



Towards Robust FastSpeech 2 by Modelling Residual Multimodality

Fabian Kögel, Bac Nguyen, Fabien Cardinaux

Sony Europe B.V., Stuttgart Laboratory 1, Germany

Fabian.Koegel@sony.com

Abstract

State-of-the-art non-autoregressive text-to-speech (TTS) models based on FastSpeech 2 can efficiently synthesise high-fidelity and natural speech. For expressive speech datasets however, we observe characteristic audio distortions. We demonstrate that such artefacts are introduced to the vocoder reconstruction by over-smooth mel-spectrogram predictions, which are induced by the choice of mean-squared-error (MSE) loss for training the mel-spectrogram decoder. With MSE loss FastSpeech 2 is limited to learn conditional averages of the training distribution, which might not lie close to a natural sample if the distribution still appears multimodal after all conditioning signals. To alleviate this problem, we introduce TVC-GMM, a mixture model of Trivariate-Chain Gaussian distributions, to model the residual multimodality. TVC-GMM reduces spectrogram smoothness and improves perceptual audio quality in particular for expressive datasets as shown by both objective and subjective evaluation.

Index Terms: expressive TTS, residual multimodality, over-smoothness, mixed density networks, gaussian mixture

1. Introduction

The distribution of natural speech signals is multimodal since the same content can be spoken in many different ways and dependent since it is continuous in time and coherent within a speaking voice [1]. Tractably modelling this complex distribution is a fundamental challenge in statistical parametric speech synthesis (SPSS) [2]. A common problem are audio artefacts due to over-smooth predictions caused by naive modelling assumptions [3]. Ultimately, SPSS balances approximations of the natural speech distribution with capabilities of available modelling methods and computing resources. This has led to various families of TTS models that achieve high-quality speech synthesis through autoregressive factorization [4, 5], end to end (E2E) training [6], generative adversarial loss [7, 8], variational autoencoding [9] or normalizing flows [10].

In the family of non-autoregressive architectures, FastSpeech 2 [11, 12] is widely popular for its high-quality and fast parallel generation, small data requirements and inherent controllability. Its two stage pipeline separates the acoustic model from the vocoder and reduces training time by re-using pre-trained vocoders. Further, as discussed by [13], its explicit modelling of pitch, energy and phoneme duration simplifies the distribution to be learned and reduces the over-smoothness and allows for manual control during generation. However, characteristic audio artefacts can still be observed with models from this family - in particular with more expressive and multi-speaker datasets. Although this problem has been treated by fine-tuning the vocoder or training it jointly with the acoustic model in an




Figure 1: Examples of residual multimodality observed in marginal distributions over time and frequency for phoneme ph after all pitch p and energy e conditioning in FastSpeech 2.

E2E architecture [7, 14, 8], we believe this only masks the underlying issue by tuning to a particular dataset using high training efforts. We find that the conditioning is not sufficient for expressive datasets and there is residual multimodality that in conjunction with the inherent assumption of unimodality of the MSE loss function is the root cause for the audio degradation. Adding more conditioning or fine-tuning for individual datasets is not sustainable for the application in controllable expressive speech synthesis with intended large diversity of prosody.

In this paper we propose a mixture model to model this residual multimodality as a compromise between the more powerful distribution mapping approaches (data- and computation-intensive and difficult to optimise approaches of GAN and normalizing flows) and the controllability, fast/data-efficient training and generation speed of the FastSpeech 2 architecture.

We make the following contributions: (1) we demonstrate how smoothness causes artefacts across vocoder architectures for FastSpeech 2, (2) we propose TVC-GMM¹, a novel trivariate-chain gaussian mixture modelling layer which can tractably model the residual multimodality in the spectrogram distribution of non-autoregressive two-stage TTS models, (3) we show that TVC-GMM reduces spectrogram smoothness and improves perceptual audio quality in objective and subjective evaluation in particular for expressive datasets, while maintaining fast and data-efficient training and generation speed.

¹code at <https://github.com/sony/ai-research-code/tvc-gmm>

2. Residual Multimodality

We introduce the term residual multimodality for the gap in complexity between the actual distribution to be learned and the assumed distribution in the loss function. Although deep neural models are universal approximators, the distribution that can practically be learned depends on the loss to be optimised [15]. FastSpeech 2 deploys MSE loss, which optimises training data likelihood only under an assumed gaussian distribution [16] and thus relies on conditioning on external control signals to simplify the complex speech distribution. In Figure 1, we visualise the remaining distribution of our training data after all conditioning in FastSpeech 2 - that is speaker label, pitch bin and energy bin. We show the estimated density of the spectrogram value distribution marginalized over time (left) or frequency (right) for selected phonemes *ph* and pitch/energy bins *p* and *e* in the train split of the single-speaker LJSpeech dataset. Even within single frequency bands, we observe multiple modes. Additionally, we observe large variances in unimodal areas, which might indicate dependencies between bins. The residual multimodality and dependencies that appear despite conditioning can cause over-smoothing artefacts as we show in Section 4.2. Adding even more conditioning is unsustainable and distribution mapping approaches like normalizing flows add computation and data requirements, therefore we commence with modelling the residual multimodality.

3. Trivariate-Chain Gaussian Mixture Modelling (TVC-GMM)




Figure 2: *Illustration of approaches. TVC-GMM (top) models dependencies of adjacent spectrogram bins by a trivariate gaussian chain while FastSpeech 2 (bottom) only models means.*

Let $\mathbf{Y} \in \mathbb{R}^{T \times F}$ be a random variable representing the mel-spectrograms, where T and F are the length and number of frequency bins. Given some conditioning linguistic features \mathbf{L} and control inputs \mathbf{C} , the goal for the acoustic model is to estimate the true distribution of mel-spectrograms by maximizing the likelihood of the training data under the parametrized distribution $p_{\theta}(\mathbf{Y}|\mathbf{L}, \mathbf{C})$. FastSpeech 2 achieves this by minimizing the mean squared error (MSE) between the predicted and ground-truth spectrograms, which implicitly assumes the values of individual time/frequency bins $y^{t,f}$ to be the means of unimodal independent gaussian distributions [16]. As discussed in section 2, this assumption does not hold and learning the mean of the residual multimodal distribution leads to the over-smoothness problem in TTS. To alleviate this problem, we propose TVC-GMM, a mixture model of Trivariate-Chain Gaussian distributions.

3.1. Trivariate-Chain Gaussian distribution

We start by addressing the dependence between adjacent bins. Ideally, we would like to model the full $T \times F$ dimensional spectrogram distribution, but this requires a prohibitively large covariance matrix and fixed T and F . Instead, we break it down and introduce a Trivariate-Chain Gaussian distribution.

Let $\tilde{\mathbf{Y}} \in \mathbb{R}^{T \times F \times 3}$ be a random variable where each element $\tilde{\mathbf{Y}}^{t,f}$ is a triplet of neighbouring spectrogram bins $\mathbf{Y}^{t,f}$, $\mathbf{Y}^{t+1,f}$, and $\mathbf{Y}^{t,f+1}$ modelled by a trivariate Gaussian distribution, i.e.

$$\tilde{\mathbf{Y}}^{t,f} \sim \mathcal{N}(\mu^{t,f}, \Sigma^{t,f}), \quad (1)$$

where $\mu^{t,f} \in \mathbb{R}^3$ and $\Sigma^{t,f} \in \mathbb{R}^{3 \times 3}$ denote the mean and covariance matrix. The $\tilde{\mathbf{Y}}^{t,f}$ are chained together by their overlapping targets and provide a way to represent local variances dependent on correlations between time and frequency steps. Our objective is to learn the parameters $\mu^{t,f}$ and $\Sigma^{t,f}$ and to sample from $\tilde{\mathbf{Y}}$ instead of directly predicting individual spectrogram values as FastSpeech 2 does.

3.2. Trivariate-Chain Gaussian mixture model

To further address residual multimodality, we increase the flexibility of $\tilde{\mathbf{Y}}$ by replacing $\tilde{\mathbf{Y}}^{t,f}$ with a mixture model as similarly proposed for low-dimensional acoustic features in [1]:

$$\tilde{\mathbf{Y}}^{t,f} \sim \sum_{k=1}^K \alpha_k^{t,f} \mathcal{N}(\mu_k^{t,f}, \Sigma_k^{t,f}), \quad (2)$$

where K denotes the number of components and $\alpha_k^{t,f}$ denotes the mixing coefficients. In practice, we adapt the last network layer to predict the parameters $\alpha_k^{t,f}$, $\mu_k^{t,f}$ and $\Sigma_k^{t,f}$ and minimise negative log-likelihood (NLL) loss using the ground-truth spectrogram and two time/frequency shifted copies as targets (see Figure 2 for an illustration). This is what we call Trivariate-Chain Gaussian Mixture Modelling (TVC-GMM).

3.3. Sampling from TVC-GMM

We propose naive sampling and conditional sampling as two ways to sample a mel-spectrogram from TVC-GMM. Naive sampling draws $\langle y^{t,f}, y^{t+1,f}, y^{t,f+1} \rangle$ from $\tilde{\mathbf{Y}}^{t,f}$ for each time/frequency bin in parallel. As a result, we have multiple values at the overlap of the trivariate-chains, which we simply average to smooth sampling noise. However, over-sharpness from noise can introduce new artefacts – although vocoders tolerate it better than over-smoothness (see Section 4.2). Thus, to further increase consistency and reduce sampling noise, conditional sampling uses an iterative algorithm which conditions each $\tilde{\mathbf{Y}}^{t,f}$ on the values drawn in previous bins. Instead of sampling $\langle y^{t,f}, y^{t+1,f}, y^{t,f+1} \rangle$ from $\tilde{\mathbf{Y}}^{t,f}$, we fix the known value of $y^{t,f}$ (predicted as $\hat{y}^{t,f}$ in the previous bin) and obtain a bivariate slice of $\tilde{\mathbf{Y}}^{t,f}$ from which we only sample $\langle y^{t+1,f}, y^{t,f+1} \rangle$ consistent with the previous timestep. The overlapping values in frequency direction are still averaged and we leave investigation into more sophisticated sampling approaches to future work.

4. Experiments

We conduct experiments on several datasets, vocoders and an adapted FastSpeech 2 model to show (1) over-smoothness causes audio degradation, (2) smoothness is exaggerated on expressive speech datasets and (3) TVC-GMM is effective in reducing smoothness and improving perceptual audio quality.




Figure 3: Aligned synthesized mel-spectrogram samples for all datasets and models. FastSpeech 2 models are visibly oversmooth, while TVC-GMM models are closer to ground-truth. Conditional sampling reduces the sampling noise introduced by naive sampling.

4.1. Datasets & Model Details

We report results for three expressive English datasets. We assess expressiveness by pitch range since FastSpeech 2 spreads its bins evenly (Figure 4). (1) LJSpeech [17] is the least expressive and contains 13k (24h) quite monotone samples from a single female speaker. It is the standard evaluation set for many TTS models. (2) VCTK [18] contains 88k (44h, 20min/speaker) samples from 109 English speakers with various accents, deliberately selected for contextual and phonetic coverage. (3) LibriTTS [19] is the standard multi-speaker dataset for TTS models. We use the train-clean-360 split covering 115k (192h, 15min/speaker) samples from 904 speakers and the largest expressive range in our comparison. We set aside a test set of 512 samples per dataset for evaluation, spoken by seen speakers.

We adapt the FastSpeech 2 implementation² of [20] as follows: We reduce encoder and decoder stacks to 4 layers with 2 attention heads and remove the postnet as it is not in the original paper and preliminary experiments did not show significant impact. We also rearrange the variance predictors to run in parallel to better disentangle pitch and energy conditioning. To further improve controllability, we feed the mean external conditioning signal into the predictor during training to encourage it to focus on the variation around this “prosody baseline”. These changes do not affect the mechanism of variance conditioning, as we still learn a fixed set of pitch and energy embeddings. However, during inference we can now control pitch and energy more targeted by setting the prosody baseline instead of scaling the predictors outputs. As vocoder we use the HiFiGAN V1 universal checkpoint and the LJSpeech checkpoint for LJSpeech experiments. We train all models 40k steps on a single Nvidia GeForce RTX 2080ti using NLL loss as described in Section 3.2. FastSpeech 2 has 29.12M parameters, TVC-GMM [k=1] 29.43M (+1%) and TVC-GMM [k=5] 30.25M (+4%). Average training time for TVC-GMM was 2.51 ± 0.97 hours and for FastSpeech 2 2.03 ± 0.87 hours. TVC-GMM [k=5] is only 0.8% slower in inference than FastSpeech 2, mainly due to the enlarged last layer and sampling.

4.2. Insufficient modelling degrades reconstruction quality

Modelling conditional averages in multi-valued mappings leads to smooth spectrogram predictions that also visually lie some-

where between two possible realisations (Figure 3). We use a filter kernel to artificially smooth (gaussian, $\sigma = 1.0$) and sharpen (laplacian, $s = 1$) the ground-truth spectrograms and calculate the change in perceptual speech quality [21] (PESQ) and perceptual audio distance [22] (CDPAM) relative to the ground-truth audio (Table 1). Following [13], we measure smoothness by the variance of the Laplacian-filtered mel-spectrogram Var_L [23]. A lower Var_L indicates more smoothness. We observe characteristic “metallic” (over-smooth) and “bubbling” (over-sharp) artefacts in the audio reconstruction across popular vocoder models³. Generally, over-sharpness is better tolerated. As a naive way to reduce artefacts, we also experiment with fine-tuning HiFiGAN (+finetuned) on smoothed LJSpeech to better tolerate smoothed spectrograms. We find it increases smoothness tolerance on all datasets, however, it also decreases performance on GT spectrograms and requires additional effort for every practical application.




Figure 4: Pitch range of datasets. LibriTTS and VCTK are more diverse than LJSpeech.

4.3. Smoothness is exaggerated for expressive speech

Intuitively, the problem of over-smoothness can be exaggerated for expressive datasets and under prosody control when distributions are under-conditioned and more or further apart modes are averaged. Indeed, in Table 2 we report lower Var_L for LibriTTS samples generated by FastSpeech 2 than for LJSpeech samples and a larger drop in CDPAM relative to GT audio. As Figure 4 demonstrates, LJSpeech has a comparably small pitch range and models trained on more expressive datasets such as VCTK and LibriTTS have to cover a larger diversity with the same number of conditioning bins.

²we adapt <https://github.com/ming024/FastSpeech2/tree/d4e79e>

³audio demo at <https://sony.github.io/ai-research-code/tvc-gmm>

Table 1: *Mel-spectrogram smoothness measured by variance of Laplacian Var_L [23] causes reduction in perceptual speech quality (PESQ) and perceptual audio distance (CDPAM) relative to the ground-truth audio. Calculated over test sets (512 samples/dataset). Over-smoothness/sharpness both induce audio artefacts across vocoder architectures, but over-sharpness is tolerated better.*

		Change in PESQ* \uparrow								CDPAM $\times 10^2$ \downarrow						
		Var_L	Griffin-Lim [24]	HiFiGAN [25]	+finetuned	MelGAN [26]	CARGAN [27]	WaveGlow [28]	WaveRNN [29]	Griffin-Lim [24]	HiFiGAN [25]	+finetuned	MelGAN [26]	CARGAN [27]	WaveGlow [28]	WaveRNN [29]
LJSpeech	GT Spec.	0.37	-1.40	-1.48	-1.66	-2.63	-1.73	-1.09	-1.56	26.1	6.0	10.6	9.3	19.6	5.5	5.7
	+ Sharpen	0.66	-1.78	-1.86	-2.73	-2.77	-2.15	-1.40	-1.89	24.0	8.2	19.4	11.6	17.1	6.5	6.9
	+ Smooth	0.07	-2.57	-2.55	-1.53	-3.07	-2.77	-2.28	-2.55	30.1	10.2	5.2	15.1	23.7	10.8	8.7
VCTK	GT Spec.	0.36	-1.37	-1.64	-1.86	-2.49	-1.28	-1.46	-2.23	27.3	13.8	19.3	14.7	8.4	20.1	20.8
	+ Sharpen	0.77	-1.76	-2.02	-2.6	-2.66	-1.82	-1.68	-2.54	26.6	16.2	25.6	14.4	9.7	18.2	18.5
	+ Smooth	0.06	-2.52	-2.63	-2.01	-3.04	-2.84	-2.33	-2.72	27.9	20.1	15.5	15.9	19.2	20.7	20.5
LibriTTS	GT Spec.	0.41	-1.51	-1.69	-2.03	-2.66	-1.67	-1.63	-2.24	27.8	12.1	17.5	10.2	13.3	17.7	17.7
	+ Sharpen	0.66	-1.79	-2.06	-2.83	-2.82	-2.11	-1.82	-2.49	27.4	14.3	24.9	11.8	13.3	16.3	16.6
	+ Smooth	0.08	-2.68	-2.71	-2.13	-3.16	-2.98	-2.59	-2.91	30.3	18.2	14.4	15.7	20.1	19.6	18.3

* PESQ wide-band (16kHz), mapped to MOS-LQO (1.02 - 4.64 scale) following ITU-T P.862.2

4.4. TVC-GMM improves perceptual audio quality

To evaluate the effectiveness of TVC-GMM, we calculate the smoothness of synthesized mel-spectrograms and evaluate the reconstructed audio in objective and subjective studies. As acoustic models we compare FastSpeech 2 and TVC-GMM models with 1 or 5 modes and naive or conditional sampling. For objective evaluation only (Table 2), we fix the phoneme duration prediction to align with the ground-truth audio and report the perceptual audio distance CDPAM [22]. This metric is explicitly designed for TTS evaluation and shown to pick up on audio artefacts missed by other metrics, but perceived by humans. For subjective evaluation (Table 3), we randomly select 5 test samples per dataset, synthesise FastSpeech 2 and TVC-GMM audio and ask 19 proficient english speakers to rate the overall quality in terms of acoustic quality and prosodic diversity on a Likert-Scale of 1-5. We let them compare between the models and include the ground-truth vocoder reconstruction as upper anchor as this is the best an acoustic model can achieve.

We find modelling with TVC-GMM reduces the smoothness of sampled spectrograms and in turn reduces the acoustic artefacts in the vocoder reconstruction. This is evident in the first block of Table 2, where a higher Var_L indicates that spectrograms generated by TVC-GMM models are less smooth. However, naive sampling from TVC-GMM yields over-sharp spectrograms with a Var_L far above ground-truth, which we attribute to sampling noise. Consequently, we use conditional sampling to reduce noise and improve perceptual audio similarity in particular on the expressive LibriTTS dataset (Table 2, last column). Despite only naive sampling being available for the MOS study, TVC-GMM outperforms in subjective evaluation as well. Note, however, that a significant gap to the GT remains, especially for LibriTTS. We believe this gap is largely due to the duration and variance prediction that control prosody.

5. Conclusion & Future Work

Despite external pitch and energy conditioning, we observe audio artefacts caused by over-smooth mel-spectrogram predictions in FastSpeech 2, in particular in expressive speech datasets. We conclude that this is due to insufficient modelling

Table 2: *Var_L and Perceptual Audio Distance (CDPAM) between ground-truth audio and HiFiGAN reconstruction from synthesis with TVC-GMM models and FastSpeech 2. Our TVC-GMM with both naive and conditional sampling outperforms. Bold values indicate Var_L closest to GT and lowest CDPAM.*

	Var_L			CDPAM $\times 10^2$ \downarrow		
	LJS	VCTK	LTTS	LJS	VCTK	LTTS
GT (HiFiGAN)	0.37	0.36	0.41	6.0	13.8	12.1
FastSpeech 2	0.23	0.25	0.21	11.0	19.7	22.1
TVC-GMM [k=1]	0.45	0.58	0.64	10.4	19.6	21.0
+ cond. sampling	0.43	0.52	0.56	10.4	18.5	18.7
TVC-GMM [k=5]	0.44	0.60	0.63	10.2	19.5	20.1
+ cond. sampling	0.43	0.56	0.55	10.1	18.5	18.5

Table 3: *In subjective evaluation TVC-GMM (naive sampling) outperforms FastSpeech 2 in Mean Opinion Score (MOS).*

	MOS \uparrow		
	LJSpeech	VCTK	LibriTTS
GT (HiFiGAN)	3.87 \pm 0.02	3.90 \pm 0.03	3.80 \pm 0.05
FastSpeech 2	3.54 \pm 0.04	3.43 \pm 0.05	3.23 \pm 0.04
TVC-GMM [k=1]	3.69 \pm 0.03	3.61 \pm 0.04	3.39 \pm 0.04
TVC-GMM [k=5]	3.72 \pm 0.03	3.64 \pm 0.04	3.41 \pm 0.04

of the inherently multimodal natural speech distribution. To model the residual multimodality after conditioning, we propose TVC-GMM which predicts trivariate-chain gaussian distributions in the mel-spectrograms and improves perceptual audio quality under objective and subjective evaluation. We target research and practical applications with a compromise between training effort and quality as TVC-GMM is naturally limited in the number of modes that can be modelled. We also acknowledge that TVC-GMM introduces a new class of sharpness artefacts, but attribute them to sampling noise, which unlike over-smoothness is not a model limitation. We thus plan to investigate more sophisticated sampling strategies in future work.

6. References

- [1] H. Zen and A. Senior, “Deep mixture density networks for acoustic modeling in statistical parametric speech synthesis,” in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Florence, Italy: IEEE, May 2014, pp. 3844–3848.
- [2] H. Zen, A. Senior, and M. Schuster, “Statistical parametric speech synthesis using deep neural networks,” in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. Vancouver, BC, Canada: IEEE, May 2013, pp. 7962–7966.
- [3] G. E. Henter, T. Merritt, M. Shannon, C. Mayo, and S. King, “Measuring the perceptual effects of modelling assumptions in speech synthesis using stimuli constructed from repeated natural speech,” in *Interspeech 2014*. ISCA, Sep. 2014, pp. 1504–1508.
- [4] A. van den 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,” Sep. 2016.
- [5] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu, “Neural Speech Synthesis with Transformer Network,” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, pp. 6706–6713, Jul. 2019.
- [6] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio, Q. Le, Y. Agiomyrgiannakis, R. Clark, and R. A. Saurous, “Tacotron: Towards End-to-End Speech Synthesis,” in *Interspeech 2017*. ISCA, Aug. 2017, pp. 4006–4010.
- [7] J. Donahue, S. Dieleman, M. Bińkowski, E. Elsen, and K. Simonyan, “End-to-End Adversarial Text-to-Speech,” in *International Conference on Learning Representations*. arXiv, Mar. 2021.
- [8] B. Nguyen, F. Cardinaux, and S. Uhlich, “Differentiable Duration Modeling for End-to-End Text-to-Speech,” in *International Conference on Acoustics, Speech, and Signal Processing*. Rhodes Island, Greece: IEEE, 2023.
- [9] J. Kim, J. Kong, and J. Son, “Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech,” in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139. PMLR, Jul. 2021, pp. 5530–5540.
- [10] J. Kim, S. Kim, J. Kong, and S. Yoon, “Glow-TTS: A Generative Flow for Text-to-Speech via Monotonic Alignment Search,” in *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc., 2020, pp. 8067–8077.
- [11] Y. Ren, Y. Ruan, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, “FastSpeech: Fast, robust and controllable text to speech,” in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett, Eds., vol. 32. Curran Associates, Inc., 2019.
- [12] 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,” in *International Conference on Learning Representations*, 2021.
- [13] Y. Ren, X. Tan, T. Qin, Z. Zhao, and T.-Y. Liu, “Revisiting Over-Smoothness in Text to Speech,” in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Dublin, Ireland: Association for Computational Linguistics, May 2022, pp. 8197–8213.
- [14] D. Lim, S. Jung, and E. Kim, “JETS: Jointly Training FastSpeech2 and HiFi-GAN for End to End Text to Speech,” in *Interspeech 2022*. ISCA, Sep. 2022, pp. 21–25.
- [15] K. Hornik, M. Stinchcombe, and H. White, “Multilayer feedforward networks are universal approximators,” *Neural Networks*, vol. 2, no. 5, pp. 359–366, Jan. 1989.
- [16] C. M. Bishop, “Mixture density networks,” *Mixture density networks*, 1994.
- [17] “The LJ speech dataset,” 2017.
- [18] J. Yamagishi, C. Veaux, and K. MacDonald, “CSTR VCTK Corpus: English Multi-speaker Corpus for CSTR Voice Cloning Toolkit (version 0.92),” *The Rainbow Passage which the speakers read out can be found in the International Dialects of English Archive*: (<http://web.ku.edu/idea/readings/rainbow.htm>), Nov. 2019.
- [19] 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 *Interspeech 2019*. ISCA, Sep. 2019, pp. 1526–1530.
- [20] C.-M. Chien, J.-H. Lin, C.-y. Huang, P.-c. Hsu, and H.-y. Lee, “Investigating on Incorporating Pretrained and Learnable Speaker Representations for Multi-Speaker Multi-Style Text-to-Speech,” in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Jun. 2021, pp. 8588–8592.
- [21] *Wideband Extension to Recommendation P.862 for the Assessment of Wideband Telephone Networks and Speech Codecs*, Standard ITU-T P.862.2, Nov. 2007.
- [22] P. Manocha, Z. Jin, R. Zhang, and A. Finkelstein, “CDPAM: Contrastive learning for perceptual audio similarity,” in *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, vol. 2021-June. Institute of Electrical and Electronics Engineers Inc., 2021, pp. 196–200.
- [23] J. Pech-Pacheco, G. Cristobal, J. Chamorro-Martinez, and J. Fernandez-Valdivia, “Diatom autofocusing in brightfield microscopy: A comparative study,” in *Proceedings 15th International Conference on Pattern Recognition. ICPR-2000*, vol. 3. Barcelona, Spain: IEEE Comput. Soc, 2000, pp. 314–317.
- [24] N. Perraudin, P. Balazs, and P. L. Søndergaard, “A fast Griffin-Lim algorithm,” in *2013 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, Oct. 2013, pp. 1–4.
- [25] J. Kong, J. Kim, and J. Bae, “HiFi-GAN: Generative Adversarial Networks for Efficient and High Fidelity Speech Synthesis,” in *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc., 2020, pp. 17 022–17 033.
- [26] K. Kumar, R. Kumar, T. de Boissiere, L. Gestin, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, and A. C. Courville, “Mel-GAN: Generative Adversarial Networks for Conditional Waveform Synthesis,” in *Advances in Neural Information Processing Systems*, vol. 32. Curran Associates, Inc., 2019.
- [27] M. Morrison, R. Kumar, K. Kumar, P. Seetharaman, A. Courville, and Y. Bengio, “Chunked Autoregressive GAN for Conditional Waveform Synthesis,” in *International Conference on Learning Representations*, Jan. 2022.
- [28] 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)*, May 2019, pp. 3617–3621.
- [29] N. Kalchbrenner, E. Elsen, K. Simonyan, S. Noury, N. Casagrande, E. Lockhart, F. Stimberg, A. Oord, S. Dieleman, and K. Kavukcuoglu, “Efficient Neural Audio Synthesis,” in *Proceedings of the 35th International Conference on Machine Learning*. PMLR, Jul. 2018, pp. 2410–2419.