauses (FP) such as “em, um, then, due to, uh, the speaker’s custom of speaking”; 3) sudden occasions such as cough, voice crack, etc. The annotators are asked to mark the corresponding time boundaries and give the stuttering label as shown in Figure 6. We then use the given timestamps in the official transcriptions to cut the audio and text into fragments ranging from 7 to 10 seconds. Finally, we convert each text sequence into phoneme sequence with an open-source grapheme-to-phoneme tool⁵. The audio samples in our SASE dataset are available at <https://speechai-demo.github.io/FluentSpeech/>.

⁵<https://github.com/Kyubyong/g2p>

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Section 6
- A2. Did you discuss any potential risks of your work?
Section 7
- A3. Do the abstract and introduction summarize the paper’s main claims?
Abstract and Section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section 4

- B1. Did you cite the creators of artifacts you used?
Section 4
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Section 4
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Section 4 and Appendix C
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Section 4 and Appendix C
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Section 4 and Appendix C
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Section 4 and Appendix C

C Did you run computational experiments?

Section 4

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 4

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Section 4
- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?
Section 4
- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?
Section 4
- D** **Did you use human annotators (e.g., crowdworkers) or research with human participants?**
Section 4, Appendix A.3, and Appendix C
- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?
Appendix A.3 and Appendix C
- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?
Appendix A.3 and Appendix C
- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?
Appendix A.3 and Appendix C
- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
Appendix A.3 and Appendix C
- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
Appendix A.3 and Appendix C