

FedSpeech: Federated Text-to-Speech with Continual Learning

Ziyue Jiang¹, Yi Ren¹, Ming Lei² and Zhou Zhao¹

¹Zhejiang University

²Alibaba Group

ziyuejiang341@gmail.com, rayeren@zju.edu.cn, lm86501@alibaba-inc.com, zhaozhou@zju.edu.cn

Abstract

Federated learning enables collaborative training of machine learning models under strict privacy restrictions and federated text-to-speech aims to synthesize natural speech of multiple users with a few audio training samples stored in their devices locally. However, federated text-to-speech faces several challenges: very few training samples from each speaker are available, training samples are all stored in local device of each user, and global model is vulnerable to various attacks. In this paper, we propose a novel federated learning architecture based on continual learning approaches to overcome the difficulties above. Specifically, 1) we use gradual pruning masks to isolate parameters for preserving speakers' tones; 2) we apply selective masks for effectively reusing knowledge from tasks; 3) a private speaker embedding is introduced to keep users' privacy. Experiments on a reduced VCTK dataset demonstrate the effectiveness of FedSpeech: it nearly matches multi-task training in terms of multi-speaker speech quality; moreover, it sufficiently retains the speakers' tones and even outperforms the multi-task training in the speaker similarity experiment.

1 Introduction

Federated learning has become an extremely hot research topic due to the "data island" issue in recent years. After the concept of federated learning is proposed by Google recently [Konečný *et al.*, 2016], most federated learning researches focus on computer vision [McMahan *et al.*, 2017], natural language processing [Hard *et al.*, 2018], and voice recognition [Leroy *et al.*, 2019]. Collaborative training brings great improvements to the performance of those tasks. However, few of them pay attention to federated text-to-speech (TTS) tasks which aim to synthesize natural speech of multiple users with a few audio training samples stored in each user's local device. Due to the nature of current neural network based systems, federated TTS tasks face several challenges below:

- **Lack of training data.** Although neural network-based TTS models [Shen *et al.*, 2018; Ren *et al.*, 2020] can generate high-quality speech samples, they suffer from the lack

of robustness (e.g., word skipping and noise) and significant quality degradation when the training dataset's size is reduced. However, in federated TTS scenarios, each user has a limited number of audio training samples, especially for low-resource languages.

- **Strict data privacy restrictions.** In federated scenarios, training samples are all stored in device of each user locally, which makes multi-task (multi-speaker) training impossible. However, in typical federated aggregation training [McMahan *et al.*, 2017], even small gradients from other speakers may greatly hurt the tone of a specific speaker due to the catastrophic forgetting issues. Accurately maintaining the tone of each speaker is difficult.
- **Global model is vulnerable to various attacks.** Typical communication architectures used in federated learning [McMahan *et al.*, 2017; Rothchild *et al.*, 2020] aggregate the information (e.g., gradients or model parameters) and train a global model. Although the local data are not exposed, the global model may leak sensitive information about users.

However, traditional federated learning methods can not address the issues above simultaneously. Recently, continual lifelong learning has received much attention in deep learning studies. Among them, Hung *et al.* have proposed a method called "Compacting, Picking and Growing" [Hung *et al.*, 2019], which uses masks to overcome catastrophic forgetting in continual learning scenarios and has achieved significant improvements in image classification tasks. Inspired by their works, we focus on building a federated TTS system using continual learning techniques.

Thus, in order to bring the advantages of collaborative training into federated multi-speaker TTS systems, in this work, we propose a federated TTS architecture called FedSpeech, in which 1) we use gradual pruning masks to isolate parameters for preserving speakers' tones; 2) we apply selective masks for effectively reusing the knowledge from tasks under privacy restrictions; 3) a private speaker embedding is introduced to apply additional information and guarantee users' privacy. Different from Hung *et al.* [2019], we perform masks in a Transformer-based TTS model and select weights from both previous and later tasks to make it more equitable for all speakers. Our proposed FedSpeech can address the above-mentioned three challenges as follows:

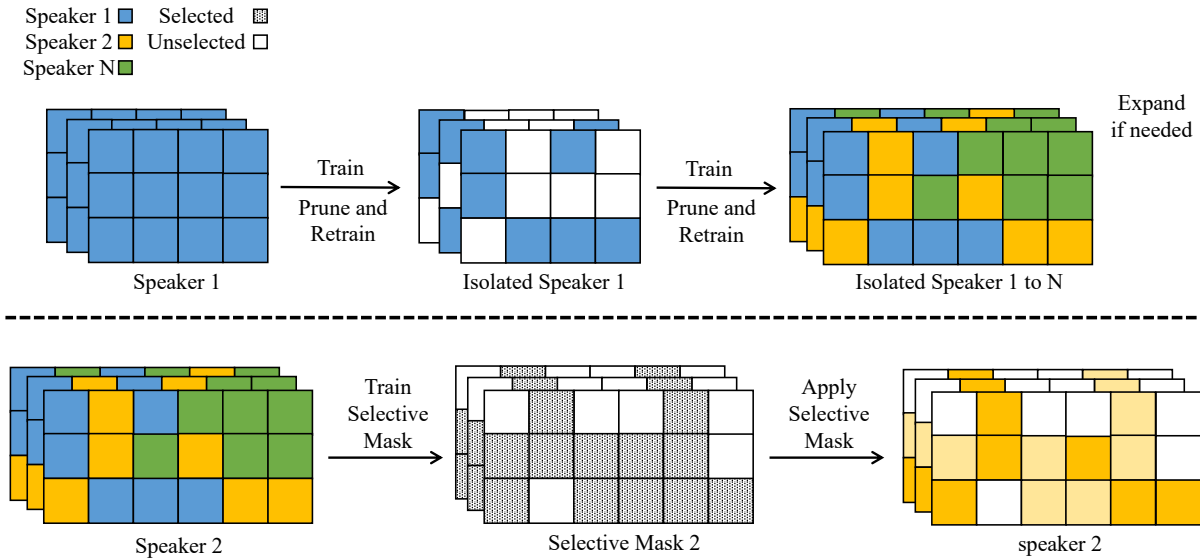


Figure 1: The two rounds of training process with FedSpeech. In Round 1, gradual pruning masks are performed to isolate weights for each speaker. If the weights preserved for a certain speaker are less than the threshold, the model will expand. In Round 2, we take speaker 2 as the example. Selective masks are trained to reuse the knowledge from the weights preserved for other speakers.

- With selective masks, FedSpeech can effectively benefit from collaborative training to lessen the influence of limited training data.
- Gradual pruning masks isolate the parameters of different speakers to overcome catastrophic forgetting issues. Thus, FedSpeech avoids the issue of tone changes for all speakers.
- The private speaker embedding is introduced coupled with two types of masks above to preserve the privacy and avoid various attacks for speakers.

We conduct experiments on a reduced VCTK dataset¹ to evaluate FedSpeech. The results show that in terms of speech quality, FedSpeech is nearly equivalent to joint training, which breaks the privacy rules. Moreover, FedSpeech achieves even higher speaker similarity scores than joint training due to our parameter isolation policy. We attach some audio files generated by our method in the supplementary materials².

2 Background

In this section, we briefly overview the background of this work, including text to speech (TTS), federated learning, and continual learning.

Text to Speech. Aiming to synthesize natural speech, text to speech (TTS) [Shen *et al.*, 2018; Yamamoto *et al.*, 2020] remains a crucial research topic. Recently, many advanced techniques have dominated this field. From concatenative synthesis with unit selection, statistical parametric synthesis to neural network based parametric synthesis and end-to-end

models [Shen *et al.*, 2018; Ren *et al.*, 2019; Ren *et al.*, 2020; Kim *et al.*, 2020], the quality of the synthesized speech is closer to the human voice. However, in federated TTS tasks, most neural TTS systems generate audios with quality decline when the training dataset’s size is significantly reduced. In this work, we efficiently reuse knowledge from different tasks, which resolves the problem above.

Federated Learning. Strictly under the privacy restrictions among distributed edge devices, federated learning aims to train machine learning models collaboratively [Li *et al.*, 2019]. In federated learning, typical communication architectures [McMahan *et al.*, 2017; Hard *et al.*, 2018; Rothchild *et al.*, 2020] usually aggregate the information (e.g., gradients or model parameters) and train a global model, which can be centralized design or decentralized design. However, the proposed methods above are vulnerable to inference attacks and may expose sensitive information about the users in TTS tasks. Moreover, the gradients from the other speakers may result in catastrophic forgetting and greatly hurt one speaker’s tone. In this work, we utilize a private speaker embedding policy to protect users’ privacy. Besides, we adopt two kinds of parameter masks in the training process and combine them in the inference process to retain tones and transfer knowledge.

Continual Learning. Continual learning aims at overcoming the catastrophic forgetting issues of neural networks when tasks arrive sequentially. In continual learning scenarios, a mechanism should be introduced to continually accumulate knowledge over different tasks without the need to retrain from scratch. Simultaneously, the mechanism should ensure a good compromise between the model’s stability and plasticity. Continual learning approaches are mainly structured into three main groups: replay, regularization-based, and pa-

¹In order to simulate the low-resource language scenarios, we randomly select 100 audio samples from each speaker for training.

²Synthesized speech samples can be found in: https://fedspeech.github.io/FedSpeech_example/

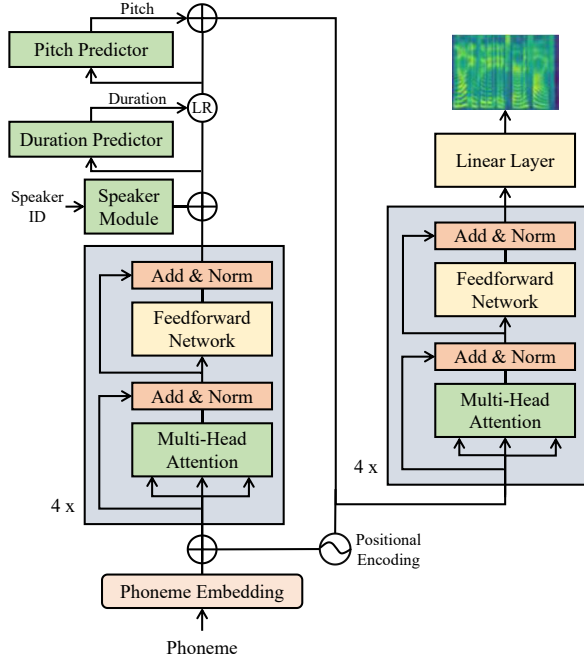


Figure 2: The overall architecture for FedSpeech. \oplus denotes the element-wise addition operation. LR denotes the length regulator proposed in FastSpeech [Ren *et al.*, 2019].

parameter isolation methods. Replay methods [Rolnick *et al.*, 2019] explicitly retrain on a subset of stored samples from previous tasks, which breaks the users’ privacy rules. Prioritizing privacy, regularization-based methods [Li and Hoiem, 2017] introduce new loss functions to distill previous knowledge or penalize important parameters’ updates. However, slight weight changes may significantly affect the voice of a specific speaker. Considering the privacy and performance, the parameter isolation method [Mallya and Lazebnik, 2018; Hung *et al.*, 2019] is a solution to the difficulties above. Our FedSpeech adopts the progressive pruning masks in Hung *et al.* [2019] and modifies the selective masks referring to Mallya *et al.* [2018] to transfer knowledge from other speakers while keeping privacy.

3 FedSpeech

In this section, we first introduce the architecture design of FedSpeech based on the novel feed-forward transformer structure proposed in Ren *et al.* [2020]. Then we introduce a private speaker embedding to apply additional information and keep users’ privacy. It is trained to capture the speaker features from the latent space and stores sensitive information, which is indispensable in the inference stage. The private speaker embedding should be preserved in each user’s device locally and can not be obtained by other users in order to protect users’ privacy. In order to solve the data scarcity issue and further address the privacy issue in federated TTS tasks, we propose a two-round sequential training process, which is a common setting in continual learning. We adopt two kinds of masks and a speaker module to isolate param-

eters and effectively reuse different speakers’ knowledge while keeping privacy. We describe the above components in detail in the following subsections.

3.1 FedSpeech Architecture

The overall model architecture of FedSpeech is shown in Figure 2. The encoder converts the phoneme embedding sequence into the phoneme hidden sequence, and then we add different variance information such as duration and pitch into the hidden sequence, finally the mel-spectrogram decoder converts the adapted hidden sequence into mel-spectrogram sequence in parallel. We adopt the feed-forward Transformer block, which is a stack of self-attention [Vaswani *et al.*, 2017] layer and 1D-convolution feed-forward network as in FastSpeech [Ren *et al.*, 2019], as the basic structure for the encoder and mel-spectrogram decoder. Besides, we adopt a pitch predictor and a duration predictor to introduce more information. Each of them consists of a 2-layer 1D-convolutional network with ReLU activation, followed by the layer normalization and the dropout layer, and an extra linear layer to project the hidden states into the output sequence. In training, we take the ground-truth values of duration and pitch extracted from the recordings as input into the hidden sequence to predict the target speech. Meanwhile, we use the ground-truth duration and pitch values as targets to train the predictors. And their outputs are used in inference to synthesize target speeches.

3.2 Speaker Module

In order to control voice by estimating the speaker features from the latent space and protect privacy, we introduce a private speaker module which is a trainable lookup table [Jia *et al.*, 2018] taking speaker’s identity number S_{id} as input and generating the speaker representation $R = \{r_1, r_2, \dots, r_n\}$, where n is the hidden size of the model. And then the speaker representation R is passed to the output of the encoder as additional key information to control the tone characteristics in both training and inference. Each speaker trains and keeps his own set of module parameters in consideration of privacy so that the others can not synthesize his voice even with his S_{id} .

3.3 Two Rounds of Sequential Training

Our work abandons the federated aggregation training setup, due to its catastrophic forgetting issue. As shown in Figure 1, we follow a sequential training setup, which is common in continual learning [Mallya and Lazebnik, 2018; Mallya *et al.*, 2018; Hung *et al.*, 2019]. It is worth noting that only knowledge from previous speakers can be used by the current speaker in Hung *et al.* [2019]. Different from their works, we propose the two rounds of sequential training. In the first round of training, the model separately learns and fixes a portion of weights for each speaker, so that in the second round of training we can selectively reuse the knowledge from both previous and later speakers. In the following, we present our two training rounds in detail.

First Round - Gradual Pruning Masks

In the first round of training, the gradual pruning masks in Figure 1 are calculated to isolate parameters for each speaker.

The speakers from 1 to N are denoted as $S_{1:N}$. The tasks for $S_{1:N}$ are $T_{1:N}$ and start sequentially. For simplicity, we take S_t as the example. When T_t starts, the global model M_g is firstly sent to S_t and trained using his private data until convergence. The learned weight matrix for layer i is denoted as W_1^i . We then gradually prune a portion of the smallest weights in W_1^i for each layer, set them to 0, and retrain the other weights to restore performance. Finally, the weights are divided into three parts: 1) the zero-valued weights released for later speakers $S_{t+1:N}$; 2) the fixed weights $W_S^{1:t-1}$ preserved by previous speakers $S_{1:t-1}$; 3) the weights W_S^t preserved by S_t . If the released weights for later speakers $S_{t+1:N}$ are less than a threshold λ , we will expand the hidden size of the model by μ . The pruning state is stored in the gradual pruning masks denoted as m_p . We then fix W_S^t and send m_p and M_g (except for the private speaker module) to the device of next speaker S_{t+1} to continue sequential training. When the first round ends, each speaker preserves a certain portion of weights denoted as $W_S^{1:N}$ represented by m_p . As the weights of each task are fixed, each speaker can perfectly retain their tones in inference. Finally, the m_p and M_g are sent to the devices of $S_{1:N}$. Thus, each speaker has m_p , M_g , and his preserved parameters of the private speaker module.

Second Round - Selective Masks

In the second round of training, we introduce the selective masks to transfer knowledge from speakers to address the data scarcity issue. The selective masks in Figure 1 are trained to automatically select useful weights preserved by speakers. Instead of selecting weights from previous tasks [Mallya *et al.*, 2018], we propose a modified selection procedure to select weights from all tasks, which is more equitable for every speaker (especially for more previous speakers) in federated TTS tasks. For a specific speaker S_t , our two rounds of training abandon the joint training of W_S^t and the selective masks, which leads to slight performance degradation. But for each speaker, we make it possible to select weights from both previous and later tasks, which leads to significant improvements overall.

Assume that when the first round ends, the weights of M_g are divided into several portions $W_S^{1:N}$ which are preserved by $S_{1:N}$. In order to benefit from collaborative training while keeping privacy, we introduce a learnable mask $m_b \in \{0, 1\}$ to transfer knowledge from the parameters preserved by other speakers. We use the piggyback approach [Mallya *et al.*, 2018] that learns a real-valued mask m_s and applies a threshold for binarization to construct m_b . For a certain speaker S_t , mask m_b is trained on his local dataset to pick the weights from other speakers by $m_b^t \odot W_S^{1:t-1} \cup W_S^{t+1:N}$.

We describe the mask training process of the selective masks in 1D convolution layer as the example. At task t , M_g (i.e., $W_S^{1:N}$) is fixed. Denote the binary masks as m_b^t . Then the equation for the input-output relationship is given by

$$\tilde{W} = m_b^t \odot W \quad (1)$$

$$y(N_i, C_{out_j}) = b(C_{out_j}) + \sum_{k=0}^{C_{in}-1} \tilde{W}(C_{out_j}, k) \star x(N_i, k) \quad (2)$$

Algorithm 1 Two Rounds of Training with FedSpeech

Input: number of tasks (speakers) N , training samples with paired text and audio $\mathcal{D}_t = \{(x_k, y_k)\}_{k=1}^K$ of speaker t .

Initialize: randomly initialize the original model M_g ; set all elements in the gradually pruning masks m_p to zeros; set all elements in the real-valued selective mask m_s^t of speaker t to α and the binary selective mask m_b^t of speaker t to zeros; initialize the threshold λ , μ , and σ .

Annotate: $T_{1:N}$ denote tasks from 1 to N ; $S_{1:N}$ denote speakers from 1 to N ; W_S^t denotes the weights in M_g preserved for speaker t .

First Round:

for task $t = T_{1:N}$ **do**

 Train the released weights until convergence

if $t \neq T_N$ **then**

 Gradually prune a portion of the smallest weights, set them to 0 and retrain other weights to restore performance

 Release the zero-valued weights for next task

end if

 store the pruning state in m_p and fix W_S^t

if the released weights are less than λ **then**

 Expand the hidden size of M_g by μ

end if

end for

 Send m_p and M_g to the local devices of $S_{1:N}$

Second Round:

fix $W_S^{1:N}$

for task $t = T_{1:N}$ **in parallel in local devices do**

 Initialize m_s^t

 Train m_s^t , m_b^t until convergence

S_t preserves m_b^t for inference

end for

where \star is the valid cross-correlation operator, N is the batch size, and C_{in}/C_{out} denote the number of in/output channels.

In the backpropagation process, the m_b^t is not differentiable. So we introduce the real-valued selective masks denoted as m_s^t . Denote σ as the threshold for selection. As in Hung *et al.* [2019], when training the binary mask m_b^t , we update the real-valued mask m_s^t in the backward pass; then m_b^t is quantized with a binarizer function β applied on m_s^t and used in the forward pass. After training, we discard m_s^t and only store m_b^t for inference. The equation for m_s^t is formulated by

$$m_b^t = \beta(m_s^t) = \begin{cases} 1 & \text{if } m_s^t > \sigma \\ 0 & \text{else} \end{cases} \quad (3)$$

$$\begin{aligned} \delta m_s^t(C_{out_j}, k) &= \frac{\partial L}{\partial m_b^t(C_{out_j}, k)} \\ &= \frac{\partial L}{\partial y(N_i, C_{out_j})} \cdot \frac{\partial y(N_i, C_{out_j})}{\partial m_b^t(C_{out_j}, k)} \\ &= \delta y(N_i, C_{out_j}) \cdot \sum_{k=0}^{C_{in}-1} \tilde{W}(C_{out_j}, k) \star x(N_i, k) \end{aligned} \quad (4)$$

Model Inference

For simplicity, we describe the inference stage using the example of S_t . Now S_t has m_p , m_b^t , M_g and his locally preserved parameters of speaker module. We pick the weights W_S^t using m_p and selectively reuse the weights in $W_S^{1:t-1} \cup W_S^{t+1:N}$ using m_b^t . The unused weights are fixed to zero in order not to hurt the tone of S_t . The overall procedure of the two rounds of training with FedSpeech is presented in Algorithm 1.

4 Experimental Setup

4.1 Datasets

We conduct experiments on the VCTK dataset [Veaux *et al.*, 2017], which contains approximate 44 hours of speech uttered by 109 native English speakers with various accents. Each speaker reads out about 400 sentences, most of which are selected from a newspaper plus the Rainbow Passage and an elicitation paragraph intended to identify the speaker’s accent. To simulate the low-resource language scenarios, we randomly select and split the samples from each speaker into 3 sets: 100 samples for training, 20 samples for validation, and 20 samples for testing. We randomly select 10 speakers denoted as task 1 to 10 for evaluation respectively. To alleviate the mispronunciation problem, we convert the text sequence into the phoneme sequence with an open-source grapheme-to-phoneme conversion tool [Park, 2019]. Following Shen *et al.* [2018], we convert the raw waveform into mel-spectrograms and set frame size and hop size to 1024 and 256 with respect to the sample rate 22050.

4.2 Model Configuration

FastSpeech 2 Model. FedSpeech is based on FastSpeech 2 [Ren *et al.*, 2020], which consists of 4 feed-forward Transformer blocks both in the encoder and the mel-spectrogram decoder. The hidden sizes of the self-attention and 1D convolution in each feed-forward Transformer block are all set to 256, which will grow if needed. The number of attention heads is set to 2. The output linear layer converts the 256-dimensional hidden into the 80-dimensional mel spectrogram.

Gradual Pruning and Selective Masks. In the training stage, the type of gradual pruning masks is short integer. We use speakers’ identity number stored in the gradual pruning masks to mark the parameters so as to isolate them. The type of selective masks is parameter when training. After training, the selective masks will be quantized to binary type. The initial value of selective masks’ unit is 0.01 and the σ in the Equation 3 is set to 0.005.

4.3 Training and Inference

We use 1 Nvidia 1080 Ti GPU, with 11GB memory. Each batch contains about 20,000 mel-spectrogram frames. We use Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.98$, $\epsilon = 10^{-9}$ and follow the learning rate schedule in [Vaswani *et al.*, 2017]. In all experiments, we choose 10 speakers. For each speaker, it takes 4k steps for FedSpeech model training (including the gradual pruning masks’ training) and 1k steps for selective

Method	MOS	γ
GT	3.97 ± 0.050	-
GT (Mel + PWG)	3.85 ± 0.051	-
Multi-task	3.82 ± 0.050	1.7
Scratch	3.71 ± 0.056	10
Finetune	3.72 ± 0.051	10
FedAvg [McMahan <i>et al.</i> , 2017]	3.58 ± 0.067	1.7
CPG [Hung <i>et al.</i> , 2019]	3.74 ± 0.053	1.7
FedSpeech	3.77 ± 0.052	1.7

Table 1: The MOS with 95% confidence intervals. γ means the model expansion rate compared with FedSpeech with 256 hidden size.

masks’ training. The total training time of FedSpeech for 10 speakers is 14 hours. For those baselines without these masks, we apply their approach to FastSpeech 2 model [Ren *et al.*, 2020] and train the model for 5k steps for a fair comparison. In the inference stage, we use the pre-trained Parallel WaveGAN (PWG) [Yamamoto *et al.*, 2020] to transform mel-spectrograms generated by FedSpeech into audio samples.

5 Results

In this section, we evaluate the performance of FedSpeech in terms of audio quality, speaker similarity, and ablation studies.

5.1 Audio Quality

We evaluate the MOS (mean opinion score) on the test set to measure the audio quality. The setting and the text content are consistent among different models so as to exclude other interference factors, only examining the audio quality. Each audio is judged by 10 native English speakers. We compare the MOS of the audio samples generated by our model with other systems, which include 1) *GT*, the ground truth audio in VCTK. 2) *GT (Mel + PWG)*, where we first convert the ground-truth audio into mel-spectrograms, and then convert the mel-spectrograms back to audio using Parallel WaveGAN (PWG) [Yamamoto *et al.*, 2020]; 3) *Multi-task*, jointly training without privacy restrictions; 4) *Scratch*, learning each task independently from scratch; 5) *Finetune*, finetuning from a previous model randomly selected and repeats the process 5 times (for task 1, Finetune is equal to Scratch); 6) *FedAvg* [McMahan *et al.*, 2017], aggregating the local information (e.g., gradients or model parameters) and train a global model. 7) *CPG* [Hung *et al.*, 2019], a parameter isolation method used in continual learning. We denote 3) as upper bounds and the others as baselines. Correspondingly, all the systems in 3), 4), 5), 6), 7) and FedSpeech use a pre-trained PWG as the vocoder for a fair comparison.

The MOS results are shown in Table 1. From the table, we can see that FedSpeech achieves the highest MOS compared with all baselines. It is worth mentioning that FedSpeech outperforms CPG, which illustrates the effectiveness of selectively reusing knowledge from both previous and later speakers. Besides, the results of FedAvg are significantly worse

Method	Task 1	Task 5	Task 10	<i>Avg.</i>	γ
Scratch	0.8600	0.8725	0.8647	0.8571	10
Finetune	0.8600	0.8782	0.8802	0.8651	10
Multi-task	0.8845	0.8738	0.8784	0.8738	1.7
FedAvg	0.7566	0.5736	0.7057	0.7020	1.7
CPG	0.8550	0.8798	0.8847	0.8688	1.7
FedSpeech	0.8861	0.8884	0.8833	0.8786	1.7

Table 2: The comparison of speaker similarity between baselines and FedSpeech. *Avg.* means the average of 10 tasks, and γ means the model expansion rate compared with FedSpeech with 256 hidden size.

than other methods, which means the gradients from other speakers greatly affect the tone of each speaker. Moreover, the MOS of FedSpeech on VCTK is close to multi-task training (the upper bound). These results demonstrate the advantages of FedSpeech for federated multi-speaker TTS tasks.

5.2 Comparison Experiments for Speaker Similarity in TTS

We conduct the speaker similarity evaluation on the test set to measure the similarity between the synthesized audio and the ground truth audio. To exclude other interference factors, we keep the text content consistent among different models. For each task, we derive the high-level representation vector that summarizes the characteristics of the speaker’s voice using the encoder³ implemented from Wan *et al.* [2018]. Specifically, the encoder is a 3-layer LSTM with projection, which is pretrained for extracting speaker’s tone embeddings. Cosine similarity is a standard measure of similarity of speaker representation vectors, and is defined as $\text{cos_sim}(A, B) = A \cdot B / \|A\| \|B\|$. The results range from -1 to 1, and higher values indicate that the vectors are more similar. We calculate the cosine similarity between the speaker representation vectors of the synthesized audios and the ground truth audios as the criterion for evaluation. We compare the results of the audio samples generated by our model with those systems described in subsection 5.1.

The results are shown in Table 2. Our FedSpeech scores the highest on average, and even higher than multi-task, the upper bound. It means FedSpeech can retain the voice of each speaker better in the inference stage and demonstrates the effectiveness of parameter isolation. Moreover, in task 1 the result of FedSpeech is significantly higher than CPG. It can be seen that selectively reusing knowledge from both previous and later speakers brings great advantages to the speakers so that all speakers in federated multi-speaker TTS task can obtain better voices.

5.3 Ablation Studies

We conduct ablation studies to demonstrate the effectiveness of several components in FedSpeech, including the gradual pruning masks and the selective masks. We conduct audio quality and speaker similarity evaluation for these ablation

Setting	MOS	Similarity
FedSpeech	3.77 ± 0.052	0.8786
- GPM	3.74 ± 0.075	0.8725
- SM	3.72 ± 0.074	0.8722
- SM - GPM	3.23 ± 0.099	0.6304

Table 3: MOS and speaker similarity comparison in the ablation studies. SM means the selective masks, GPM means the gradual pruning masks and similarity is the cosine similarity described in subsection 5.2.

studies. In this experiment, the model is well trained with our proposed first round of training so that we can focus on the effectiveness of the proposed masks.

Audio Quality. For measuring audio quality, we conduct the MOS evaluation in which each audio is judged by 10 native English speakers. As shown in Table 3, removing the gradual pruning masks or removing the selective masks does not result in significant quality decline, which means the selective masks have the ability to automatically select weights preserved by the gradual pruning masks. However, removing both types of masks leads to catastrophic quality degradation.

Similarity. We conduct the speaker similarity evaluation as subsection 5.2. As shown in Table 3, simply removing the selective masks or the gradual pruning masks results in slight performance degradation, while removing both of them leads to a catastrophic decline. It can be seen that the gradual pruning masks perfectly preserve the tone of each speaker. Besides, the selective masks have the ability to automatically select weights preserved by the gradual pruning masks and combining them leads to a better result.

6 Conclusions

In this work, we have proposed FedSpeech, a high-quality multi-speaker TTS system, to address the “data island” issue in federated multi-speaker TTS tasks. FedSpeech is implemented based on the two rounds of training with the feed-forward transformer network proposed in Ren *et al.* [2020] and consists of several key components including the selective masks, the progressive pruning masks, and the private speaker module. Experiments on a reduced VCTK dataset (the training set is reduced to a quarter for each speaker to simulate low-resource language scenarios) demonstrate that our proposed FedSpeech can nearly match the upper bound, multi-task training in terms of speech quality, and even significantly outperforms all systems in speaker similarity experiments. For future work, we will continue to improve the quality of the synthesized speech and propose a new mask strategy to compress the model and speed up training. Besides, we will also apply FedSpeech to zero-shot multi-speaker settings by using the private speaker module to generate our proposed masks.

³<https://github.com/resemble-ai/Resemblyzer>

References

- [Hard *et al.*, 2018] Andrew Hard, Kanishka Rao, Rajiv Mathews, Swaroop Ramaswamy, Françoise Beaufays, Sean Augenstein, Hubert Eichner, Chloé Kiddon, and Daniel Ramage. Federated learning for mobile keyboard prediction. *arXiv preprint arXiv:1811.03604*, 2018.
- [Hung *et al.*, 2019] Ching-Yi Hung, Cheng-Hao Tu, Cheng-En Wu, Chien-Hung Chen, Yi-Ming Chan, and Chu-Song Chen. Compacting, picking and growing for unforgetting continual learning. In *Advances in Neural Information Processing Systems*, pages 13669–13679, 2019.
- [Jia *et al.*, 2018] Ye Jia, Yu Zhang, Ron Weiss, Quan Wang, Jonathan Shen, Fei Ren, Patrick Nguyen, Ruoming Pang, Ignacio Lopez Moreno, Yonghui Wu, et al. Transfer learning from speaker verification to multispeaker text-to-speech synthesis. *Advances in neural information processing systems*, 31:4480–4490, 2018.
- [Kim *et al.*, 2020] Jaehyeon Kim, Sungwon Kim, Jungil Kong, and Sungroh Yoon. Glow-tts: A generative flow for text-to-speech via monotonic alignment search. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [Konečný *et al.*, 2016] Jakub Konečný, H Brendan McMahan, Daniel Ramage, and Peter Richtárik. Federated optimization: Distributed machine learning for on-device intelligence. *arXiv preprint arXiv:1610.02527*, 2016.
- [Leroy *et al.*, 2019] David Leroy, Alice Coucke, Thibaut Lavril, Thibault Gisselbrecht, and Joseph Dureau. Federated learning for keyword spotting. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6341–6345. IEEE, 2019.
- [Li and Hoiem, 2017] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.
- [Li *et al.*, 2019] Qinbin Li, Zeyi Wen, Zhaomin Wu, Sixu Hu, Naibo Wang, and Bingsheng He. A survey on federated learning systems: vision, hype and reality for data privacy and protection. *arXiv preprint arXiv:1907.09693*, 2019.
- [Mallya and Lazebnik, 2018] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7765–7773, 2018.
- [Mallya *et al.*, 2018] Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 67–82, 2018.
- [McMahan *et al.*, 2017] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pages 1273–1282. PMLR, 2017.
- [Park, 2019] Jongseok Park, Kyubyong Kim. g2pe. <https://github.com/Kyubyong/g2p>, 2019. v1.0.0.
- [Ren *et al.*, 2019] Yi Ren, Yangjun Ruan, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. FastSpeech: Fast, robust and controllable text to speech. In *Advances in Neural Information Processing Systems*, pages 3171–3180, 2019.
- [Ren *et al.*, 2020] Yi Ren, Chenxu Hu, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. FastSpeech 2: Fast and high-quality end-to-end text-to-speech. *arXiv preprint arXiv:2006.04558*, 2020.
- [Rolnick *et al.*, 2019] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. In *Advances in Neural Information Processing Systems*, pages 350–360, 2019.
- [Rothchild *et al.*, 2020] Daniel Rothchild, Ashwinee Panda, Enayat Ullah, Nikita Ivkin, Ion Stoica, Vladimir Braverman, Joseph Gonzalez, and Raman Arora. Fetchsgd: Communication-efficient federated learning with sketching. In *International Conference on Machine Learning*, pages 8253–8265. PMLR, 2020.
- [Shen *et al.*, 2018] Jonathan Shen, Ruoming Pang, Ron J Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, Rj Skerrv-Ryan, et al. Natural tts synthesis by conditioning wavenet on mel spectrogram predictions. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4779–4783. IEEE, 2018.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [Veaux *et al.*, 2017] Christophe Veaux, Junichi Yamagishi, Kirsten MacDonald, et al. Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit. University of Edinburgh. The Centre for Speech Technology Research (CSTR), 2017.
- [Wan *et al.*, 2018] Li Wan, Quan Wang, Alan Papir, and Ignacio Lopez Moreno. Generalized end-to-end loss for speaker verification. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4879–4883. IEEE, 2018.
- [Yamamoto *et al.*, 2020] Ryuichi Yamamoto, Eunwoo Song, and Jae-Min Kim. 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)*, pages 6199–6203. IEEE, 2020.