

Enhancing Zero-Shot Multi-Speaker TTS with Negated Speaker Representations

Yejin Jeon¹, Yunsu Kim², Gary Geunbae Lee^{1,3}

¹Graduate School of Artificial Intelligence, POSTECH, Republic of Korea

²aiXplain Inc. Los Gatos, CA, USA

³Department of Computer Science and Engineering, POSTECH, Republic of Korea

{jeonyj0612, gblee}@postech.ac.kr, yunsu.kim@aixplain.com

Abstract

Zero-shot multi-speaker TTS aims to synthesize speech with the voice of a chosen target speaker without any fine-tuning. Prevailing methods, however, encounter limitations at adapting to new speakers of out-of-domain settings, primarily due to inadequate speaker disentanglement and content leakage. To overcome these constraints, we propose an innovative negation feature learning paradigm that models decoupled speaker attributes as deviations from the complete audio representation by utilizing the subtraction operation. By eliminating superfluous content information from the speaker representation, our negation scheme not only mitigates content leakage, thereby enhancing synthesis robustness, but also improves speaker fidelity. In addition, to facilitate the learning of diverse speaker attributes, we leverage multi-stream Transformers, which retain multiple hypotheses and instigate a training paradigm akin to ensemble learning. To unify these hypotheses and realize the final speaker representation, we employ attention pooling. Finally, in light of the imperative to generate target text utterances in the desired voice, we adopt adaptive layer normalizations to effectively fuse the previously generated speaker representation with the target text representations, as opposed to mere concatenation of the text and audio modalities. Extensive experiments and validations substantiate the efficacy of our proposed approach in preserving and harnessing speaker-specific attributes vis-à-vis alternative baseline models.

Introduction

Recent neural network-based models have revolutionized the field of text-to-speech (TTS), generating natural speech that closely resembles human-like qualities (Shen et al. 2018; Ren et al. 2021). This has fueled a demand for *personalized* TTS capabilities, particularly voice imitation using multi-speaker TTS. Yet, the effective implementation of such a task presents formidable challenges. First, multi-speaker TTS models must be able to adapt to diverse speakers, including those unseen in the data, while maintaining fidelity across all text inputs. Second, there exist practical limitations in gathering sufficient speech data for each speaker for accurate imitation by TTS models. Thus, the ability to accurately replicate the voices of both seen and unseen

speakers with minimal audio data samples, or even from just one audio input, has gained paramount significance.

One avenue of investigation involves fine-tuning a pre-trained model initially trained on an extensive multi-speaker dataset, with a limited set of speech samples of the desired target speaker (Chen et al. 2019, 2021). While this few-shot approach has demonstrated good speaker adaptability, it is not impervious to inherent limitations. In particular, it necessitates additional training tailored to each target speaker to ensure accurate synthesis of speech in their specific voice. Moreover, when confronted with a scarcity of audio samples for fine-tuning, the resulting voice quality exhibits a noticeable degradation in speaker similarity.

In response to the challenges observed in few-shot synthesis, another branch of research have focused on extracting target speaker embeddings from a single reference audio sample in a zero-shot manner. Some investigations utilize x-vectors that are generated from pretrained speaker verification models, which in turn are trained with large voice datasets (Jia et al. 2018). However, not only does this method entail substantial data requisites, but also experiences limited speaker fidelity for out-of-domain speakers. Consequently, other endeavors have gravitated towards strategies of reduced data expenditure while conceptualizing speaker attributes as latent characteristics. Specifically, target speaker attributes are learned through unsupervised methods like Global Style Tokens (GST) (Choi et al. 2022) or variational autoencoders (VAE) (Zhao et al. 2022). While a certain degree of speaker adaptation is achievable through the aforementioned methods, speaker-dependent factors are not sufficiently disentangled from other components such as the uttered linguistic content. Thus, models learn to extract speaker embeddings that include undesired content information. Not only does this lead to discrepancies in speaker timbre, but also towards unstable synthesis since the TTS model learns to rely on the contextual cues from the intended target input text, as well as those provided via the reference audio.

Towards accurate and robust multi-speaker TTS, we propose a novel zero-shot and end-to-end model that incorporates the following fundamental aspects. (1) We introduce for the first time, a negated feature learning strategy that effectively generates speaker representations by computing it as the differentials between the full representation of the reference audio and its linguistic contents. Not

only does this strategy facilitate better disentanglement of target timbre information from other irrelevant speech factors, but also decrease content leakage that is introduced into target speaker representations. (2) We further implement multi-stream Transformers in order to generate multiple hypotheses such that various speaker attributes can be learned and preserved. By using attention pooling, a more optimal speaker representation is computed from the unification of various speaker hypotheses. (3) Knowledge of the ultimate speaker representations are injected into the backbone TTS pipeline using adaptive layer normalizations to enhance target text and speaker fusion. Through in-depth experimentation and verification, we demonstrate that our zero-shot unsupervised approach is more robust and achieves higher speaker similarity compared to previous baseline studies.

Related Work

Controllable Speech Generation Tasks

The pursuit of controllable speech generation has given rise to a range of diverse tasks. Among these, a pertinent task to multi-speaker TTS is voice conversion, wherein an utterance articulated by a source speaker is modified to make it sound as though spoken by a different target speaker. Since this transformation requires that the linguistic contents of the source audio remain unchanged, this necessitates the availability of parallel speech data where diverse speakers articulate identical text. Given this setting, several investigations have concentrated on achieving alignments between parallel spectral features by employing techniques like dynamic frequency warping (Helander et al. 2008). Nevertheless, the scarcity of readily available parallel datasets has prompted certain studies to resort to non-parallel data. Yet, these endeavors are constrained to voice conversion between a predefined set of speakers (Kaneko and Kameoka 2018; Kameoka et al. 2018) and are contingent on one-hot speaker representations (Kameoka et al. 2019; Hsu et al. 2016). In contrast, multi-speaker TTS demands precise speaker modelling that can generalize to any textual inputs, including those outside the established linguistic domain of the source audio. For this purpose, unlike unimodal acoustic modeling seen in voice conversion, multi-speaker TTS requires dual-modality modeling, and fusion of text and audio domains. Furthermore, model training is not confined to parallel speaker datasets.

Attribute-Specific Disentanglement

Due to its potential to govern the generative process of outcomes, the disentanglement of target features is a pivotal aspect that transcends various research domains. Foremost in this pursuit were endeavors within computer vision for style transfer; the semantic contents of the source image are imbued with the stylistic intricacies of the target image using pairwise correlation matrices (Gatys, Ecker, and Bethge 2016). This seminal work initiated an array of investigations such as those using different variants of instance normalizations (Dumoulin, Shlens, and Kudlur 2017; Ulyanov, Vedaldi, and Lempitsky 2017). This thematic trend subsequently diffused into adjacent areas such as machine trans-

lation, where each source-target translation pairs were annotated by the labels of the translators who generated that particular translation (Wang, Hoang, and Federico 2021). By doing so, translations could be generated in a translator-specific linguistic style pertaining to their grammatical preferences and verbosity. Similarly, the domain of speech has witnessed a large amount of work investigating the manipulation of both discrete and continuous attributes such as speaker accents (Zhang et al. 2022), emotion nuances (Lei et al. 2022), and prosodic variations (Hodari, Watts, and King 2019), each forming their own research sub-domains.

Towards multi-speaker TTS, speaker modeling from a reference audio is mainly conducted in two ways: 1) direct, fixed-length extraction of timbral attributes via a speaker encoder; and 2) additional disentanglement of target speaker embeddings from other observable factors. For the former, many works assume that timbre is a global feature in speech, and thus utilize coarse-grained methodologies such as pre-trained wav2vec models (Chen and Rudnicky 2022), hierarchical GSTs (An et al. 2019), and multivariate Gaussian mixture VAEs (Kulkarni, Colotte, and Jouviet 2022). The latter paradigm involves maximizing the distance between the directly generated speaker vectors with that of contextual representations. This is done by first pretraining the backbone TTS system with a single speaker dataset, which allows for the content encoder to encapsulate just contextual information. A Mutual-Information Neural Estimator (MINE) is then used to maximize the lower bound for mutual information through Donsker-Varadhan representation (Hu et al. 2020). Similarly in (Paul et al. 2021), Rényi Divergence is used to create more nuanced speaker representations by establishing distance between timbre and content information.

Methodology

Preliminaries

Given training dataset $D = (X_n, Y_n)_{n=1}^N$ of N samples and where X_n and Y_n form respective pairs of text and audio, the objective of conventional single-speaker TTS frameworks lies in generating acoustic representation \hat{Y}_n that faithfully corresponds to the input text X_n . This process comprises of an acoustic model and a vocoder. The acoustic model is further consisted of three components: a linguistic encoder responsible for generating a contextual representation for X_n , an alignment module, and a generator that produces intermediate acoustic representations in the form of mel-spectrograms \hat{M}_n . The vocoder then converts \hat{M}_n into \hat{Y}_n . In this work, we focus solely on the former acoustic model.

For the baseline acoustic model, we leverage the non-autoregressive FastSpeech model (Ren et al. 2021). Specifically, a text sequence $X_n = \{x_1, x_2, \dots, x_l\}$ of length L is converted to a context embedding via linguistic encoder \mathcal{E}_{text} , which is a feed-forward Transformer block. The alignment module is formulated as the variance adapter, which adds variance information to help facilitate one-to-many mapping necessary for text-to-speech. The variance adapter, which is comprised of pitch, energy, and duration predictors, are each made up of a 2-layer, 1-D convolutional network followed by ReLU activation, layer normalizations, dropout,

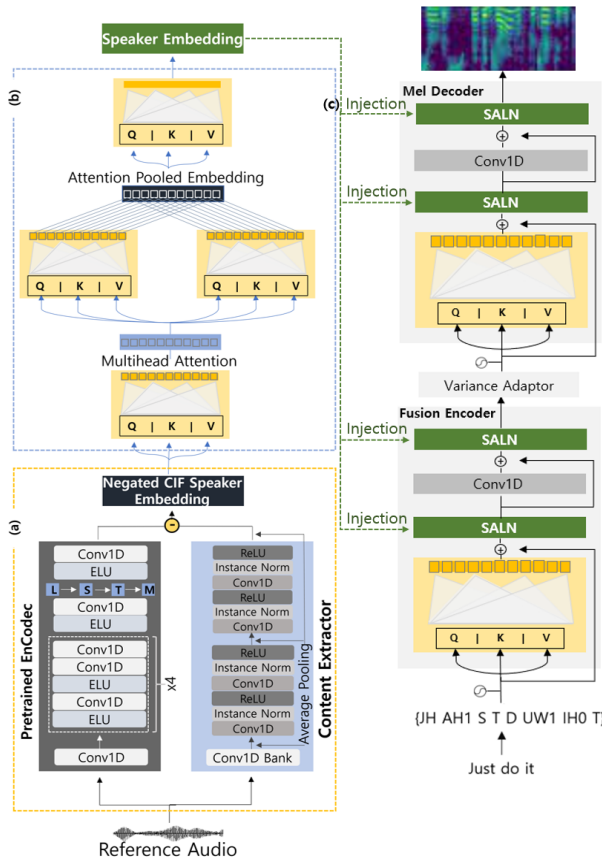


Figure 1: Main architecture: (a) Negated feature generation, which results in CIF embeddings, are passed to (b) multi-stream Transformers to formulate multi-perspective speaker representations, which are then (c) injected into the TTS backbone. Target speaker embeddings are fused together with input text representations at the fusion encoder, and re-injected into the mel decoder for further speaker preservation.

and a linear layer. The resulting embedding is then utilized by generator \mathcal{E}_{mel} , which has the same structure as \mathcal{E}_{text} , to produce \hat{M}_n corresponding to X_n .

Given the imperative to generate audio in the distinctive timbre of the intended speaker for multi-speaker TTS, the singular textual requisite for single-speaker TTS shifts to dual inputs (i.e., reference audio and text utterance to control speaker voice and the uttered content, respectively). Thus, in conjunction with the original text-to-audio training flow, we propose and incorporate a novel concurrent speaker-specific pipeline that estimates disentangled speaker-specific attributes from reference audio Y_n in a zero-shot manner. This is explained in the subsequent sections.

Negated Feature Extraction

Unlike prior methodologies that rely on a single encoder block to directly learn speaker features (Kang, Min, and Hwang 2023) and maximize the distance between speaker and content features (Hu et al. 2020; Paul et al. 2021), we

introduce a novel negation strategy that entails *subtracting* content information from a comprehensive audio representation, which results in disentangled content-information-free (CIF) speaker-specific embeddings (Fig. 1-(a)). While this arithmetic property has been validated in computer vision research for facial expression learning (Zhang et al. 2021), this is the first time in which negation has been studied for zero-shot multi-speaker TTS. Thus, our negation strategy requires two distinct components: full representation of the reference audio, and the content embeddings extracted from the same acoustic input. In order to generate the total audio representation, we use a pretrained audio compression EnCodec model (Défossez et al. 2022), as it is capable of compressing audio without loss of information into a compact representation¹. The reference audio undergoes a series of transformations, including passing through a 1D convolutional layer with a kernel size of 7, and then four convolution blocks. Each convolution block consists of non-linear ELU activations and a strided convolution with a kernel size twice that of the stride. Strides of sizes 2, 4, 5, and 8 are employed.

Simultaneously, to extract content information, we employ instance normalization (IN) - a well-proven method in stylistic transfer in images (Huang and Belongie 2017), which normalizes artistic styles (e.g., textures, brushstrokes, etc.) so that only the underlying image content is preserved. Similarly, since timbre can be thought of as the *style* applied to the text content, we implement IN in our content extractor to extract content information from the reference audio. This is done by independently calculating the mean μ_c and standard deviation σ_c of each W -dim channel c of the feature map F , which is generated from the previous convolution layer for each reference sample. μ_c and σ_c are then used to normalize each element of F_{cw} . To maintain numerical stability, a small value ϵ is added as seen in Equation 2. IN can be mathematically formulated as follows:

$$\mu_c(M) = \frac{1}{W} \sum_{w=1}^W F_{cw} \quad (1)$$

$$\sigma_c(F) = \sqrt{\frac{1}{W} \sum_{w=1}^W (F_{cw} - \mu_c(F))^2 + \epsilon} \quad (2)$$

$$IN(F) = \frac{F_{cw} - \mu_c(F)}{\sigma_c(F)} \quad (3)$$

Overall, the content extractor incorporates a convolutional bank, whose output is then passed onto two blocks of twice repeating series of convolution 1D, IN, and ReLU. Additional information in the form of average pooling (Chou, Yeh, and Lee 2022) is introduced between the two blocks.

Multi-stream Transformers

In order to achieve accurate speech synthesis in any speaker’s voice, it is crucial to capture various timbre patterns inherent in the input acoustic speaker prompt. However, existing approaches in multi-speaker TTS solely rely

¹We used the official 24kHz EnCodec model from <https://github.com/facebookresearch/encodec>




Figure 2: Comparative visualizations of distinct speaker embeddings. From left to right are baselines Valle et al., Min et al., and our proposed model.

Model	Speaker	MCD (\downarrow)	MOS (\uparrow)	WER (\downarrow)
Valle et al.	Seen	6.60	3.34 ± 0.20	21.82%
	Unseen	8.97	3.25 ± 0.19	26.76%
Min et al.	Seen	5.97	3.56 ± 0.19	17.21%
	Unseen	9.52	3.51 ± 0.19	18.34%
Ours	Seen	5.59	3.86 ± 0.18	15.37%
	Unseen	8.57	3.84 ± 0.18	16.14%

Table 1: Main results for baselines and proposed models. Better performance is determined through lower MCD and WER, and higher MOS scores.

on a single linear information pathway to generate just one speaker embedding, which restricts the model from leveraging alternative hypotheses that capture different perspectives of the target speaker’s timbre. Drawing inspiration from ensemble models (Sagi and Rokach 2018), which combine multiple hypotheses together to enhance predictive performance, we utilize a multi-stream Transformer (Burtsev and Rumshisky 2021) architecture originally used for machine translation (Fig. 1-(b)).

In the previous subsection, a novel negation strategy generated a CIF speaker representation. This representation is subsequently fed into a standard Transformer block to effectively capture long-range dependencies, and then passed to the multi-stream Transformer. Unlike the conventional Transformer (Vaswani et al. 2017), our design features two separate and parallel encoder streams. Each stream conducts independent computations on the full CIF embedding, enabling the exploration and preservation of distinct timbre hypotheses. Furthermore, a Transformer stream employs identical blocks as those found in the original Transformer encoder. All blocks feature 2 heads for the multi-head attention and is of depth 1, ensuring that a single Transformer encoder block is used for each parallel stream. The resulting speaker embedding has a dimension of 128. While the construction of these multiple alternative perspectives is essential, leveraging these hypotheses effectively is another important aspect for improving timbre predictions. Towards this end, we conduct attention pooling across all tokens occupying the same position within each independent Transformer stream, facilitating integration of information from multiple perspectives.




Figure 3: ABX comparisons between generated audio from different models. Best seen in color.

Speaker and Text Fusion

Effective injection and fusion of the previously generated speaker embedding into the original TTS pipeline is crucial, as input text must be spoken in the target speaker’s voice. Conventionally, the speaker embedding is concatenated with either the linguistic encoder outputs or the mel-spectrogram generator inputs of the TTS backbone. However, simple concatenation of target speaker information with target input text may not be enough to sufficiently fuse the two components together.

To address this limitation, we employ Style-Adaptive Layer Normalizations (SALN; (Min et al. 2021)) that involves two inputs: speaker embedding s , and hidden vector i with dimensionality H obtained from the linguistic encoder. First, i is normalized using mean and standard deviation calculations, yielding normalized output x . Subsequently, a fully connected layer predicts the gain g and bias b with regards to the speaker embedding s (Fig. 1-(c)). This process is summarized as follows:

$$SALN(i, s) = g(s) \cdot x + b(s) \quad (4)$$

Since the target text x is adaptively adjusted in accordance with target speaker s , this facilitates a more effective fusion of the input text with the speaker’s distinctive traits, ulti-




Figure 4: Mel-spectrogram comparisons between Valle et al., Min et al., and the proposed system (from top to bottom). Regions of notable disparities are highlighted in boxes. Notably, the uppermost sub-figure accentuates the instance of erroneously inserted enunciation. In addition, compared to the other models, voiceless alveolar fricative /s is correctly generated in the last sub-figure.

mately leading to an improved quality of voice generation.

Loss Function

Our model is trained end-to-end in an unsupervised manner using a summation of different losses. Specifically, L1 distance is used to minimize the difference between the ground truth M_n and synthesized mel-spectrogram \hat{M}_n representations. In addition, we use L2 losses between the ground truth and predicted pitch (P), energy (E), and duration (D) information, which are introduced via the variance adapter. The full reconstruction loss is denoted as follows:

$$\mathcal{L}_{recon} = \mathcal{L}_{mel}(M_n, \hat{M}_n) + \mathcal{L}_{pitch}(P, \hat{P}) + \mathcal{L}_{energy}(E, \hat{E}) + \mathcal{L}_{dur}(D, \hat{D}) \quad (5)$$

Experimental Settings

Dataset

The benchmark LibriTTS train-clean-100 dataset (Zen et al. 2019) is used to conduct training and validation. This dataset is comprised of 53 hours of speech contributed by 247 people, with nearly equal representation from both genders (123 female and 124 male speakers). All audios are resampled at 22050 Hz and 16 bits, and mel-spectrograms are extracted with short-time Fourier transform (Griffin and Lim 1984) using filter length of 1024, hop size of 256, window size

	MCD (↓)	MOS (↑)	WER (↓)
Without Negation	7.44	3.30 ± 0.23	16.68%
With Negation	7.08	3.85 ± 0.12	15.53%

Table 2: Experimental results to verify the efficacy of generating CIF speaker embeddings with negated learning. Seen and unseen speaker scores are averaged across all metrics.

	MCD (↓)	MOS (↑)	WER (↓)
Wav2Vec Layer 1	7.51	3.43 ± 0.15	16.42%
Wav2Vec Layer 7	7.80	3.12 ± 0.19	16.79%
Wav2Vec Layer 12	7.77	3.02 ± 0.21	16.65%
Instance Norm	7.08	3.85 ± 0.12	15.56%

Table 3: Ablation experiments and comparisons between various content extractor models.

of 1024, and 80 frequency bins. Prior to conducting experiments with this dataset, we applied a preprocessing step using the open-source Librosa library², which involved trimming out any initial and final silences present in each audio sample. Additionally, we utilized the G2P library³ to convert the grapheme-based text into phonemes.

Evaluations

We compare our model with two well-established baseline models, Mellotron (Valle et al. 2020) and Meta-StyleSpeech (Min et al. 2021). Mellotron is based on the Tacotron2 (Shen et al. 2018) model, which is augmented with GSTs, speaker embeddings, and pitch contours to perform multi-speaker voice synthesis. Meta-StyleSpeech adopts meta-learning in order to better adapt to unseen speakers, thereby simulating episodic training. In order to assess our model in terms of these two baseline models, we employ both subjective and objective metrics to assess our methodology.

The Mean Opinion Score (MOS) is a subjective evaluation metric in which human participants rate the synthesized audio’s similarity to the ground truth audio using a Likert scale ranging from 1 to 5. A score of 5 indicates the highest degree of speaker similarity between the synthesized audio and the ground truth audio. For MOS evaluation, we utilize the Amazon Mechanical Turk⁴ platform, where participants compare the speaker similarity of an anonymized audio that is synthesized from one of three models (i.e., two baseline models and our proposed model) to a specified ground truth audio. Each participant answered 53 comparison questions, with an average participating time of approximately 40 minutes. Final MOS scores are calculated with a 95% confidence interval. In addition, as a complement to model-specific MOS validations, we further conducted intra-model A or B tests (ABX) in which the same pool of evaluators were presented with three audio files, each synthesized from a different model. They were then asked to choose the audio file that best matched the speaker’s voice in comparison to the reference ground truth audio.

²<https://librosa.org/doc/latest/index.html>

³<https://github.com/Kyubyong/g2p>

⁴<https://www.mturk.com/>

Stream	Fusion	MCD (\downarrow)	MOS (\uparrow)	WER (\downarrow)
Single	N/A	7.52	3.49 \pm 0.18	15.65%
Multi	Concat	7.31	3.51 \pm 0.27	15.84%
	Attn. Pool	7.08	3.85 \pm 0.12	15.53%

Table 4: Single and multi-stream comparisons with different hypotheses fusion methods.

The objective metrics that are employed in our evaluations are Mel Cepstral Distortion (MCD) and Word Error Rate (WER). MCD quantifies the difference between the k th Mel Frequency Cepstral Coefficient (MFCC) vectors of the synthesized \hat{S} and the ground truth audio signals S for frame T (Equation 6). Specifically, we use the dynamic time warping (DTW) version to take differences in audio durations into account. Smaller MCD scores indicates that the synthesized audio is more similar overall with the ground truth audio. We utilize the open source pymcd library⁵.

$$MCD = \frac{10\sqrt{2}}{\ln 10} \frac{1}{T} \sum_t \sqrt{\sum_{k=1}^K (S_{t,k} - \hat{S}_{t,k})^2} \quad (6)$$

Word Error Rate (WER) is a widely utilized metric in Automatic Speech Recognition (ASR) to evaluate the accuracy of word pronunciation in synthesized audio. While not typically used in TTS research, we include this metric in our evaluations in order to examine the potential amount of content information that has leaked into speaker representations during the acquisition of speaker features from the reference audio. This is feasible since lower WER rates are indicative of improved pronunciation, thereby implying reduced content leakage within our acquired speaker representations. To compute WER, we leverage OpenAI’s pretrained Whisper large-v2 model (Radford et al. 2022) to transcribe synthetic audios generated from different models, and subsequently employ the jiwer package⁶ to precisely calculate the error rates.

Training Details

Adam optimization is utilized with hyperparameters $\beta_1 = 0.9$, $\beta_2 = 0.98$, and $\epsilon = 10^{-9}$. To convert the generated mel-spectrograms into audio, we employ the HiFi-GAN (Kong, Kim, and Bae 2020) vocoder. All experiments are conducted on a single RTX A6000 GPU with a batch size of 16, until step 300,000. On average, an experiment takes approximately 28 hours to complete. The total number of parameters of our proposed model is 30M.

Results and Discussions

The results in Table 1 unequivocally demonstrate the superiority of our model, which surpasses the baseline systems across all three metrics for both seen and unseen speakers. Moreover, we provide visual evidence of the effectiveness of our negated feature learning process by conducting uniform manifold projections (McInnes, Healy, and Melville

⁵<https://github.com/chenqi008/pymcd>

⁶<https://github.com/jitsi/jiwer>

Model	MCD (\downarrow)			
Valle et al.	6.60			
Min et al.	5.97			
Proposed		H2	H4	H8
	D1	5.59	5.71	5.64
	D2	5.61	5.71	5.74
	D4	5.69	5.77	5.64

Table 5: Multi-stream Transformer optimizations for seen speakers. The numbers next to head H and block depth D denote the precise parameters used.

Model	MCD (\downarrow)			
Valle et al.	8.97			
Min et al.	9.52			
Proposed		H2	H4	H8
	D1	8.57	8.85	8.70
	D2	8.70	8.67	8.80
	D4	8.84	8.82	8.95

Table 6: Multi-stream Transformer optimizations for unseen speakers.

2018) of the final speaker embeddings (Fig. 2). These visualizations clearly depict the enhanced clustering of utterances from the same speakers, validating our model’s ability to better capture distinct speaker-specific representations, which ultimately aligns with higher speaker fidelity (MOS).

While our model exhibits the highest MOS scores, evaluations were done separately for each model. Consequently, to achieve a more comprehensive assessment, we conduct inter-model comparisons through ABX across four possible conditions involving speaker visibility and parallelism of text utterances (Fig. 3). We find that while our proposed method generally has higher preference, clear model preferences manifest when both ground truth and synthesized audio share the same spoken context. Notably, the variance in preference between the two leading models diminishes from 12.50% to 6.25%, and from 18.75% to 0% for seen and unseen speaker conditions, respectively. In addition, we show in Figure 4 the acoustic fluctuations that occur across different models, which influences MCD and WER scores.

Impact of Negated Feature Generation

Disentangled CIF speaker embeddings were achieved by negating content information from the full acoustic representation of the reference audio. To assess the impact of negation, we trained a separate model that bypassed this process and directly fed the input reference audio into the multi-stream Transformer architecture. However, as seen in Table 2, this yields lower voice similarity (MOS) and overall audio fidelity (MCD). Moreover, significantly increased WER infers increased content leakage in speaker representations, which leads to pronunciation instability due to conflicting contextual information from the reference audio and the actual target text. Hence, negation plays a crucial role in generating accurate CIF embeddings.

While the effectiveness of negation has been established, the adequacy of instance normalization in extracting con-




Figure 5: Long-tail distributions of the audio durations in the training dataset. The average duration for approximately 26,500 audio samples is 5.37 seconds (indicated in dark blue). Minimum and maximum audio lengths are 0 and 33 seconds, respectively.

tent information remains to be validated. To address this, we explore an alternative approach using a pretrained wav2vec 2.0 base model (Baevski et al. 2020) as the context extractor. Specifically, embeddings outputted from the 1st, 7th, and 12th layers are used since each layer of wav2vec encapsulate different speech properties such as local acoustic features and phonetic content (Pasad, Chou, and Livescu 2021). Table 3 reveals that using the first layer of wav2vec is most effective. Nevertheless, our method of using instance normalization proves more effective than the different variants of the pretrained wav2vec model.

Utilization of Multiple Hypotheses

To assess the efficacy of multi-stream Transformers, we train three separate models. The first model incorporates a linear single-stream pathway to produce a single hypothesis. Recognizing the potential effect that multi-hypotheses fusion methods may have on performance, for fairer comparison, we developed a second model. This model employs multi-stream Transformers, but uses naïve concatenation to combine multiple hypotheses together to create the ultimate speaker embedding. A comparison between these two models demonstrates that even with naïve concatenation, multi-stream Transformers tends to surpass single-stream model performance (Table 4). Furthermore, to determine the most effective way to fuse multiple hypotheses, we compare the second model, which uses concatenation, with a third model (i.e., our proposed system) that applies attention pooling. Our findings reveal that attention pooling proves to be a more effective approach compared to simple concatenation.

We further conducted optimization experiments for the multi-stream Transformer component of our scheme by exploring various head settings, as well as the number of blocks used within each parallel stream (Tables 5, 6). Interestingly, we discovered that the best performance was achieved with minimal head numbers and blocks. This observation could be attributed to the characteristics of our training dataset. As illustrated in Fig. 5, a significant por-

	MCD (\downarrow)	MOS (\uparrow)	WER (\downarrow)
Enc. Only	8.66	3.00 ± 0.19	17.47%
Dec. Only	8.59	2.92 ± 0.19	17.56%
Enc. & Dec.	7.08	3.85 ± 0.12	15.53%

Table 7: Speaker knowledge integration experiments. Optimal performance is observed when text and speaker are combined together at the fusion encoder, and speaker representations are additionally infused in the mel decoder.

portion (53.57%) of the dataset have durations less than 5 seconds. Given the brevity of the audios and the fact that multi-stream Transformers conduct computations on a CIF speaker embedding that has already undergone content information pruning, it becomes more effective to adopt a smaller number of heads and blocks. Despite the increase in MCD scores when increasing the number of heads and/or blocks, it is important to note that all scores, both for seen and unseen speakers, outperform the baseline models.

Speaker Knowledge Infusion

Strategic placement and frequency of fusing speaker knowledge within the basic TTS pipeline are crucial considerations. Thus, we additionally experiment with injecting speaker representations either exclusively into the TTS encoder or decoder. Our findings reveal that the most effective method is to incorporate speaker knowledge into both the encoder and decoder (Table 7). Also notable is that when speaker information is only integrated at the decoder, MCD slightly decreases, but MOS and WER scores are worse than encoder-exclusive fusion. This indicates that early fusion of text and speaker attributes is more important.

Conclusion

In this paper, we proposed for the first time, a novel negation feature learning scheme that could effectively extract speaker representations from a reference audio with less content leakage. Furthermore, we improve speaker fidelity through multi-stream Transformers by generating multiple hypotheses, which are all taken into account with attention pooling to generate the final speaker embedding. Through adaptive layer normalizations, target speaker attributes were fused into both the encoder and decoder of the baseline TTS model. Experiments have proven that our negated feature generation and learning methodology outperforms other well-known baseline models in terms of zero-shot speaker adaptation.

Acknowledgments

This research is supported by Culture Technology R&D Program through the Korea Creative Content Agency funded by Ministry of Culture, Sports and Tourism (Development of contents meta-verse based on XR and AI, R2021040136) and by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2019-0-01906, Artificial Intelligence Graduate School Program(POSTECH)).

References

- An, X.; Wang, Y.; Yang, S.; Ma, Z.; and Xie, L. 2019. Learning Hierarchical Representations for Expressive Speaking Style in End-to-End Speech Synthesis. In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 184–191.
- Baevski, A.; Zhou, Y.; Mohamed, A.; and Auli, M. 2020. wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, 12449–12460. Curran Associates, Inc.
- Burtsev, M.; and Rumshisky, A. 2021. Multi-Stream Transformers. In *arXiv preprint arXiv:2107.10342*.
- Chen, L.-W.; and Rudnicky, A. 2022. Fine-Grained Style Control In Transformer-Based Text-To-Speech Synthesis. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 7907–7911.
- Chen, M.; Tan, X.; Li, B.; Liu, Y.; Zhao, S.; and Liu, T.-Y. 2021. AdaSpeech: Adaptive Text to Speech for Custom Voice. In *International Conference on Learning Representations (ICLR)*.
- Chen, Y.; Assael, Y.; Shillingford, B.; Budden, D.; Reed, S.; Zen, H.; Wang, Q.; Cobo, L. C.; Trask, A.; Laurie, B.; Gulcehre, C.; van den Oord, A.; Vinyals, O.; and de Freitas, N. 2019. Sample Efficient Adaptive Text-to-Speech. In *International Conference on Learning Representations (ICLR)*.
- Choi, B. J.; Jeong, M.; Lee, J. Y.; and Kim, N. S. 2022. SNAC: Speaker-Normalized Affine Coupling Layer in Flow-Based Architecture for Zero-Shot Multi-Speaker Text-to-Speech. *IEEE Signal Processing Letters*, 29: 2502–2506.
- Chou, J.-c.; Yeh, C.-c.; and Lee, H.-y. 2022. One-shot Voice Conversion by Separating Speaker and Content Representations with Instance Normalization. In *ArXiv e-prints*.
- Dumoulin, V.; Shlens, J.; and Kudlur, M. 2017. A Learned Representation For Artistic Style. *ICLR*.
- Défossez, A.; Copet, J.; Synnaeve, G.; and Adi, Y. 2022. High Fidelity Neural Audio Compression. In *arXiv preprint arXiv:2210.13438*.
- Gatys, L. A.; Ecker, A. S.; and Bethge, M. 2016. Image Style Transfer Using Convolutional Neural Networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2414–2423.
- Griffin, D.; and Lim, J. 1984. Signal estimation from modified short-time Fourier transform. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 32(2): 236–243.
- Helander, E.; Schwarz, J.; Nurminen, J.; Silén, H.; and Gabouj, M. 2008. On the impact of alignment on voice conversion performance. In *Interspeech*.
- Hodari, Z.; Watts, O.; and King, S. 2019. Using generative modelling to produce varied intonation for speech synthesis. *10th ISCA Speech Synthesis Workshop (SSW10)*, abs/1906.04233.
- Hsu, C.-C.; Hwang, H.-T.; Wu, Y.-C.; Tsao, Y.; and Wang, H.-M. 2016. Voice conversion from non-parallel corpora using variational auto-encoder. In *2016 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA)*, 1–6.
- Hu, T.-Y.; Shrivastava, A.; Tuzel, O.; and Dhir, C. 2020. Unsupervised Style and Content Separation by Minimizing Mutual Information for Speech Synthesis. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 3267–3271.
- Huang, X.; and Belongie, S. 2017. Arbitrary Style Transfer in Real-time with Adaptive Instance Normalization. In *International Conference on Computer Vision (ICCV)*.
- Jia, Y.; Zhang, Y.; Weiss, R. J.; Wang, Q.; Shen, J.; Ren, F.; Chen, Z.; Nguyen, P.; Pang, R.; Moreno, I. L.; and Wu, Y. 2018. Transfer learning from speaker verification to multi-speaker text-to-speech synthesis. In *International Conference on Neural Information Processing Systems (NeurIPS)*.
- Kameoka, H.; Kaneko, T.; Tanaka, K.; and Hojo, N. 2018. StarGAN-VC: Non-parallel many-to-many voice conversion with star generative adversarial networks. In *IEEE Spoken Language Technology Workshop (SLT)*.
- Kameoka, H.; Kaneko, T.; Tanaka, K.; and Hojo, N. 2019. ACVAE-VC: Non-parallel voice conversion with auxiliary classifier variational autoencoder. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(9): 1432–1443.
- Kaneko, T.; and Kameoka, H. 2018. CycleGAN-VC: Non-parallel Voice Conversion Using Cycle-Consistent Adversarial Networks. In *26th European Signal Processing Conference (EUSIPCO)*.
- Kang, M.; Min, D.; and Hwang, S. J. 2023. GradStyleSpeech: Any-Speaker Adaptive Text-to-Speech Synthesis with Diffusion Models. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- Kong, J.; Kim, J.; and Bae, J. 2020. HiFi-GAN: Generative Adversarial Networks for Efficient and High Fidelity Speech Synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Kulkarni, A.; Colotte, V.; and Jouviet, D. 2022. Multi-stage attention for fine-grained expressivity transfer in multispeaker text-to-speech system. In *2022 30th European Signal Processing Conference (EUSIPCO)*, 180–184.
- Lei, Y.; Yang, S.; Wang, X.; and Xie, L. 2022. MsEmoTTS: Multi-Scale Emotion Transfer, Prediction, and Control for Emotional Speech Synthesis. *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, 30: 853–864.
- McInnes, L.; Healy, J.; and Melville, J. 2018. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. In *ArXiv e-prints*.
- Min, D.; Lee, D. B.; Yang, E.; and Hwang, S. J. 2021. MetaStyleSpeech : Multi-speaker adaptive text-to-speech generation. In *International Conference on Learning Representations (ICLR)*.
- Pasad, A.; Chou, J.-C.; and Livescu, K. 2021. Layer-wise Analysis of a Self-supervised Speech Representation Model.

- In *Workshop on Automatic Speech Recognition and Understanding (ASRU)*.
- Paul, D.; Mukherjee, S.; Pantazis, Y.; and Stylianou, Y. 2021. A Universal Multi-Speaker Multi-Style Text-to-Speech via Disentangled Representation Learning Based on Rényi Divergence Minimization. In *Proc. Interspeech 2021*, 3625–3629.
- Radford, A.; Kim, J. W.; Xu, T.; Brockman, G.; McLeavey, C.; and Sutskever, I. 2022. Robust Speech Recognition via Large-Scale Weak Supervision. In *ArXiv e-prints*.
- Ren, Y.; Hu, C.; Tan, X.; Qin, T.; Zhao, S.; Zhao, Z.; and Liu, T.-Y. 2021. FastSpeech 2: Fast and High-Quality End-to-End Text to Speech. In *International Conference on Learning Representations (ICLR)*.
- Sagi, O.; and Rokach, L. 2018. Ensemble learning: A survey. *WIREs Data Mining and Knowledge Discovery*, 8(4): e1249.
- Shen, J.; Pang, R.; Weiss, R. J.; Schuster, M.; Jaitly, N.; Yang, Z.; Chen, Z.; Zhang, Y.; Wang, Y.; Skerry-Ryan, R.; Saurous, R. A.; Agiomyriannakis, Y.; and Wu, Y. 2018. Natural tts synthesis by conditioning wavenet on mel spectrogram predictions. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4779 – 4783.
- Ulyanov, D.; Vedaldi, A.; and Lempitsky, V. 2017. Improved Texture Networks: Maximizing Quality and Diversity in Feed-Forward Stylization and Texture Synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Valle, R.; Li, J.; Prenger, R.; and Catanzaro, B. 2020. Melotron: Multispeaker expressive voice synthesis by conditioning on rhythm, pitch and global style tokens. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention Is All You Need. In *International Conference on Neural Information Processing Systems (NeurIPS)*.
- Wang, Y.; Hoang, C.; and Federico, M. 2021. Towards Modeling the Style of Translators in Neural Machine Translation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1193–1199. Online: Association for Computational Linguistics.
- Zen, H.; Dang, V.; Clark, R.; Zhang, Y.; Weiss, R. J.; Jia, Y.; Chen, Z.; and Wu, Y. 2019. LibriTTS: A Corpus Derived from LibriSpeech for Text-to-Speech. In *Conference of the International Speech Communication Association (Interspeech)*.
- Zhang, W.; Ji, X.; Chen, K.; Ding, Y.; and Fan, C. 2021. Learning a Facial Expression Embedding Disentangled From Identity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6759–6768.
- Zhang, Y.; Wang, Z.; Yang, P.; Sun, H.; Wang, Z.; and Xie, L. 2022. AccentSpeech: Learning Accent from Crowdsourced Data for Target Speaker TTS with Accents. In *2022 13th International Symposium on Chinese Spoken Language Processing (ISCSLP)*, 76–80.
- Zhao, B.; Zhang, X.; Wang, J.; Cheng, N.; and Xiao, J. 2022. nnSpeech: Speaker-Guided Conditional Variational Autoencoder for Zero-Shot Multi-speaker text-to-speech. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4293–4297.