



# Fine-grained Noise Control for Multispeaker Speech Synthesis

Karolos Nikitaras<sup>1</sup>, Georgios Vamvoukakis<sup>1</sup>, Nikolaos Ellinas<sup>1</sup>, Konstantinos Klapsas<sup>1</sup>,  
Konstantinos Markopoulos<sup>1</sup>, Spyros Raptis<sup>1</sup>, June Sig Sung<sup>2</sup>, Gunu Jho<sup>2</sup>, Aimilios Chalamandaris<sup>1</sup>,  
Pirros Tsiakoulis<sup>1</sup>

<sup>1</sup>Innoetics, Samsung Electronics, Greece

<sup>2</sup>Mobile Communications Business, Samsung Electronics, Republic of Korea

k.nikitaras@partner.samsung.com, {g.vamvouk, n.ellinas}@samsung.com,  
k.klapsas@partner.samsung.com, {k.markop, s.raptis, js6.sung, gunu.jho, aimilios.ch,  
p.tsiakoulis}@samsung.com

## Abstract

A text-to-speech (TTS) model typically factorizes speech attributes such as content, speaker and prosody into disentangled representations. Recent works aim to additionally model the acoustic conditions explicitly, in order to disentangle the primary speech factors, i.e. linguistic content, prosody and timbre from any residual factors, such as recording conditions and background noise. This paper proposes unsupervised, interpretable and fine-grained noise and prosody modeling. We incorporate adversarial training, representation bottleneck and utterance-to-frame modeling in order to learn frame-level noise representations. To the same end, we perform fine-grained prosody modeling via a Fully Hierarchical Variational AutoEncoder (FVAE) which additionally results in more expressive speech synthesis. Experimental results support our claims and ablation studies verify the importance of each proposed component. Audio samples are available in our demo page<sup>1</sup>.

**Index Terms:** multispeaker text-to-speech synthesis, noise modeling, fine-grained control

## 1. Introduction

Advances in text-to-speech (TTS) neural systems have resulted in models that are capable of synthesizing high quality speech from many different speakers. Such models can be autoregressive (AR) [1, 2] or non-autoregressive (NAR) [3, 4, 5], attentive [1] or duration-informed [3, 2, 4], end-to-end [5, 4] or not [1, 2, 3]. Non end-to-end systems depend on a vocoder [6, 7, 8] that synthesizes speech from the predicted acoustic features.

Multi-speaker TTS systems are commonly trained on publicly available corpora, such as VCTK [9] and LibriTTS [10]. VCTK contains recordings from 109 speakers of different English dialects, whereas LibriTTS contains recordings from about 2,500 English speakers and is divided into disjoint subsets based on data quality. These corpora are collected in a large scale, having as a result multiple artifacts and noise. It is worth mentioning that one fourth of LibriTTS corresponds to signal-to-noise ratio (SNR) of 20 dBs or less [10], when approximated by waveform amplitude distribution analysis (WADA) [11]. Despite that fact, they are used widely in multiple tasks, such as voice cloning [12, 13] and prosody modeling [14].

In order to produce clean speech from such corpora it is important to explicitly model the noise [15, 16, 17]. Furthermore, given that speaker identities and recording/noise conditions are often correlated [15], it is essential to learn speaker representations disentangled from acoustic conditions. Recent works

aim to learn speech representations disentangled from noise-related factors [15] or explicitly model the acoustic conditions [15, 16, 17]. In both cases, speech and noise corpus mixing is a common practice [15, 16, 17]. Typically, if the signal is mixed with artificial noise, it is considered as noisy, otherwise as clean. As a result, the inherent noise that potentially exists within signals is considered as part of the clean speech and thus, it cannot be explicitly modeled. Given that high-quality clean speech data is costly to collect, the only way to counteract such behavior is to incorporate unsupervised noise modeling.

### 1.1. Related Work

A simple way to improve the quality of a speech corpus is to use enhancement methods such as noise suppression and dereverberation prior to text-to-speech training [18]. Recent works train the TTS model directly on the noisy data. Utterance-level speech and noise corpus mixing is a common technique and has been widely used. In [15], speaker timbre is disentangled from acoustic conditions and noise is modeled in a coarse grained way.

Fine-grained noise modeling in a two-step way is performed in [17], where a noise condition module based on U-Net [19] is trained jointly with the TTS model in a semi-supervised way. In Styler [16], frame-level noise is considered as one of several style factors and is modeled in a fully supervised way. The latter two works both apply supervision over the noise representations as they require knowledge of whether the input signal contains clean or noisy speech.

In terms of prosody, [17] models phoneme-level F0 and duration as it follows the FastSpeech architecture [3], while [16] also extends to phoneme-level energy modeling. Both methods use prosody labels extracted by external modules. Regarding the phoneme alignment, they work either in a supervised way or by simple linear compression or expansion of audio to the text's length respectively. We hope to extend previous work using a fully unsupervised training scheme both in terms of noise and prosody modeling.

### 1.2. Proposed Method

We base our TTS model on a Tacotron-like architecture [20], which is an attentive, autoregressive seq-to-seq model and we closely follow [15] in terms of the data augmentation process. In order to learn speech attributes other than content and speaker, we perform fine-grained noise and prosody modeling. Our conditional generative model can be written as  $p(\text{speech}|z_s, z_r, z_p, \text{text})$ , where  $p(z_r|\text{speech})$  and  $p(z_p|\text{speech}, \text{text})$  are the variational posterior distributions

<sup>1</sup><https://innoetics.github.io/publications/fv-noise/index.html>

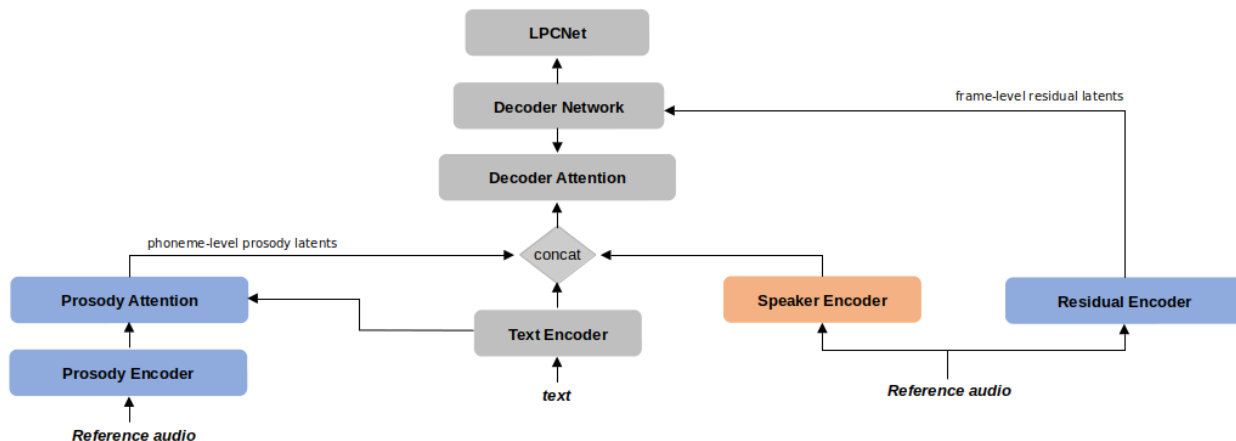


Figure 1: The overall architecture of the proposed system.

that correspond to the residual and prosody encoder, respectively, while  $p(z_s|speech)$  comes from pre-trained speaker encoder network (see section 2.2). All the posterior distributions are defined as diagonal covariance Gaussian whose mean and variance are parameterized by each encoder network, while all the prior distributions as isotropic Gaussian.

We propose a system that performs unsupervised, interpretable frame-level noise and phoneme-level prosody modeling. In order to achieve disentanglement between speech and noise, we incorporate adversarial training, representation bottleneck and utterance-to-frame modeling. Our prosody modeling also enhances this disentanglement, while results in more expressive speech synthesis. Experimental results show that unsupervised, fine-grained noise modeling outperforms other methods and this verifies the effectiveness of our system in controlling and so, removing even the inherent noise within the speech corpus.

All in all, the contributions of this work are the following:

- unsupervised fine-grained noise modeling
- a novel utterance-to-frame training scheme that enables frame level noise modeling
- a low-dimensional (bottleneck) residual latent representation in addition to adversarial training in order to discourage linguistic content leakage into the residual latent space
- fine-grained prosody modeling to further enhance the disentanglement of the noise latent representations from speech related factors as well as synthesize expressive speech
- synthesizing speech with higher estimated SNR than the inherently noisy training data

The rest of the paper is organized as follows. In Section 2, we present the proposed system and its main components in terms of speaker 2.2, noise 2.3 and prosody 2.4 modeling. Section 3 provides ablation studies and objective evaluation comparing our proposed system to other alternatives regarding SNR estimation measured by WADA 3.2.3 as well as the reconstruction performance 3.3. Section 4 concludes the paper indicating directions for future work.

## 2. Method

### 2.1. Model Architecture

The overall model architecture is shown in Figure 1, which is based on [20]. Text encoder converts the phoneme sequences into the text hidden sequence. Prosody encoder maps the mel-spectrogram sequence into frame-level hidden prosody sequence which in turn, is being aligned with the text hidden sequence. We then use the extracted alignment in order to produce the average, raw, phoneme-level mel representation. Fully Hierarchical Variational AutoEncoder (FVAE) takes as input the latter representation and produces hierarchically the three latent dimensions as in [14]. The projection of these three latent dimensions as well as the speaker embedding are concatenated to the text hidden sequence. The noise encoder operates in frame-level and produces the residual hidden sequence. Adversarial training is applied over the intermediate, frame-level residual representation. Also, the Variational Autoencoder (VAE) framework is being utilized over the latent representations of the prosody and residual encoder. Posterior distributions are defined as diagonal covariance Gaussian whose mean and variance are parameterized by each encoder network, while all the prior distributions as isotropic Gaussian with zero mean and unit variance. Finally, we condition the Decoder Network to the residual embedding at each step.

### 2.2. Speaker Encoder

We use transfer learning from speaker verification as in [12]. The pre-trained speaker encoder outputs fixed-dimensional embeddings called d-vectors, which are trained via the Generalized End2End (GE2E) loss [21]. Trainable speaker encoder embeddings could have also been used<sup>2</sup>, however the focus of this work is noise modeling, hence we selected the pre-trained network for speaker embedding.

### 2.3. Residual Encoder

Unsupervised, frame-level modeling entails the risk of overfitting. Thus, we employ several techniques to counteract such behavior. The Residual Encoder has two blocks, each consisting of a convolutional and a batch normalization layer each. It maps the mel-spectrogram sequence into frame-level hidden residual

<sup>2</sup>Some preliminary experiments showed that trainable speaker embeddings like [15] are more robust to noise than pretrained d-vectors but are worse in terms of speaker similarity

representation. We impose a bottleneck over the residual representation by using a small number of dimensions. This prohibits the model to use the latent residual representation in order to directly pass frame-level spectral information to the decoder module. Additionally, adversarial training is being imposed over the hidden representation in order to further discourage the residual latent space from modeling linguistic content, by applying adversarial CTC loss [22]. We also perform utterance-to-frame modeling over the hidden sequence by applying mean pooling with a probability which starts from 1, at the early stage of training, and gradually decreases to 0. Otherwise, the system learns correlated speech and noise representation within the residual latent space, even when using just a single latent dimension. We further enhance this disentanglement by performing prosody modeling. Ablation studies in Section 3.2.3 verify the importance of each proposed component.

#### 2.4. Prosody Modeling

We use three latent dimensions which are trained in a hierarchical, unsupervised way and correspond to phoneme-level f0, energy and duration as in [14]. Given that the noise existence (5-25 dB) within the augmented dataset is prohibitive to the phoneme-to-mel alignment extraction, we used the original version of the speech signals in the input of the prosody encoder.

The Prosody Encoder passes the mel-spectrogram sequence through two blocks, each consisting of a convolutional and a batch normalization layer, followed by an LSTM layer. It outputs a hidden, frame-level prosody representation which is being aligned to the hidden text sequence via the Prosody Attention module. This alignment is being used along with the raw mel sequence resulting to an average, phoneme-level mel representation. FVAE takes in the latter representation and produces, in a hierarchical way, the posterior distributions of three phoneme-level, prosody latents.

### 3. Experiments

Train-clean-100 is chosen as our train set which is a LibriTTS subset consisting of utterances with at least 20 dB SNR, measured by WADA [10]. The train set is augmented with background noise from the CHIME 4 challenge [23] following the data augmentation process in [15]. The motivation here is not to counteract the inherent speaker-noise correlation in order to learn noise-invariant speaker representations as in [15], but to simulate real noisy data, impose noise-related attributes as the predominant residual information and be able to control and then remove the inherent noise existing in our data by not using noise labels in our fine-grained, noise modeling process.

#### 3.1. Model and training setup

The input text is first normalized and converted into a phoneme sequence by a traditional TTS front-end module. The acoustic features used for training the synthesizer are matching the ones by the LPCNet vocoder [7], i.e. 20 bark-scale cepstral coefficients plus pitch period and pitch correlation, while sequences of 80-dim mel-scale filterbank frames are used as the input of the prosody and the residual encoders.

The synthesizer network is based on the attentive, sequence-to-sequence architecture [20], with extra input  $z_p$  as well as  $z_s$  concatenated to the text encoder network output, while  $z_r$  concatenated and passed to the decoder at each step. Assuming a batch sample of B utterances having P as the maximum number of phonemes and F as the maximum number of

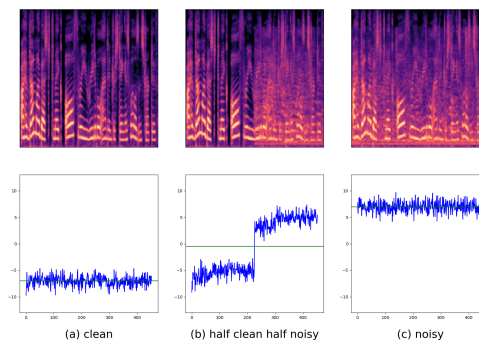


Figure 2: Utterances with varying acoustic conditions along with their residual latent representation

acoustic frames, then  $z_s \in \mathbb{R}^{B \times 256}$ ,  $z_p \in \mathbb{R}^{B \times P \times 3}$  and  $z_r \in \mathbb{R}^{B \times F \times 1}$ .

The network parameters are trained with the Adam optimizer [24], which minimizes the average L1 loss before and after the post-net, batch size 32 and an initial learning rate of  $3 \times 10^{-3}$  linearly decaying to  $3 \times 10^{-5}$  after 100,000 iterations. L2 regularization with weight  $10^{-5}$  is also used.

#### 3.2. Residual Encoder Evaluation

The following three sections aim to quantify the ability of our system to model and control noise. First, we explore the mapping of several utterances with varying acoustic conditions into the residual latent space. We also estimate the SNR of the generated speech by traversing the frame-level latent space. Finally, we compare our system to several baselines in the clean speech production.

##### 3.2.1. Training phase

Figure 2 illustrates a speech signal under different acoustic conditions, along with each residual latent representation. Left column shows the clean version of the signal, at the center column artificial noise is added to the second half of the signal, while at the right column noise is added to the entire length of the signal. It is shown that the 1-dimensional latent distribution consists of a clean and a noisy tail which correspond to the negative and the positive axis, respectively.

##### 3.2.2. Inference phase

During inference time, our system needs to assign one residual value per frame that is being synthesized. Figure 3 shows that by traversing the residual latent space in the range  $[-12, 12]$ , while broadcasting each value across all synthesizing frames, we generate speech and then estimate the mean SNR using WADA. The noisy tail corresponds to about 6 dBs which is close to 5 dB, the lowest SNR level that used in the data augmentation process. The high variance in the clean tail indicates that there is room for further improvements in order to achieve more robust noise modeling.

##### 3.2.3. Performance Evaluation

Results shown in Table 1 quantify the ability to control noise through WADA SNR estimation, by comparing the full proposed model with several alternative models, while also conducting ablation studies:

1. “Ground-truth” corresponds to the ground truth record-

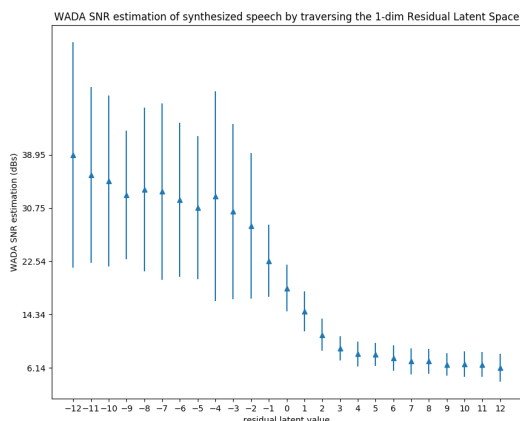


Figure 3: WADA-SNR estimation of synthesized speech by traversing the 1-dim Residual Latent Space

ings

2. “Ground-truth enhanced” are the ground truth recordings enhanced by a speech-enhancement model [25]
3. “Noise-unaware” is a model trained on the ground-truth data
4. “Noise-unaware enhanced” is a model trained on the enhanced ground-truth data
5. “Utt-level” performs utterance-level noise modeling by closely following [15]
6. “Proposed - CTC” is the proposed system without the adversarial CTC loss
7. “Proposed - utt-to-frame” is the proposed system without utterance-to-frame level modeling, that performs directly frame-level noise modeling
8. “Proposed - fvae” is the proposed system without the prosody modeling, trained on the artificially noisy corpus
9. “Proposed” is the proposed system.

The systems (6,7,8,9) perform frame-level noise modeling, (5) performs utterance level noise modeling, while the rest do not perform any noise modeling. All systems are trained on the artificially noisy speech corpus, except from the two noise-unaware systems (4,5) that are trained on the ground-truth speech corpus.

Without the adversarial CTC loss over the intermediate, frame-level residual latents, (6) fails to learn disentangled speech and noise representations. The same holds for the alternative that directly performs frame-level noise modeling (7). These two verify that unsupervised, frame-level modeling entails the risk of overfitting. Comparing the full proposed model (9) with the alternative without the FVAE (8), we show that by performing fine-grained prosody modeling we encourage the residual encoder to capture other factors than prosody and so, result in more clean speech. Supervised noise modeling (2,4) seems incapable of removing the inherent noise, since the performance between both (1,2) and (3,4) is quite close. Fine-grained noise modeling (8,9) results in more clean speech than coarse-grained (5). Finally, the difference between the proposed

Table 1: Average WADA-SNR of synthesized speech across several systems

Model	mean WADA SNR
ground-truth	28.3
ground-truth enhanced	28.7
noise-unaware	21.5
noise-unaware enhanced	21.1
utt-level [15]	21.3
proposed - CTC	1.1
proposed - utt-to-frame	7.6
proposed - fvae	23.2
proposed	<b>36.0</b>

Table 2: Reconstruction performance results

Model	GPE	VDE	FFE	MCD
utt-level	0.32	0.14	0.34	6.87
proposed-fvae	0.38	0.14	0.36	6.96
proposed	<b>0.30</b>	<b>0.11</b>	<b>0.28</b>	<b>6.03</b>

system (9) and the ground-truth recordings (1) illustrates the effectiveness of the unsupervised, fine-grained noise modeling in controlling and removing the inherent noise present in some of the ground-truth recordings.

### 3.3. Prosody encoder evaluation

Table 2 compares the reconstruction performance of the proposed system with several alternative models. To quantify the reconstruction performance we use mel-cepstral distortion (MCD) which evaluates the timbral distortion and F0 Frame Error (FFE) which evaluates the reconstruction of the F0 track. Lower is better for both metrics. The latter is a combination of the Gross Pitch Error (GPE) and the Voicing Decision Error (VDE).

The “Proposed - fvae” has slightly worse performance than “utt-level” and this might be a result of the correlated speech and noise representation within the residual latent space. The Proposed model outperforms the two alternatives across all metrics. As in [14], FVAE captures phoneme-level F0, energy and duration. The significant FFE drop indicates that F0 information is captured, while the drop in VDE that energy and duration are captured.

## 4. Conclusions

In this work, we propose a TTS system that achieves clean and expressive speech synthesis in an unsupervised and interpretable manner. We incorporate adversarial training, representation bottleneck and utterance-to-frame noise modeling in order to factorize noise and speech into disentangled representations. We also enhance this disentanglement by performing fine-grained prosody modeling, in the sense that we encourage the frame-level residual encoder to learn attributes other than prosody. Experimental results show that our system outperforms other methods in clean speech synthesis, indicating that our unsupervised, fine-grained noise modeling method can control and remove the inherent noise that exists in crowd-sourced data. In addition, for future work, we would also like to investigate the capability of the proposed method to learn richer noise representations in order to further improve the synthesized speech quality.

## 5. References

- [1] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. 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)*.
- [2] J. Shen, Y. Jia, M. Chrzanowski, Y. Zhang, I. Elias, H. Zen, and Y. Wu, “Non-attentive tacotron: Robust and controllable neural tts synthesis including unsupervised duration modeling,” *arXiv preprint arXiv:2010.04301*.
- [3] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, “Fastspeech 2: Fast and high-quality end-to-end text to speech,” *arXiv preprint arXiv:2006.04558*.
- [4] J. Donahue, S. Dieleman, M. Bińkowski, E. Elsen, and K. Simonyan, “End-to-end adversarial text-to-speech,” *arXiv preprint arXiv:2006.03575*.
- [5] 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*.
- [6] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, “Wavenet: A generative model for raw audio,” *arXiv preprint arXiv:1609.03499*.
- [7] J.-M. Valin and J. Skoglund, “Lpcnet: Improving neural speech synthesis through linear prediction,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- [8] R. Prenger, R. Valle, and B. Catanzaro, “Waveglow: A flow-based generative network for speech synthesis,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- [9] J. Yamagishi, C. Veaux, K. MacDonald *et al.*, “Cstr vtck corpus: English multi-speaker corpus for cstr voice cloning toolkit (version 0.92).”
- [10] 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,” *arXiv preprint arXiv:1904.02882*.
- [11] C. Kim and R. M. Stern, “Robust signal-to-noise ratio estimation based on waveform amplitude distribution analysis,” in *Ninth Annual Conference of the International Speech Communication Association*.
- [12] 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*.
- [13] S. Choi, S. Han, D. Kim, and S. Ha, “Attentron: Few-shot text-to-speech utilizing attention-based variable-length embedding,” *arXiv preprint arXiv:2005.08484*.
- [14] G. Sun, Y. Zhang, R. J. Weiss, Y. Cao, H. Zen, and Y. Wu, “Fully-hierarchical fine-grained prosody modeling for interpretable speech synthesis,” in *ICASSP 2020-2020 IEEE international conference on acoustics, speech and signal processing (ICASSP)*.
- [15] W.-N. Hsu, Y. Zhang, R. J. Weiss, Y.-A. Chung, Y. Wang, Y. Wu, and J. Glass, “Disentangling correlated speaker and noise for speech synthesis via data augmentation and adversarial factorization,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- [16] K. Lee, K. Park, and D. Kim, “Styler: Style factor modeling with rapidity and robustness via speech decomposition for expressive and controllable neural text to speech,” *arXiv preprint arXiv:2103.09474*.
- [17] C. Zhang, Y. Ren, X. Tan, J. Liu, K. Zhang, T. Qin, S. Zhao, and T.-Y. Liu, “Denoispeech: Denoising text to speech with frame-level noise modeling,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- [18] C. Valentini-Botinhao and J. Yamagishi, “Speech enhancement of noisy and reverberant speech for text-to-speech,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- [19] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention*.
- [20] N. Ellinas, G. Vamvoukakis, K. Markopoulos, A. Chalamandaris, G. Maniati, P. Kakoulidis, S. Raptis, J. S. Sung, H. Park, and P. Tsiakoulis, “High quality streaming speech synthesis with low, sentence-length-independent latency,” *arXiv preprint arXiv:2111.09052*.
- [21] 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)*.
- [22] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *Proceedings of the 23rd international conference on Machine learning*.
- [23] E. Vincent, S. Watanabe, A. A. Nugraha, J. Barker, and R. Marxer, “An analysis of environment, microphone and data simulation mismatches in robust speech recognition,” *Computer Speech & Language*.
- [24] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*.
- [25] E. Tzinis, Z. Wang, and P. Smaragdis, “Sudo rm-rf: Efficient networks for universal audio source separation,” in *2020 IEEE 30th International Workshop on Machine Learning for Signal Processing (MLSP)*.