



Automatic Tuning of Loss Trade-offs without Hyper-parameter Search in End-to-End Zero-Shot Speech Synthesis

Seongyeon Park¹, Bohyung Kim¹, Tae-hyun Oh^{2,3}

¹CNAI, Korea

²Institute for Convergence Research and Education in Advanced Technology, Yonsei University, Korea

³Dept. of EE and GSAI, POSTECH, Korea.

eaglesy221@snu.ac.kr

Abstract

Recently, zero-shot TTS and VC methods have gained attention due to their practicality of being able to generate voices even unseen during training. Among these methods, zero-shot modifications of the VITS model have shown superior performance, while having useful properties inherited from VITS. However, the performance of VITS and VITS-based zero-shot models vary dramatically depending on how the losses are balanced. This can be problematic, as it requires a burdensome procedure of tuning loss balance hyper-parameters to find the optimal balance. In this work, we propose a novel framework that finds this optimum without search, by inducing the decoder of VITS-based models to its full reconstruction ability. With our framework, we show superior performance compared to baselines in zero-shot TTS and VC, achieving state-of-the-art performance. Furthermore, we show the robustness of our framework in various settings. We provide an explanation for the results in the discussion.

Index Terms: Zero-shot, Voice Conversion, Text-to-speech, Speech Synthesis, Efficient Optimum Discovery

1. Introduction

The recent advance of neural networks has enhanced the quality of speech synthesis models in the areas of Text-to-Speech (TTS) and Voice Conversion (VC), resulting in the production of realistic and naturally sounding speech. Among multiple categories of TTS and VC, zero-shot approaches, *e.g.*, [1–11], have obtained much attention. These approaches are practical, because they can synthesize speech of given voices even unseen during training.

While models for zero-shot VC and TTS were developed separately, Glow-TTS [12] enables VC and multi-speaker TTS with a single model, utilizing the invertibility of Normalizing Flows (NFs) [13]. Building upon this, VITS [14] extends it with Variational Autoencoder (VAE) [15] structure where its decoder is HiFi-GAN [16]. By jointly training its components, VITS is not only practical in terms that it works in a single stage (*i.e.*, can predict waveform directly), but can also produce ground truth level fidelity speech of multiple speakers.

By virtue of these favorable properties, there have been attempts [17–20] to extend VITS or Glow-TTS to zero-shot VC and TTS. They use speaker encoders instead of the speaker embedding table in VITS or Glow-TTS, so that the models can extract and synthesize diverse unseen voices in the inference phase, *i.e.*, enabling zero-shot. While extending the useful properties of

Acknowledgements: T.-H. Oh was supported by IITP grants funded by the Korea government(MSIT) (No.2021-0-02068, Artificial Intelligence Innovation Hub; No.2022-0-00124, Development of Artificial Intelligence Technology for Self-Improving Competency-Aware Learning Capabilities).




Figure 1: Word Error Rate of speech synthesized through VC or TTS, from VITS and YourTTS [20] according to the loss balance hyper-parameter α (the loss weight parameter of the reconstruction loss). Both axes are in log scale. The WER has a noticeable variance depending on the balance hyper-parameter. Interestingly, the results obtained with the parameter suggested by the original authors show a notable gap from the best points we found. These graphs suggest that there is likely to exist an optimal value of the hyper-parameter α . One might exhaustively tune balance hyper-parameters for better performance. In this work, we propose a framework that alleviates the exhaustive hyper-parameter tuning process, and instead enables a near-optimal trade-off selection without search.

Glow-TTS and/or VITS, these zero-shot models reported superior performance over previous zero-shot VC and TTS models. In particular, VITS-based models [19, 20] show superior performance compared to Glow-TTS-based models [17, 18], as the models can internally optimize the latent representations instead of using human-designed features such as mel-spectrograms.

However, we found that VITS and zero-shot modifications of VITS tend to be sensitive to how losses are balanced. Specifically, performance depending on the balance hyper-parameter for reconstruction loss can be seen in Figure 1. A naive approach to finding the optimal point would involve searching and tuning the loss balance hyper-parameters. In this work, we propose a framework that alleviates the exhaustive tuning process and instead enables a near-optimal trade-off selection without search.

Our proposed method is built on our hypothesis that the converged reconstruction loss value of separately training the HiFi-GAN vocoder would provide a favorable reference for the reconstruction loss in training VITS-based models. More specifically, as a preliminary task, we first train just HiFi-GAN as a mel-spectrogram-to-waveform vocoder. Denoting the converged reconstruction loss value as ϵ^* , we obtain this value from the preliminary task. Then, we train VITS-based models for the main speech synthesis task, such that its reconstruction loss is

converged at ε^* , which would be sufficient to get high fidelity audios even in the main task. We realize this learning process by the Modified Differential Multiplier Method (MDMM) [21].

With this method, our framework solves the pre-mentioned trade-off problem in VITS-based models without hyper-parameter tuning. We experimentally show the superior performance achieved by this framework in zero-shot VC and TTS, achieving state-of-the-art performance. We also show its robustness across various model and audio configurations. We further explain the results in the discussion. Source code and audio samples are available at: https://github.com/cnaigitHub/Auto_Tuning_Zeroshot_TTS_and_VC

2. Method

In this section, we first describe the base models used in this work. We explain VITS [14], the state-of-the-art model for multi-speaker speech synthesis, and introduce a simple variant of VITS that enables zero-shot synthesis. Then we describe our framework and how to apply it to the base models.

2.1. Base Model for Speech Synthesis

VITS basically integrates Glow-TTS [12] and HiFi-GAN [16].

Glow-TTS is a Normalizing Flow (NF) [13]-based TTS model, which uses the monotonic alignment search algorithm to obtain a hard attention alignment between input text and target mel-spectrogram timesteps. The original Glow-TTS is designed to conduct text-to-mel-spectrogram conversion, *i.e.*, TTS.

HiFi-GAN is a waveform synthesizer, which typically decodes mel-spectrograms into speech waveforms via adversarial training. Instead of mel-spectrograms as input, the HiFi-GAN used as the decoder in VITS decodes latent features obtained from Glow-TTS and directly converts to waveforms, which enables achieving high fidelity audio synthesis.

With these modules, VITS integrates them and adds the VAE structure. The latent representation of this VAE follows a posterior distribution, encoded from linear spectrograms of speech audio. Also, a text-conditioned prior distribution is produced by the Glow-TTS module. These posterior and prior are regularized by the KL divergence during training. Noticeably, the invertibility of NF in VITS allows latent features to go forward and backward, which enables performing both VC and TTS in a single model. In addition, to achieve multi-speaker speech synthesis, VITS adopts a trainable speaker embedding table. The speaker embeddings are injected to condition the posterior encoder, the NF layer, and the HiFi-GAN decoder.

Training of VITS. VITS employs adversarial training in the waveform domain by training both a generator and discriminator network. The generator and discriminator are optimized by the following respective losses,

$$\begin{aligned} L_g &= \alpha L_{recon} + L_{KL} + L_{dur} + L_{adv}(G; D) + L_{fm}(G; D), \\ L_d &= L_{adv}(D; G), \end{aligned} \quad (1)$$

where the reconstruction loss term L_{recon} is defined by an L_1 loss between the mel-spectrogram of the generated waveform and ground truth waveform, L_{adv} is the adversarial loss, and L_{KL} is the KL divergence between the posterior obtained from audio and the prior obtained from text. For detail on the other losses L_{dur} and L_{fm} , refer to the literature [14]. The balance hyper-parameter α weights the reconstruction loss against other losses, and is the same one used in Figure 1. As the model outputs waveform, a differentiable STFT and linear projection

to the mel-scale is required to back-propagate from L_{recon} [14].

Note that, when training HiFi-GAN as a mel-spectrogram-to-waveform vocoder, it uses the same losses for its generator and discriminator as Equation (1), except for excluding L_{KL} and L_{dur} . This enables a direct comparison of L_{recon} values in HiFi-GAN training with the one in VITS training despite noticeable differences in modeling and tasks.

Zero-shot Modification of VITS. We modify and extend the original multi-speaker VITS model with simple changes in order to enable end-to-end zero-shot speech synthesis. To prevent performance drop on non-English datasets, we use a jointly trained speaker encoder instead of a pre-trained one. We replace the speaker embedding table of VITS with a speaker encoder consisting of convolutional layers. The architecture of the speaker encoder is similar to the one in [19], except that our speaker encoder takes linear spectrograms as the input source. Second, we do not condition the posterior encoder by speaker embedding. This is meant to prevent performance drop on voice conversion with unseen source speakers. Third, we use 10 transformer blocks for the text encoder following YourTTS [20]. Fourth, we use the Deterministic Duration Predictor (DDP) instead of the Stochastic Duration Predictor (SDP). This is due to reports on the instability of SDP in [20, 22]. For simplicity, we call this modified model *Zero-shot VITS*. Through experiments, we show that this model shows comparable performance to YourTTS [20] (a baseline zero-shot modification of VITS) without our framework applied, but outperforms it when our framework is applied.

2.2. Loss Value of Separate Vocoder Training as a Guide

As shared in all VAE-based methods, both VITS and zero-shot modifications of VITS show trade-offs between their reconstruction loss and the other losses. By changing the balance ratio of these losses, these models converge on different Pareto equilibria. As mentioned before, model performance varies drastically depending on which Pareto equilibrium it converged (Figure 1). A vanilla method to find the optimum would be to conduct exhaustive iterative experiments with different balance parameters, which is computationally burdensome. Instead of directly attempting different balance hyper-parameters α , we tune a specific target value of L_{recon} to find an optimal balance.

The motivation of this approach is that the decoder of a VITS-based model is a HiFi-GAN; thus, if HiFi-GAN works well for its own vocoder task, the level of the converged value of L_{recon} may indicate the desired quality and may be transferred to other tasks. To test this hypothesis, we separately train HiFi-GAN for a vocoder task alone, we call a preliminary task. Then, we obtain the converged L_{recon} value, denoted as ε^* , and do not use any other results, including the trained vocoder model from the preliminary task. In our framework, our idea is to bring and use ε^* as the target value to converge L_{recon} in VITS training. That is, we propose to use this particular L_{recon} value as the desired point for a trade-off between L_{recon} and other losses.

2.3. Constraint Optimization of VITS to a Specific L_{recon}

Constrained Optimization for Neural Networks. Platt and Barr [21] introduces methods for constrained optimization on neural networks. We briefly explain the Modified Differential Multiplier Method (MDMM) proposed in [21]. Consider a neural network, whose parameters are θ . Suppose we would like to solve the following optimization problem:

$$\min_{\theta} F(\theta) \quad \text{s.t.} \quad G(\theta) = 0 \quad (2)$$

To optimize such problem, the Lagrange multiplier method can be used. To be more specific, the solution θ of Equation (2) should be a critical point of the following Lagrangian function:

$$\mathcal{L}(\theta, \lambda) = F(\theta) + \lambda G(\theta) \quad (3)$$

for some Lagrange multiplier λ . Note that the solution is also a critical point for λ , as $\frac{\partial \mathcal{L}}{\partial \lambda} = G(\theta) = 0$. The Lagrange multiplier λ can be regarded as an additional variable and can be updated via gradient descent or ascent. The MDMM method adds a penalty term for the constraint as:

$$\mathcal{L}(\theta, \lambda) = F(\theta) + \lambda G(\theta) + \frac{c}{2} G(\theta)^2, \quad (4)$$

where c is the damping constant parameter. The MDMM updates θ with gradient descent and λ with gradient ascent over \mathcal{L} . Please refer to more details. [21].

Selecting a VITS' Trade-Off Point with MDMM. We can select a trade-off point of VITS by enforcing the reconstruction loss L_{recon} to a user-chosen constant value ε . To be more specific, we define $G(\theta) = L_{recon}(\theta) - \varepsilon$. The loss $F(\theta)$ is set to be the sum of all the remaining loss terms. Then, we can update θ and λ according to the following gradients:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \lambda} &= G(\theta) = L_{recon}(\theta) - \varepsilon \\ \frac{\partial \mathcal{L}}{\partial \theta} &= \frac{\partial F}{\partial \theta} + \lambda \frac{\partial L_{recon}}{\partial \theta} + c(L_{recon} - \varepsilon) \frac{\partial L_{recon}}{\partial \theta} \end{aligned} \quad (5)$$

By this optimization, the model converges on the specified Pareto equilibrium, where L_{recon} is equal to ε . In our proposed framework, we set $\varepsilon = \varepsilon^*$, but for the sake of proving its empirical optimality, we also report model performance on nearby values, specifically $\varepsilon = \varepsilon^* \pm 0.1$.

3. Experiments, Results and Discussion

3.1. Experiment Settings and Configurations

Evaluation. We use Word Error Rate (WER) / Character Error Rate (CER) of the synthesized speech and Resemblyzer Embedding Cosine Similarity (RECS) [23]¹ between the target speech and synthesized speech. We use a pre-trained ASR model² for obtaining the WER / CER between synthesized speech and the ground truth transcript.

Datasets. We either use the VCTK [24] dataset or a multilingual dataset mostly composed of public data³. When using the VCTK dataset, we hold out 4 male and 4 female speakers from its training set and use them for evaluation on unseen speakers. For evaluation on seen speakers, we randomly hold out 128 utterances from the remaining training set. The rest not used for evaluation is used for training. Evaluation for models trained with the multilingual dataset was conducted using the *test-clean* split of LibriTTS [25], consisting of speakers unseen during training. The rest is used for training.⁴ We use International Phonetic Alphabet as transcription, following VITS [14].

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

²<https://huggingface.co/facebook/hubert-large-ls960-ft>

³The dataset consists of 15 different languages, mostly publicly available data. English [24, 25], Dutch / French / German / Italian / Polish / Portuguese / Spanish [26, 27], Thai / Cantonese [28], Korean [29], Japanese [30, 31], Mandarin [32], Vietnamese [33], and Arabic [34]. We also add some proprietary data containing 4 Korean speakers, which take up less than 0.5% of the entire training set. We exclude any data longer than 13 seconds, and trim silences with librosa [35].

⁴For training *Zero-shot VITS* with the multilingual dataset, we add a language embedding lookup table and condition the model on language embeddings, by following the multilingual training of YourTTS.

Detail Configurations. We downsampled audio to 16kHz, and used linear/mel spectrograms with a hop size of 320 and a window size of 1280. The sampling rate and hop size were selected according to the input and output frequencies of WavLM [36] features, which are used in one of the baselines [11]. We trained each model with a batch size of 32 for 500k steps. We use the AdamW optimizer for our methods, with a learning rate of 2e-4 and a weight decay of 0.01. For the preliminary task, we train HiFi-GAN for 500k steps, as improvement in perceptual quality is small afterward. We used the optimizers and learning rates specified in the original works for the other methods.

3.2. Effectiveness of Our Framework

To see the effectiveness of our framework, we apply our framework to VITS and zero-shot modifications of VITS, and compare them with the state-of-the-arts in zero-shot VC / TTS on VCTK [24]. We denote this experiment as Exp1.

Compared Methods. YourTTS [20] is a VITS-based model, that uses a language embedding table to be trained as a multilingual TTS model. It uses a pre-trained speaker encoder to enable zero-shot synthesis, and uses the speaker consistency loss to improve the similarity of the synthesized and reference speeches. Its VITS-based structure enables both TTS and VC. C-DSVAE [11] is a VAE-based VC model, and has separate encoders for speaker information and content. Due to its fast convergence, we train C-DSVAE for 150k steps. We refer the readers to the original papers [11, 20] for more details. C-DSVAE is a two-stage model, which predicts the mel-spectrogram and thus needs an external vocoder to produce waveform. We use the separate HiFi-GAN vocoder [16] trained on VCTK for 500k steps with our configurations.

Results. The HiFi-GAN trained with our settings converged to $L_{recon} = 0.25$. Thus we use $\varepsilon = \varepsilon^* = 0.25$ for the MDMM optimization framework in Exp1⁵. As shown in Table 1, *Zero-shot VITS* with our framework shows favorable performance in all settings of VC and TTS compared to previously proposed state-of-the-art zero-shot models. Also, VITS and YourTTS' performance is significantly improved by our framework.

3.3. Robustness of Our Framework

We experiment our framework with different models and data configurations to show the robustness of our method. For each experiment, we first train HiFi-GAN vocoders with different conditions to obtain ε^* . Then, we train VITS-based models following each setting with $\varepsilon = \varepsilon^*$ and nearby values $\varepsilon = \varepsilon^* \pm 0.1$. Unless specified otherwise, VCTK is used as the default dataset, and *Zero-shot VITS* as the default model. We train models for 300k steps as metric improvements were small afterward.

- Exp2 (Different dataset): Use the multilingual dataset.
- Exp3 (Decoder with less reconstruction ability): Reduce the internal channel numbers of HiFi-GAN and *Zero-shot VITS*' decoder by a factor of 8.
- Exp4 (Different model with same decoder): Use YourTTS, which has the same decoder as *Zero-shot VITS*.
- Exp5 (Different model with different dataset): Similar to Exp4, but use the multilingual dataset.
- Exp6 (Different audio configuration): Use audio with 22050 Hz sampling rate, and use spectrograms with hop size 256, window size 1024.

Results. As shown in Table 2, $\varepsilon = \varepsilon^*$ shows the best trade-off performance compared to nearby values. In Exp3, although

⁵The VITS-based models used in Exp1, *i.e.* VITS, YourTTS, and *Zero-shot VITS*, have identical decoders, and thus share the same ε^* value.

Table 1: Comparisons between the prior arts and the methods with our framework. Zero-shot VITS combined with our framework shows favorable performance compared to the prior arts without our framework in zero-shot VC and TTS. Further, the performance of VITS and YourTTS is improved significantly when our framework is applied, especially in the VC task. The best results within the same setting are marked in bold.

Method	VC (Seen-to-seen)		VC (Unseen-to-seen)		VC (Unseen-to-unseen)		TTS (Seen)		TTS (Unseen)	
	WER/CER (%) (↓)	RECS (↑)	WER/CER (%)	RECS	WER/CER (%)	RECS	WER/CER (%)	RECS	WER/CER (%)	RECS
GT (VCTK)					6.86 / 2.24	0.803 ± 0.008				
VITS [14]	16.03 / 7.12	0.792 ± 0.010	×	×	×	×	29.87 / 16.21	0.801 ± 0.001	×	×
C-DSVAE [11]	44.05 / 24.80	0.709 ± 0.012	46.89 / 27.20	0.693 ± 0.009	45.23 / 25.69	0.614 ± 0.010	×	×	×	×
YourTTS [20]	16.78 / 7.94	0.814 ± 0.010	26.61 / 13.40	0.791 ± 0.007	24.04 / 11.90	0.734 ± 0.009	37.64 / 19.83	0.817 ± 0.010	36.28 / 20.27	0.751 ± 0.009
Zero-shot VITS (Section 2.1)	20.79 / 10.53	0.815 ± 0.010	29.53 / 15.54	0.804 ± 0.007	26.53 / 13.23	0.744 ± 0.008	28.06 / 13.87	0.818 ± 0.010	27.38 / 14.38	0.747 ± 0.009
VITS with our framework	10.52 / 4.24	0.786 ± 0.011	×	×	×	×	26.22 / 13.22	0.794 ± 0.010	×	×
YourTTS with our framework	11.34 / 5.16	0.811 ± 0.010	14.19 / 6.60	0.768 ± 0.008	13.72 / 6.09	0.730 ± 0.009	33.27 / 17.16	0.812 ± 0.010	32.34 / 17.26	0.755 ± 0.008
Zero-shot VITS with our framework	12.18 / 5.94	0.800 ± 0.010	15.07 / 7.20	0.782 ± 0.007	14.46 / 6.69	0.740 ± 0.008	19.25 / 9.14	0.814 ± 0.010	19.36 / 9.04	0.762 ± 0.009

Table 2: Performance variation according to different parameter settings of $\varepsilon = \{\varepsilon^* - 0.1, \varepsilon^*, \varepsilon^* + 0.1\}$. The performance of unseen TTS and unseen-to-unseen VC tasks is reported in the format of WER / RECS. VC and TTS results are compared within each. The setting $\varepsilon = \varepsilon^*$ shows the best results for both VC and TTS in most of the experiments. The best results within each experiment are marked in bold.

(WER / RECS)	GT (LibriTTS test)	YourTTS (LibriTTS test)	Exp1 (VCTK)	Exp2 (LibriTTS test)	Exp3 (VCTK)	Exp4 (VCTK)	Exp5 (LibriTTS test)	Exp6 (VCTK)
ε^*	N.A	N.A	0.25	0.25	0.43	0.25	0.25	0.275
Unseen-to-unseen VC	$\varepsilon = \varepsilon^* - 0.1$		21.09 / 0.720	16.99 / 0.758	24.20 / 0.727	18.04 / 0.703	8.08 / 0.753	13.32 / 0.719
	$\varepsilon = \varepsilon^*$		14.46 / 0.740	5.14 / 0.811	31.19 / 0.727	13.10 / 0.726	4.95 / 0.782	12.26 / 0.737
	$\varepsilon = \varepsilon^* + 0.1$	3.78 / 0.862	11.79 / 0.790	22.14 / 0.751	7.36 / 0.809	54.83 / 0.700	18.81 / 0.722	7.63 / 0.788
Unseen TTS	$\varepsilon = \varepsilon^* - 0.1$		26.47 / 0.743	58.27 / 0.775	25.47 / 0.733	54.22 / 0.715	79.71 / 0.769	21.03 / 0.748
	$\varepsilon = \varepsilon^*$		19.36 / 0.762	22.15 / 0.821	37.18 / 0.736	33.24 / 0.741	48.15 / 0.810	16.15 / 0.764
	$\varepsilon = \varepsilon^* + 0.1$		28.30 / 0.757	24.39 / 0.815	55.77 / 0.708	38.07 / 0.751	54.25 / 0.805	19.53 / 0.743

Table 3: Mean opinion scores (MOS) for unseen-to-unseen VC. The best results within the same training set are marked in bold. The VITS-based models are further trained to 800k steps.

Model	Test dataset	Framework applied	MOS (Nat / Sim)
			VC (Unseen-to-unseen)
GT	VCTK	N.A	3.99 ± 0.15 / 3.67 ± 0.18
GT	LibriTTS	N.A	4.01 ± 0.14 / 3.63 ± 0.17
C-DSVAE	VCTK	×	3.78 ± 0.15 / 3.37 ± 0.18
YourTTS	VCTK	×	3.81 ± 0.16 / 3.64 ± 0.16
Zero-shot VITS	VCTK	×	3.77 ± 0.15 / 3.82 ± 0.16
Zero-shot VITS	VCTK	○	3.64 ± 0.16 / 3.71 ± 0.17
YourTTS	LibriTTS	×	3.68 ± 0.14 / 3.19 ± 0.18
YourTTS	LibriTTS	○	3.70 ± 0.14 / 3.26 ± 0.18
Zero-shot VITS	LibriTTS	○	3.91 ± 0.12 / 3.41 ± 0.18

$\varepsilon = \varepsilon^* - 0.1$ showed better WER / CER than $\varepsilon = \varepsilon^*$, the model produces noisy artifacts that clearly degrade perceptual quality. The samples can be found in the supplementary material. Also, the models trained with the multilingual dataset with $\varepsilon = \varepsilon^*$ (Exp2 & 5) show superior results over the multilingual baseline (YourTTS in Table 2). They even show VC performance quite similar to the ground truth, in terms of WER.

The Agnosticity of the ε^* Value. Interestingly, despite different audio configurations, datasets, and models, similar values of ε^* were effective for Exp1, 2, 4~6. This would be a useful result for future work, as it even dismisses the necessity of the preliminary task of training a separate HiFi-GAN.

3.4. Subjective Evaluation Results

For the Mean Opinion Score (MOS) test, we randomly sampled 10 sentences as source speeches from the evaluation set of unseen speakers. The target voices were also randomly sampled but selected to be different from the source. We used Amazon Mechanical Turk, and asked 15 native English speakers to rate the naturalness of the synthesized speech (Nat MOS), and how similar the synthesized voice and target voice is (Sim MOS). For limited resources and simplicity of user study, we evaluate perceptual quality only for the VC task. Ratings ranged from 1 to 5 in integers, 5 being the best.

Results. In Table 3, the methods using our framework show better or comparable perceptual naturalness and speaker similarity against other methods, while having significantly better WER /

CER according to Tables 1 and 2. When using the multilingual dataset, using Zero-shot VITS with MDMM for unseen-to-unseen VC shows Nat / Sim MOS comparable to the ground truth audios.

3.5. Discussion

We explain why using the ε^* value obtained from HiFi-GAN is effective. When training HiFi-GAN, the perceptual quality improved as L_{recon} got smaller. This hints that converging a VITS-based model to a smaller L_{recon} induces the decoder to a state where it produces high-quality audio, and thus provides a reason not to use ε bigger than ε^* . However, the audio quality improvement of HiFi-GAN was not significant after L_{recon} got smaller past ε^* . This means that pushing the model towards a L_{recon} value smaller than this does not lead to improvement, but rather hinders the model’s ability to regularize the posterior to the text conditioned prior. This provides a reason not to use ε smaller than ε^* .

4. Conclusion

In this work, we first introduced the importance of tuning the trade-off between the reconstruction loss and the other losses of VITS-based models. We hypothesized that the converged reconstruction loss value ε^* obtained from the preliminary vocoder task with HiFi-GAN might guarantee a sufficient level of reconstruction quality across other tasks. To train models with a specific target reconstruction loss ε^* , we used MDMM to enforce the model to be trained with the constraint. This allows us to obtain a model superior to competitive baselines without searching the balance hyper-parameter tuning. We showed our framework is generalized well across various scenarios of different datasets and models. Moreover, with a bigger multilingual dataset, our method achieves a quality close to ground truth in zero-shot VC. These results hint that the performance of VITS-based models can be improved by pushing towards a certain low reconstruction loss value determined by the quality of the decoder. While we show this with the HiFi-GAN vocoder in this work, it would be an interesting future direction to investigate whether our framework can be applied to other decoders.

5. References

- [1] Y. Jia, Y. Zhang, R. Weiss, Q. Wang, J. Shen, F. Ren, P. Nguyen, R. Pang, I. Lopez Moreno, Y. Wu *et al.*, “Transfer learning from speaker verification to multispeaker text-to-speech synthesis,” *Advances in neural information processing systems*, vol. 31, 2018.
- [2] E. Cooper, C.-I. Lai, Y. Yasuda, F. Fang, X. Wang, N. Chen, and J. Yamagishi, “Zero-shot multi-speaker text-to-speech with state-of-the-art neural speaker embeddings,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6184–6188.
- [3] S. Choi, S. Han, D. Kim, and S. Ha, “Attentron: Few-Shot Text-to-Speech Utilizing Attention-Based Variable-Length Embedding,” in *Proc. Interspeech 2020*, 2020, pp. 2007–2011.
- [4] D. Min, D. B. Lee, E. Yang, and S. J. Hwang, “Meta-StyleSpeech: Multi-speaker adaptive text-to-speech generation,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 7748–7759.
- [5] J. chieh Chou and H.-Y. Lee, “One-Shot Voice Conversion by Separating Speaker and Content Representations with Instance Normalization,” in *Proc. Interspeech 2019*, 2019, pp. 664–668.
- [6] K. Qian, Y. Zhang, S. Chang, X. Yang, and M. Hasegawa-Johnson, “AUTOVC: Zero-shot voice style transfer with only autoencoder loss,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 5210–5219.
- [7] D.-Y. Wu, Y.-H. Chen, and H. yi Lee, “VQVC+: One-Shot Voice Conversion by Vector Quantization and U-Net Architecture,” in *Proc. Interspeech 2020*, 2020, pp. 4691–4695.
- [8] D.-Y. Wu and H.-y. Lee, “One-shot voice conversion by vector quantization,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7734–7738.
- [9] D. Wang, L. Deng, Y. T. Yeung, X. Chen, X. Liu, and H. Meng, “VQMIVC: Vector Quantization and Mutual Information-Based Unsupervised Speech Representation Disentanglement for One-Shot Voice Conversion,” in *Proc. Interspeech 2021*, 2021, pp. 1344–1348.
- [10] S. Yang, M. Tantrawenith, H. Zhuang, Z. Wu, A. Sun, J. Wang, N. Cheng, H. Tang, X. Zhao, J. Wang, and H. Meng, “Speech Representation Disentanglement with Adversarial Mutual Information Learning for One-shot Voice Conversion,” in *Proc. Interspeech 2022*, 2022, pp. 2553–2557.
- [11] J. Lian, C. Zhang, G. K. Anumanchipalli, and D. Yu, “Towards Improved Zero-shot Voice Conversion with Conditional DSVAE,” in *Proc. Interspeech 2022*, 2022, pp. 2598–2602.
- [12] J. Kim, S. Kim, J. Kong, and S. Yoon, “Glow-TTS: A generative flow for text-to-speech via monotonic alignment search,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 8067–8077, 2020.
- [13] D. Rezende and S. Mohamed, “Variational inference with normalizing flows,” in *International conference on machine learning*. PMLR, 2015, pp. 1530–1538.
- [14] J. Kim, J. Kong, and J. Son, “Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 5530–5540.
- [15] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” in *International Conference on Learning Representations*, 2014.
- [16] J. Kong, J. Kim, and J. Bae, “HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 17022–17033, 2020.
- [17] E. Casanova, C. Shulby, E. Gölge, N. M. Müller, F. S. de Oliveira, A. Candido Jr., A. da Silva Soares, S. M. Aluisio, and M. A. Ponti, “SC-GlowTTS: An Efficient Zero-Shot Multi-Speaker Text-To-Speech Model,” in *Proc. Interspeech 2021*, 2021, pp. 3645–3649.
- [18] J. Cong, S. Yang, L. Xie, and D. Su, “Glow-WaveGAN: Learning Speech Representations from GAN-Based Variational Auto-Encoder for High Fidelity Flow-Based Speech Synthesis,” in *Proc. Interspeech 2021*, 2021, pp. 2182–2186.
- [19] Y. Lei, S. Yang, J. Cong, L. Xie, and D. Su, “Glow-WaveGAN 2: High-quality Zero-shot Text-to-speech Synthesis and Any-to-any Voice Conversion,” in *Proc. Interspeech 2022*, 2022, pp. 2563–2567.
- [20] E. Casanova, J. Weber, C. D. Shulby, A. C. Junior, E. Gölge, and M. A. Ponti, “YourTTS: Towards zero-shot multi-speaker tts and zero-shot voice conversion for everyone,” in *International Conference on Machine Learning*. PMLR, 2022, pp. 2709–2720.
- [21] J. Platt and A. Barr, “Constrained differential optimization,” in *Neural Information Processing Systems*, D. Anderson, Ed., vol. 0. American Institute of Physics, 1987.
- [22] H. Cho, W. Jung, J. Lee, and S. H. Woo, “SANE-TTS: Stable And Natural End-to-End Multilingual Text-to-Speech,” in *Proc. Interspeech 2022*, 2022, pp. 1–5.
- [23] L. Wan, Q. Wang, A. Papir, and I. L. Moreno, “Generalized end-to-end loss for speaker verification,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4879–4883.
- [24] C. Veaux, J. Yamagishi, K. MacDonald *et al.*, “Superseded-CSTR VCTK corpus: English multi-speaker corpus for CSTR Voice Cloning Toolkit,” 2016.
- [25] H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, “LibriTTS: A Corpus Derived from LibriSpeech for Text-to-Speech,” in *Proc. Interspeech 2019*, 2019, pp. 1526–1530.
- [26] V. Pratap, Q. Xu, A. Sriram, G. Synnaeve, and R. Collobert, “MLS: A Large-Scale Multilingual Dataset for Speech Research,” in *Proc. Interspeech 2020*, 2020, pp. 2757–2761.
- [27] “M-AILABS dataset,” <https://www.caito.de/2019/01/03/the-m-ail-abs-speech-dataset/>.
- [28] R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. Tyers, and G. Weber, “Common voice: A massively-multilingual speech corpus,” in *Proceedings of The 12th Language Resources and Evaluation Conference*. Marseille, France: European Language Resources Association, May 2020, pp. 4218–4222.
- [29] K. Park, “KSS dataset: Korean single speaker speech dataset,” <https://www.kaggle.com/datasets/bryanpark/korean-single-speaker-speech-dataset>, 2018.
- [30] S. Takamichi, K. Mitsui, Y. Saito, T. Koriyama, N. Tanji, and H. Saruwatari, “JVS corpus: free japanese multi-speaker voice corpus,” *arXiv preprint arXiv:1908.06248*, 2019.
- [31] R. Sonobe, S. Takamichi, and H. Saruwatari, “JSUT corpus: free large-scale japanese speech corpus for end-to-end speech synthesis,” *arXiv preprint arXiv:1711.00354*, 2017.
- [32] Y. Shi, H. Bu, X. Xu, S. Zhang, and M. Li, “AISHELL-3: A multi-speaker mandarin tts corpus and the baselines,” *arXiv preprint arXiv:2010.11567*, 2020.
- [33] H.-T. Luong and H.-Q. Vu, “A non-expert Kaldi recipe for Vietnamese speech recognition system,” in *Proceedings of the Third International Workshop on Worldwide Language Service Infrastructure and Second Workshop on Open Infrastructures and Analysis Frameworks for Human Language Technologies (WLSI/OIAF4HLT2016)*. Osaka, Japan: The COLING 2016 Organizing Committee, Dec. 2016, pp. 51–55.
- [34] “Tunisian modern standard arabic,” <https://www.openslr.org/46/>.
- [35] B. McFee, C. Raffel, D. Liang, D. P. Ellis, M. McVicar, E. Battenberg, and O. Nieto, “librosa: Audio and music signal analysis in python,” in *Proceedings of the 14th python in science conference*, vol. 8, 2015.
- [36] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, “WavLM: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.