

Revisiting Over-Smoothness in Text to Speech

Yi Ren

Zhejiang University
rayeren@zju.edu.cn

Xu Tan

Microsoft Research
xuta@microsoft.com

Tao Qin

Microsoft Research
taoqin@microsoft.com

Zhou Zhao*

Zhejiang University
zhaozhou@zju.edu.cn

Tie-Yan Liu

Microsoft Research
tyliu@microsoft.com

Abstract

Non-autoregressive text to speech (NAR-TTS) models have attracted much attention from both academia and industry due to their fast generation speed. One limitation of NAR-TTS models is that they ignore the correlation in time and frequency domains while generating speech mel-spectrograms, and thus cause blurry and over-smoothed results. In this work, we revisit this over-smoothing problem from a novel perspective: the degree of over-smoothness is determined by the gap between the complexity of data distributions and the capability of modeling methods. Both simplifying data distributions and improving modeling methods can alleviate the problem. Accordingly, we first study methods reducing the complexity of data distributions. Then we conduct a comprehensive study on NAR-TTS models that use some advanced modeling methods. Based on these studies, we find that 1) methods that provide additional condition inputs reduce the complexity of data distributions to model, thus alleviating the over-smoothing problem and achieving better voice quality. 2) Among advanced modeling methods, Laplacian mixture loss performs well at modeling multimodal distributions and enjoys its simplicity, while GAN and Glow achieve the best voice quality while suffering from increased training or model complexity. 3) The two categories of methods can be combined to further alleviate the over-smoothness and improve the voice quality. 4) Our experiments on the multi-speaker dataset lead to similar conclusions as above and providing more variance information can reduce the difficulty of modeling the target data distribution and alleviate the requirements for model capacity.

1 Introduction

Non-autoregressive text to speech (NAR-TTS) models (Ren et al., 2019, 2020; Peng et al., 2020; Vainer and Dušek, 2020; Łańcucki, 2020; Kim

et al., 2020; Miao et al., 2020) have shown much faster inference speed than their autoregressive counterparts (Wang et al., 2017; Shen et al., 2018; Ping et al., 2018), while achieving comparable or even better voice quality (Ren et al., 2019, 2020). The text-to-speech mapping can be formulated as a conditional distribution $P(y|x)$ where x and y are the text and speech sequences, respectively. Text-to-speech mapping is a one-to-many mapping problem (Wang et al., 2017), since multiple possible speech sequences correspond to a text sequence due to speech variations such as pitch, duration and prosody. Furthermore, speech mel-spectrograms are strongly correlated in time and frequency dimensions (see Section 2.1 for detailed analyses). Therefore, $P(y|x)$ is actually a dependent and multimodal distribution (Ling et al., 2013; Zen and Senior, 2014)¹.

Early non-autoregressive TTS models (Ren et al., 2019; Peng et al., 2020) use mean absolute error (MAE) or mean square error (MSE) as loss function to model speech mel-spectrograms, implicitly assuming that data points in mel-spectrograms are independent to each other and follow a unimodal distribution². Consequently, the mel-spectrograms following dependent and multimodal distributions cannot be well modeled by the MAE or MSE loss, which presents great challenges in non-autoregressive TTS modeling and causes over-smoothed (blurred) predictions in mel-spectrograms (Vasquez and Lewis, 2019; Sheng and Pavlovskiy, 2019).

In this work, we conduct a comprehensive study on the over-smoothing problem in TTS. We find that the over-smoothness is closely related to the mismatch between the complexity of data distri-

¹Here "dependent" means that the different dimensions (in either temporal domain or frequency domain) of y are dependent to each other.

²MAE can be derived from the Laplace distribution and MSE from the Gaussian distribution, both of which are unimodal.

* Corresponding author

butions (*e.g.*, dependent and multimodal distributions are more complex than independent and unimodal distributions) and the power of modeling methods (*e.g.*, simple losses such as MAE and MSE are less powerful than GAN or Glow-based methods). Both simplifying data distributions and enhancing modeling methods can alleviate the over-smoothing problem. From this perspective, we categorize recent methods combating over-smoothness into two classes: **1. Simplifying data distributions:** The data distribution $P(y|x)$ can be simplified by providing more conditional input information. We review two main methods: 1) Providing the previous mel-spectrogram frames $y_{<t}$ to predict current frame y_t , *i.e.*, factorizing the complex dependent distribution $P(y|x)$ into a simpler conditional distribution $\prod_t P(y_t|y_{<t}, x)$, as used in autoregressive TTS models (Wang et al., 2017; Li et al., 2019). 2) Providing more variance information v^3 such as pitch, duration, and energy to predict mel-spectrogram in parallel, *i.e.*, modeling $P(y|x, v)$ rather than $P(y|x)$, as done in some non-autoregressive TTS models (Ren et al., 2020; Łańcucki, 2020). **2. Enhancing modeling methods:** Generally speaking, the modeling method should be powerful enough to fit complex data distributions. We review methods based on different distribution assumptions, including Laplacian mixture loss, structural similarity index (SSIM) (Wang et al., 2004) loss, generative adversarial network (GAN) (Lee et al., 2020) and Glow (Kim et al., 2020).

By studying those methods, we have the following findings. We hope that our studies and findings can inspire the community to design better models for TTS.

- By either autoregressive factorization or providing more variance information as input, complex distributions can be simplified to be less dependent and multimodal, which clearly alleviates the over-smoothing problem and improves the generated voice quality. Among them, providing more variance information such as FastSpeech 2 enjoys the advantages of fast generation due to its non-autoregressive nature.
- Enhanced modeling methods outperform MAE in synthesized voice quality. Laplacian mixture

³The term "variance information" is first mentioned in FastSpeech 2 (Ren et al., 2020), which refers some speech-related conditional information

loss significantly improves the quality of generated mel-spectrograms and enjoys the simplicity of modeling. GAN and Glow achieve the best quality under both subjective and objective evaluations (since they make no assumptions about output distributions), but at the cost of increased training or model complexity.

- To further analyze the effectiveness of combining the basic ideas of the above two categories, we enhance FastSpeech 2 (considering its fast inference speed and good quality in the first category) with Laplacian mixture loss, SSIM, GAN and Glow, respectively. We find that the enhanced FastSpeech 2 generates speech with even better quality and the over-smoothing problem is further alleviated, which shows that the methods in the two categories are complementary to each other.
- We also extend our experiments to multi-speaker TTS task and obtain similar conclusions as above. Besides, we find that Glow has poor modeling ability in multi-speaker scenarios due to limited model capacity and more complex target data distributions compared with the single-speaker scenario, while it can be greatly alleviated by simplifying data distributions (introducing more variance information).

2 Preliminary Study

Text-to-speech mapping is a one-to-many mapping since multiple speech sequences can possibly correspond to a text sequence with different pitch, duration and prosody, making the mel-spectrograms distribution $P(y|x)$ multimodal. And due to the continuous nature of the mel-spectrograms, adjacent data points are dependent to each other. In this section, we first empirically characterize the distribution of $P(y|x)$ in TTS⁴ through visualization (Section 2.1), and then provide a novel perspective to study the over-smoothing problem in TTS (Section 2.2).

2.1 Visualizations

We first visualize the distribution of $P(y|x)$ to see whether it is dependent and multimodal⁵. Denote

⁴In this work, we mainly focus on text (phoneme) to mel-spectrogram mapping, and leave mel-spectrogram to waveform mapping and text to waveform mapping to future work.

⁵We approximate the mel-spectrogram distribution on LJSpeech dataset. The detailed data processing procedure is the same as that in Section 3.2.




Figure 1: The marginal distributions $P(y(t, f)|x = ph)$ for several different phonemes. We choose 4 different phonemes ($ph = AE2, IY1, R, T$) and 4 frequency bins ($f = 10, 20, 50, 70$) in this case study. More marginal distributions are added in Appendix D.1.




Figure 2: Joint distributions of two data points in mel-spectrogram. (a) Along the frequency axis $P(y(t, f_1), y(t, f_2)|x = ph)$, where $f_1 = 10$ and $f_2 = 11$. (b) Along the time axis $P(y(t_1, f), y(t_2, f)|x = ph)$. We choose 2 phonemes ($ph = AO2, AH2$) in this visualization. More visualizations are included in Appendix D.1.

the data point in the t -th frame and the f -th frequency bin in the ground-truth mel-spectrogram as $y(t, f)$, where $t \in [1, T]$, $f \in [1, F]$, and T, F represent the total length and the number of frequency bins of mel-spectrograms respectively. Since different phonemes have different mel-spectrograms, we analyze the distribution of each phoneme separately. Specifically, for all the mel-spectrogram frames corresponding to each phoneme ph , we calculate three distributions: 1) marginal distribution $P(y(t, f)|x = ph)$; 2) joint distribution between two different frequency bins $P(y(t, f_1), y(t, f_2)|x = ph)$; 3) joint distribution between two different time frames $P(y(t_1, f), y(t_2, f)|x = ph)$ ⁶. For each distribution, we first compute the histograms and smooth into probability density functions with kernel density estimation (Dehnad, 1987) for better visualization.

The marginal distributions $P(y(t, f)|x = ph)$ for several different phonemes are shown in Figure 1. It can be seen that the shape of the marginal distribution of each data point $y(t, f)$ in mel-sepctrogram is multimodal, especially for data points in high-frequency bins. The joint distribution $P(y(t, f_1), y(t, f_2)|x = ph)$ and $P(y(t_1, f), y(t_2, f)|x = ph)$ are shown in Figure 2a and Figure 2b, respectively. Obviously, those

⁶We denote one of the frame index among all mel-spectrograms corresponding to the phoneme x as t_1 and set $t_2 = t_1 + 1$.

joint distributions are also multimodal and neighboring points in mel-spectrograms are strongly correlated. From these observations, we can see that the distribution $P(y|x)$ of mel-sepctrograms is multimodal and dependent across time and frequency.

2.2 A Novel Perspective

Table 1: The two categories of methods to combat over-smoothness in TTS.

Categories	Methods
Simplifying data distributions	AR modeling (Li et al., 2019), FastSpeech 2 (Ren et al., 2020), FastPitch (Łańcucki, 2020)
Enhancing modeling methods	SSIM (Wang et al., 2004), Laplacian mixture, GAN (Lee et al., 2020), Glow (Kim et al., 2020)

The dependent and multimodal distribution of $P(y|x)$ increases the difficulty of TTS modeling and causes over-smoothing problem if it is not correctly handled. We provide a novel perspective to depict this problem: the degree of over-smoothness is closely related to the gap between the complexity of data distributions and the power of modeling methods. A larger gap between the power of a modeling method and the complexity of a data distribution results in more severe over-smoothing problem. Consequently, simplifying data distributions and enhancing modeling methods can alleviate the over-smoothing problem. From this perspective,

we list the methods to combat over-smoothness in two categories in Table 1.

In the following sections, we first explore the effectiveness of simplifying data distributions (Section 3) and then that of enhancing modeling methods (Section 4). Finally, we combine the basic ideas of these two categories to improve an existing model (Section 5.1) and conduct further exploration on the multi-speaker dataset (Section 5.2).

3 Simplifying Data Distributions $P(y|x)$

Simplifying data distributions $P(y|x)$ is usually achieved by providing more conditional input information in the TTS literature (Wang et al., 2017; Li et al., 2019; Ren et al., 2020)⁷. In this way, more conditional information in input can alleviate the one-to-many mapping issue, and thus the distribution becomes less multimodal and the correlation along time and frequency is reduced given more condition. There are mainly two methods to provide more conditional input information: 1) autoregressive factorization along the time (Wang et al., 2017; Li et al., 2019) or frequency axis; 2) providing more variance information to predict mel-spectrogram in parallel, as used in some non-autoregressive TTS models (Ren et al., 2020; Łańcucki, 2020). In this section, we first overview these two kinds of methods in detail and conduct the experiment analyses to measure their effectiveness in solving the over-smoothing problem.

3.1 Methods

In this subsection, we overview the two kinds of methods to simplify data distributions, including autoregressive factorization and providing more variance information as input.

Autoregressive Factorization The joint probability $P(y|x)$ can be factorized according to the chain rule in two ways along time and frequency dimensions respectively:

- $P(y|x) = \prod_{t=1}^T P(y_t|y_{<t}, x)$, where $y_{<t}$ is the proceeding frames before the t -th frame and T is the total frames.
- $P(y|x) = \prod_{f=1}^F P(y_f|y_{<f}, x)$, where $y_{<f}$ is the proceeding frequency bins before the f -th

⁷Ren et al. (2019) use knowledge distillation to simplify the target mel-spectrogram itself, which also simplifies the data distribution but affects data quality as analyzed in Ren et al. (2020). Thus, we do not consider this method in our study.

frequency bin and F is the total number of frequency bins (e.g., 80).

More Variance Information Another way to simplify data distributions $P(y|x)$ is to provide more variance information v such as pitch, duration, and energy, to convert $P(y|x)$ into $P(y|x, v)$, as used in previous works (Ren et al., 2020; Łańcucki, 2020). In this way, the distribution becomes less multimodal and the correlation along time and frequency is reduced.

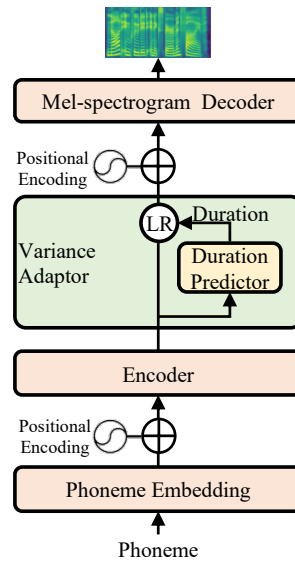


Figure 3: The overall architecture for our baseline model.

3.2 Experiments and Analyses

Experimental Settings We conduct all experiments⁸ on LJSpeech dataset (Ito, 2017). We use ParallelWaveGAN (PWG) (Yamamoto et al., 2020) as the vocoder to convert mel-spectrograms to waveforms. To evaluate the voice quality of the synthesized speech subjectively, we conduct the MOS (Loizou, 2011) tests. To measure the degree of over-smoothness of mel-spectrograms objectively, we calculate the variation of the Laplacian (Pech-Pacheco et al., 2000) (Var_L) on the generated mel-spectrograms and compare with that on the ground-truth mel-spectrograms. We use FastSpeech (Ren et al., 2019) trained with MAE loss as the baseline model as shown in Figure 3. For autoregressive modeling along time, we directly use TransformerTTS (Li et al., 2019). For

⁸The corresponding audio samples are available at <https://revisittts.github.io/revisittts/>.

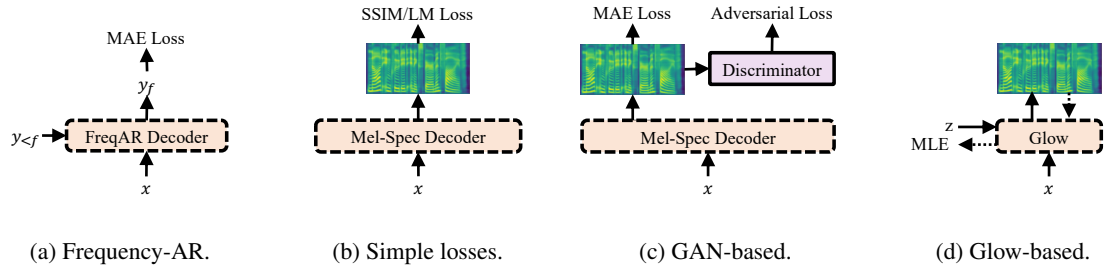


Figure 4: The architecture of mel-spectrogram decoder used in each mel-spectrogram modeling method. All these methods use the same architectures of phoneme embedding, encoder and variance adaptor as the baseline model in Figure 3. x denotes the output hidden of the encoder.




Figure 5: Visualizations of the ground-truth and generated mel-spectrograms by different methods for simplifying data distributions.

autoregressive modeling along frequency, we modify the vanilla mel-spectrogram decoder in baseline model to support autoregressive generation along frequency (which is called FreqAR decoder) and feed $y_{<f}$ to the FreqAR decoder to model $P(y_f|y_{<f})$ as shown in Figure 4a. For the method providing more variance information, we use FastSpeech 2 (Ren et al., 2020), which adds pitch, duration and energy information to the variance adaptor of the baseline model. We put more detailed model descriptions in Appendix A and experimental settings in Appendix B.

Table 2: Results for different methods for simplifying data distributions for TTS. The best scores are in bold.

Methods	MOS	Var _L
<i>GT</i>	4.18±0.08	0.367
<i>GT (PWG)</i>	3.90±0.08	/
<i>MAE</i>	3.64±0.10	0.072
<i>FreqAR</i>	3.80±0.09	0.175
<i>TimeAR</i>	3.81±0.08	0.186
<i>FastSpeech 2</i>	3.83±0.09	0.184

Results and Analyses We conduct MOS evaluation and compute Var_L to compare methods including the baseline model (denoted as *MAE*), autoregressive modeling along frequency (denoted as *FreqAR*) and along time (denoted as *TimeAR*) and *FastSpeech 2*. The results are shown in Table 2. We also visualize the mel-spectrograms gener-

ated by all modeling methods in Figure 5. From the results, we have some observations: 1) Autoregressive modeling along frequency (*FreqAR*) and time (*TimeAR*) dimensions both outperform *MAE* in terms of MOS and Var_L, which shows that simplifying data distributions using frequency or time dimension factorization can ease the over-smoothing problem. However, the autoregressive models suffer from slow inference. 2) FastSpeech 2 also greatly outperforms the baseline model, further indicating that simplifying data distributions by providing more variance information is another way to alleviate the over-smoothing problem. In conclusion, autoregressive modeling and providing more variance information can both simplify the complex distribution to be less dependent and multimodal and thus alleviate the over-smoothing problem. Besides, methods that provide more variance information such as FastSpeech 2 also enjoy the fast inference speed.

4 Enhancing Modeling Methods

Most previous non-autoregressive TTS models (Ren et al., 2019; Wang et al., 2019; Ren et al., 2020) use mean absolute error (MAE) or mean square error (MSE) as training loss. However, they fail to capture dependent and multimodal distributions. MAE loss is derived from the Laplace distribution and MSE from the Gaussian distribution (Chai and Draxler, 2014), which means min-

imizing MAE/MSE will maximize the data log-likelihood under a Laplace/Gaussian distribution. Both of these distributions are unimodal and thus encourage the model to predict a single mode in each data point. As a result, the model just learns an average of all modes, which leads to over-smoothed results. Another problem brought by MAE and MSE is that they are independent across time and frequency for mel-spectrogram output, which ignores the correlation across time and frequency axes in mel-spectrogram.

In this section, we first introduce several enhanced modeling methods to directly model the dependent and multimodal distribution $P(y|x)$ (Section 4.1), and then conduct experiments to compare and analyze these methods (Section 4.2).

4.1 Methods

We list the enhanced methods, including SSIM loss, Laplacian mixture loss, GAN⁹ and Glow-based method and their distribution assumptions in Table 3. We put the details of each method in Appendix A.

Table 3: The distribution assumptions of different methods.

Assumptions	Modeling Methods
Independent & unimodal	Mean absolute error
Independent & multimodal	Laplacian mixture
No assumption	SSIM (Wang et al., 2004), GAN (Goodfellow et al., 2014), Glow (Kingma and Dhariwal, 2018)

Structural Similarity Index (SSIM) Structural Similarity Index (SSIM) (Wang et al., 2004) is one of the state-of-the-art perceptual metrics to measure image quality, which can capture structural information and texture. The value of SSIM is between 0 and 1, where 1 indicates perfect perceptual quality relative to the ground truth. The model architecture of SSIM loss follows the baseline model in Figure 3 and we directly replace the MAE loss

⁹Although GAN is well-known to suffer from the mode collapse issue, practically it performs very well in modeling the multi-modal distribution through well-tuning (Mao et al., 2017). GAN can avoid the average frame prediction in MAE loss that has a strong unimodal assumption, and can generate high-quality and reasonable results when the data distribution is multimodal. Therefore, we regard GAN as "no distribution assumption" method.

in the baseline model with SSIM loss as shown in Figure 4b.

Laplacian Mixture (LM) Loss Laplacian mixture loss¹⁰ can model samples independently with multimodal distribution. As shown in Figure 4b, the basic architecture of mel-spectrogram decoder follows baseline model and we modify the output layer of the baseline model to predict the multimodal distribution of each mel-spectrogram bin.

Generative Adversarial Network (GAN) We introduce adversarial training to better model the dependent and multimodal distribution. Inspired by Wu and Luan (2020); Binkowski et al. (2020), we adopt multiple random window discriminators. We use the LSGAN (Mao et al., 2017) loss to train the TTS model and multi-window discriminators.

Glow Glow (Kingma and Dhariwal, 2018) is a kind of normalizing flow, which maps data into a known and simple prior (e.g., spherical multivariate Gaussian distribution). As shown in Figure 4d, our Glow-based decoder models the distribution of mel-spectrograms conditioned on the encoder output hidden states x .

4.2 Experiments and Analyses




Figure 6: Visualizations of the ground-truth and generated mel-spectrograms by different modeling methods.

¹⁰We choose Laplace distribution as the mixture distribution since the distribution of the magnitude of spectrogram is Laplacian (Tits et al., 2019; Usman et al., 2018; Gazor and Zhang, 2003). We have also tried other mixture distributions (e.g., mixture of logistic and mixture of Gaussian) and have similar findings.

The dataset, baseline model and evaluation metrics are the same as Section 3.2. We conduct MOS evaluation and compute Var_L to compare different modeling methods including MAE (denoted as *MAE*), Laplacian mixture loss (denoted as *LM*), structural similarity index (denoted as *SSIM*), generative adversarial network (denoted as *GAN*), Glow (denoted as *Glow*). The results are shown in Table 4. We also visualize mel-spectrograms generated by all modeling methods in Figure 6. From the results, we can find that:

Table 4: Results for different modeling methods for TTS. The best scores are in bold.

Methods	MOS	Var_L
<i>GT</i>	4.18±0.08	0.367
<i>GT (PWG)</i>	3.90±0.08	/
<i>MAE</i>	3.64±0.10	0.072
<i>SSIM</i>	3.68±0.10	0.201
<i>LM</i>	3.72±0.08	0.296
<i>GAN</i>	3.76±0.08	0.326
<i>Glow</i>	3.78±0.08	0.310

1) *LM* and *SSIM* outperform *MAE* in terms of voice quality according to MOS evaluation and the mel-spectrogram visualizations, which shows that even simply replacing the loss function with those without strong unimodal and independent assumptions can significantly alleviate the over-smoothing problem and improve the generated mel-spectrograms. Among simple loss functions, *LM* performs the best and generates sharper and clearer mel-spectrograms since its Var_L is closer to *GT*, which demonstrates that Laplacian mixture loss can model the multimodal distribution well.

2) *GAN* and *Glow* show more superior performances compared with other modeling methods, indicating that modeling mel-spectrogram with both dependent and multimodal distribution can significantly ease the over-smoothing problem and improve the generated speech. From the visualizations, we can see that *GAN* and *Glow* generate formants with rich details in the middle/high-frequency region. However, GAN-based and Glow-based methods suffer from training complexity and large model footprint respectively: GAN relies on the discriminator to fit the distribution, which causes unstable training and difficult hyperparameters tuning; Glow imposes strong architectural constraints and requires a large model footprint (about 2x model parameters) to keep the bijection between the simple independent latent dis-

tribution (e.g., spherical multivariate Gaussian distribution) and the complex and dependent data distribution.

5 Further Explorations and Extensions

In this section, we first explore the advantages of combining methods from two categories and then perform extensional analyses on multi-speaker dataset.

5.1 Combining Methods from Two Categories

Table 5: Results of different models that combine the basic ideas of two categories. The best scores are in bold.

Methods	MOS	Var_L
<i>GT</i>	4.27±0.10	0.367
<i>GT (PWG)</i>	4.00±0.10	/
<i>FastSpeech 2</i>	3.81±0.10	0.184
<i>FastSpeech 2 + SSIM</i>	3.84±0.09	0.205
<i>FastSpeech 2 + LM</i>	3.90±0.10	0.306
<i>FastSpeech 2 + GAN</i>	3.86±0.10	0.341
<i>FastSpeech 2 + Glow</i>	3.92±0.10	0.315

After studying the two categories to combat over-smoothing problems, we have demonstrated that both simplifying data distributions and enhancing modeling methods can alleviate this problem and improve the voice quality of TTS. A natural thought is to combine the methods from two categories, which may integrate the advantages of both aspects to reduce the gap between the complexity of data distribution and the power of the modeling method, resulting in better voice quality and further alleviating the over-smoothing problem. To demonstrate this idea, we choose the FastSpeech 2 as the model from the first category since it achieves better perceptual voice quality than autoregressive modeling models according to Table 2 and importantly, it enjoys the fast and robust inference advantages due to its non-autoregressive nature. Then we combine two categories by applying enhanced modeling methods to FastSpeech 2 and obtain the following systems: 1) *FastSpeech 2 + SSIM*, which replaces MAE loss with SSIM loss; 2) *FastSpeech 2 + LM*, which predicts the k-component mixture of Laplace distribution and uses LM loss for training; 3) *FastSpeech 2 + GAN*, which adds the adversarial loss to FastSpeech 2; and 4) *FastSpeech 2 + Glow*, which replaces the mel-spectrogram decoder with Glow. We conduct subjective and objective eval-

Table 6: Results for of different multi-speaker TTS models on multi-speaker dataset. The best scores are in bold.

Methods	MOS	Var _L
<i>GT</i>	3.97±0.13	0.338
<i>GT (PWG)</i>	3.67±0.10	/
<i>FastSpeech</i>	2.96±0.10	0.050
<i>FastSpeech + GAN</i>	3.14±0.15	0.251
<i>FastSpeech + Glow</i>	2.94±0.10	0.232
<i>FastSpeech 2</i>	3.18±0.16	0.132
<i>FastSpeech 2 + GAN</i>	3.35±0.13	0.257
<i>FastSpeech 2 + Glow</i>	3.30±0.12	0.264

uations to compare these combined systems with *FastSpeech 2*.

The results are shown in Table 5. We can see that combining *FastSpeech 2* with more powerful modeling methods can further alleviate the over-smoothing problem in terms of Var_L and improve the generated speech quality in terms of MOS, which demonstrate our idea that simplifying data distributions and enhancing modeling methods can be used together and it is complementary to each other to further improve TTS performance. Among these methods, *GAN* performs best in alleviating the over-smoothing problem and *Glow* achieves the best perceptual voice quality, while they suffer from the cost of increased training or model complexity as we described in Section 4.2; compared with *GAN* and *Glow*, *LM* can generate mel-spectrograms with comparable clearness and quality while enjoying its simplicity.

5.2 Analyses on Multi-Speaker Dataset

To demonstrate the generalization of our findings and provide more insights, we conduct experiments on a multi-speaker LibriTTS (Zen et al., 2019) dataset. We modify our models to support multiple speakers by adding speaker embeddings to the encoder outputs to indicate the speaker identity. We put more details of our multi-speaker TTS models and the dataset in Appendix C. We compare the following systems: 1) *FastSpeech*; 2) *FastSpeech + GAN*; 3) *FastSpeech + Glow*; 4) *FastSpeech 2*; 5) *FastSpeech 2 + GAN*; and 6) *FastSpeech 2 + Glow*. The results are shown in Table 6. We can see that 1) simplifying data distributions by providing more variance information and enhancing modeling method can alleviate the over-smoothing problem and improve the generated mel-spectrogram, and combining them together can achieve further better audio quality, which are consistent with the

findings on single-speaker dataset. 2) *FastSpeech + Glow* leads to inferior performance compared with the baseline model (*FastSpeech*), because the multi-speaker dataset has more complex target data distributions and *Glow* requires a large model footprint to capture them as described in Section 4.2. When providing more variance information, *FastSpeech 2 + Glow* achieves much better performance, which can reduce the difficulty of modeling the target data distribution and thus alleviate the requirements for model capacity.

6 Related Works

Non-autoregressive Text to Speech Previous TTS systems such as Tacotron (Wang et al., 2017), Tacotron 2 (Shen et al., 2018), Deep Voice 3 (Ping et al., 2018) and TransformerTTS (Li et al., 2019) synthesize speech sequentially, which suffer from slow inference speed. To solve these shortcomings, various non-autoregressive TTS models are proposed to synthesize spectrogram frames in parallel. *FastSpeech* (Ren et al., 2019) and *ParaNet* (Peng et al., 2020) are early non-autoregressive TTS model which both adopt a fully parallel model architecture and rely on an autoregressive teacher model to provide the alignment between phonemes and mel-spectrograms. *FastSpeech* introduces knowledge distillation for mel-spectrograms to simplify data distributions. *FastSpeech 2* (Ren et al., 2020) and *FastPitch* (Łańcucki, 2020) introduce more variance information as input to further reduce the output uncertainty and ease the one-to-many mapping problem. However, they are trained with MAE loss, which fits independent and unimodal Laplace distribution and results in blurry and over-smoothed mel-spectrograms in inference. *SpeedySpeech* (Vainer and Dušek, 2020) use the combination of MAE and structural similarity index (SSIM) losses to avoid blurry mel-spectrograms. *Glow-TTS* (Kim et al., 2020) and *Flow-TTS* (Miao et al., 2020) both use flow-based decoder to apply some invertible transforms between mel-spectrograms and noise data sampled from simple distribution. Sheng and Pavlovskiy (2019) employ a cascaded Tacotron 2 and *GAN* pipeline to reduce the over-smoothness of synthesized speech. *Multi-SpectroGAN* (Lee et al., 2020) introduces generative adversarial network (*GAN*) and a multi-scale discriminator and is trained with only the adversarial feedback by conditioning hidden states with variance information (e.g., duration,

pitch and energy) to a discriminator. GAN and Flow-based methods can model dependent and multimodal distribution well, while they suffer from training or model complexity. In this work, we conduct systematic studies on several modeling methods in both NAR-TTS and AR-TTS from a novel perspective.

Handling Dependent and Multimodal Distributions

Dependent and multimodal distributions increase the uncertainty for model training and lead to blurry results, which is observed in many generation tasks (Gu et al., 2018; Isola et al., 2017; Mathieu et al., 2016). There are some common ways to handle dependent and multimodal distributions: 1) using loss functions or modeling methods that can well fit the distributions; and 2) introducing some input variables or transforming the target data to simplify data distributions. In neural machine translation, Gu et al. (2018) tackle this problem by introducing knowledge distillation to simplify target data distributions and using fertilities extracted by an external aligner to directly model the non-determinism in the translation process. In image translation task, Isola et al. (2017) compare the generated images by their proposed adversarial generative method with those generated by using MAE and MSE losses and conclude that adversarial loss can help avoid blurry results. In video prediction task, Mathieu et al. (2016) propose a multi-scale architecture, an adversarial training method, and an image gradient difference loss function to deal with the inherently blurry predictions obtained from the MSE loss function. However, there is no systematic analysis on multimodal and dependent distributions in TTS task as far as we know. In this paper, we conduct comprehensive analyses and studies on handling dependent and multimodal distributions in TTS from a novel perspective.

7 Conclusion

In this paper, we revisited the over-smoothing problem in TTS with a novel perspective: the degree of over-smoothness is determined by the gap between the complexity of data distribution and the capability of the modeling method. Under this perspective, we classified existing methods combating over-smoothness into two categories: simplifying data distributions and enhancing modeling methods, and conducted comprehensive analyses and studies on these methods. For simplifying data distributions, we found that both AR

factorization and providing more variance information as input (e.g., FastSpeech 2) can alleviate the over-smoothing problem, and FastSpeech 2 enjoys the advantage of fast generation over AR factorization. For enhancing modeling methods, we found that Laplacian mixture loss can improve the generation quality and enjoy its simplicity, while GAN and Glow can further achieve better quality at the cost of increased training or model complexity. Based on the above findings, we further combined these two categories of methods and found that the over-smoothing problem is further alleviated, and the generated speech quality is further improved, which shows that these two categories are complementary to each other. When performing our analyses on the multi-speaker dataset, we drew similar conclusions and found that providing more variance information can reduce the difficulty of modeling the target data distribution and alleviate the requirements for model capacity.

We hope that our studies can inspire the community and industry to develop more powerful TTS models. Besides, since we are the first to discuss the over-smoothing problem systematically in the speech domain. We hope our analysis methodology as well as the findings can be extended to other tasks and inspire other domains.

Acknowledgments

This work was supported in part by Zhejiang Natural Science Foundation under Grant LR19F020006, National Natural Science Foundation of China under Grant No.61836002, No.62072397.

References

- Sercan Ömer Arik, Mike Chrzanowski, Adam Coates, Gregory Frederick Diamos, Andrew Gibiansky, Yongguo Kang, Xian Li, John Miller, Andrew Y. Ng, Jonathan Raiman, Shubho Sengupta, and Mohammad Shoeybi. 2017. [Deep voice: Real-time neural text-to-speech](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 195–204. PMLR.
- Mikolaj Binkowski, Jeff Donahue, Sander Dieleman, Aidan Clark, Erich Elsen, Norman Casagrande, Luis C. Cobo, and Karen Simonyan. 2020. [High fidelity speech synthesis with adversarial networks](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

- Tianfeng Chai and Roland R Draxler. 2014. Root mean square error (rmse) or mean absolute error (mae)? –arguments against avoiding rmse in the literature. *Geoscientific model development*, 7(3):1247–1250.
- Khosrow Dehnad. 1987. Density estimation for statistics and data analysis.
- Saeed Gazor and Wei Zhang. 2003. Speech probability distribution. *IEEE Signal Processing Letters*, 10(7):204–207.
- Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. 2014. **Generative adversarial nets**. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 2672–2680.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor O. K. Li, and Richard Socher. 2018. **Non-autoregressive neural machine translation**. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- John A Hartigan, Pamela M Hartigan, et al. 1985. The dip test of unimodality. *Annals of statistics*, 13(1):70–84.
- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. 2017. **Image-to-image translation with conditional adversarial networks**. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 5967–5976. IEEE Computer Society.
- Keith Ito. 2017. The lj speech dataset. <https://keithito.com/LJ-Speech-Dataset/>.
- Jaehyeon Kim, Sungwon Kim, Jungil Kong, and Sungroh Yoon. 2020. **Glow-tts: A generative flow for text-to-speech via monotonic alignment search**. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Diederik P. Kingma and Jimmy Ba. 2015. **Adam: A method for stochastic optimization**. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Diederik P. Kingma and Prafulla Dhariwal. 2018. **Glow: Generative flow with invertible 1x1 convolutions**. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 10236–10245.
- Adrian Łańcucki. 2020. **Fastpitch: Parallel text-to-speech with pitch prediction**. *arXiv preprint arXiv:2006.06873*.
- Sang-Hoon Lee, Hyun-Wook Yoon, Hyeong-Rae Noh, Ji-Hoon Kim, and Seong-Whan Lee. 2020. **Multi-spectrogon: High-diversity and high-fidelity spectrogram generation with adversarial style combination for speech synthesis**. *arXiv preprint arXiv:2012.07267*.
- Naihan Li, Shujie Liu, Yanqing Liu, Sheng Zhao, and Ming Liu. 2019. **Neural speech synthesis with transformer network**. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 6706–6713. AAAI Press.
- Zhen-Hua Ling, Li Deng, and Dong Yu. 2013. Modeling spectral envelopes using restricted boltzmann machines and deep belief networks for statistical parametric speech synthesis. *IEEE transactions on audio, speech, and language processing*, 21(10):2129–2139.
- Philippos C Loizou. 2011. Speech quality assessment. In *Multimedia analysis, processing and communications*, pages 623–654. Springer.
- Xudong Mao, Qing Li, Haoran Xie, Raymond Y. K. Lau, Zhen Wang, and Stephen Paul Smolley. 2017. **Least squares generative adversarial networks**. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 2813–2821. IEEE Computer Society.
- Michaël Mathieu, Camille Couprie, and Yann LeCun. 2016. **Deep multi-scale video prediction beyond mean square error**. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Chenfeng Miao, Shuang Liang, Minchuan Chen, Jun Ma, Shaojun Wang, and Jing Xiao. 2020. **Flow-tts: A non-autoregressive network for text to speech based on flow**. In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*, pages 7209–7213. IEEE.
- José Luis Pech-Pacheco, Gabriel Cristóbal, Jesús Chamorro-Martínez, and Joaquín Fernández-Valdivia. 2000. Diatom autofocusing in brightfield microscopy: a comparative study. In *Proceedings 15th International Conference on Pattern Recognition, ICPR-2000*, volume 3, pages 314–317. IEEE.
- Kainan Peng, Wei Ping, Zhao Song, and Kexin Zhao. 2020. **Non-autoregressive neural text-to-speech**. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 7586–7598. PMLR.

- Wei Ping, Kainan Peng, Andrew Gibiansky, Serkan O. Arik, Ajay Kannan, Sharan Narang, Jonathan Raiman, and John Miller. 2018. Deep voice 3: 2000-speaker neural text-to-speech. In *International Conference on Learning Representations*.
- Yi Ren, Chenxu Hu, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. 2020. [FastSpeech 2: Fast and high-quality end-to-end text-to-speech](#). *arXiv preprint arXiv:2006.04558*.
- Yi Ren, Yangjun Ruan, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. 2019. [FastSpeech: Fast, robust and controllable text to speech](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 3165–3174.
- Jonathan Shen, Ruoming Pang, Ron J. Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, RJ-Skerrv Ryan, Rif A. Saurous, Yannis Agiomyrgiannakis, and Yonghui Wu. 2018. [Natural TTS synthesis by conditioning wavenet on MEL spectrogram predictions](#). In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2018, Calgary, AB, Canada, April 15-20, 2018*, pages 4779–4783. IEEE.
- Leyuan Sheng and Evgeniy N Pavlovskiy. 2019. Reducing over-smoothness in speech synthesis using generative adversarial networks. In *2019 International Multi-Conference on Engineering, Computer and Information Sciences (SIBIRCON)*, pages 0972–0974. IEEE.
- Noé Tits, Kevin El Haddad, and Thierry Dutoit. 2019. [The theory behind controllable expressive speech synthesis: a cross-disciplinary approach](#). *arXiv preprint arXiv:1910.06234*.
- Mohammed Usman, Mohammed Zubair, Mohammad Shiblee, Paul Rodrigues, and Syed Jaffar. 2018. Probabilistic modeling of speech in spectral domain using maximum likelihood estimation. *Symmetry*, 10(12):750.
- Jan Vainer and Ondřej Dušek. 2020. [Speedyspeech: Efficient neural speech synthesis](#). *arXiv preprint arXiv:2008.03802*.
- Aäron Van Den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W Senior, and Koray Kavukcuoglu. 2016. Wavenet: A generative model for raw audio. *SSW*, 125.
- Sean Vasquez and Mike Lewis. 2019. [Melnet: A generative model for audio in the frequency domain](#). *arXiv preprint arXiv:1906.01083*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Yiren Wang, Fei Tian, Di He, Tao Qin, ChengXiang Zhai, and Tie-Yan Liu. 2019. Non-autoregressive machine translation with auxiliary regularization. In *AAAI*.
- Yuxuan Wang, RJ Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, et al. 2017. [Tacotron: Towards end-to-end speech synthesis](#). *arXiv preprint arXiv:1703.10135*.
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612.
- Jie Wu and Jian Luan. 2020. [Adversarially trained multi-singer sequence-to-sequence singing synthesizer](#). *arXiv preprint arXiv:2006.10317*.
- Ryuichi Yamamoto, Eunwoo Song, and Jae-Min Kim. 2020. [Parallel wavegan: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram](#). In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*, pages 6199–6203. IEEE.
- Heiga Zen, Viet Dang, Rob Clark, Yu Zhang, Ron J Weiss, Ye Jia, Zhifeng Chen, and Yonghui Wu. 2019. [Libritts: A corpus derived from librispeech for text-to-speech](#). *arXiv preprint arXiv:1904.02882*.
- Heiga Zen and Andrew Senior. 2014. Deep mixture density networks for acoustic modeling in statistical parametric speech synthesis. In *2014 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 3844–3848. IEEE.

A Details in Modeling Methods

A.1 Baseline Model

Our baseline model is based on FastSpeech (Ren et al., 2019). The dimension of phoneme embeddings and the hidden size of the self-attention are set to 256; the number of attention heads is set to 2 and the kernel sizes of the 1D-convolution in the 2-layer convolutional network after the self-attention layer are set to 9 and 1, with input/output size of 256/1024 for the first layer and 1024/256 in the second layer; the output linear layer converts the 256-dimensional hidden states into 80-dimensional mel-spectrograms. The size of the phoneme vocabulary is 76, including punctuations. In the duration predictor, the kernel sizes of the 1D-convolution are set to 3, with input/output sizes of 256/256 for both layers and the dropout rate is set to 0.5.

A.2 Autoregressive Modeling along Frequency

Autoregressive modeling along frequency (denoted as FreqAR) factorizes $P(y|x)$ to $\prod_{f=1}^F P(y_f|y_{<f}, x)$, where $y_{<f}$ is the proceeding frequency bins before the f -th frequency bin and F is the total number of frequency bins (e.g., 80). We implement FreqAR by adding an extra small LSTM to the top of the four feedforward Transformer blocks¹¹ as shown in Figure 7. For timestep $f = 1$, we set the input hidden h_0 of LSTM to the output hidden of the mel-spectrogram decoder q , which has a size of $T \times H$, where T is the total number of the frames, and H is the hidden size of the LSTM. For $1 < f < F$, we concatenate the previous output frequency bins y_{f-1} with the size of $T \times 1$ and q with the size of $T \times H$ along the channel axis and project the channel to H with a linear layer as the input hidden h_f of LSTM. The output hidden of LSTM is projected to o_f using another linear layer. Finally, we concatenate y_1, y_2, \dots, y_{80} along the channel axis and obtain the output mel-spectrogram. We train the FreqAR model with teacher forcing. In this work, we set H to 32.




Figure 7: The architecture of the decoder for autoregressive modeling along frequency. “FFT blocks” denotes the feedforward Transformer Blocks (Ren et al., 2019).

A.3 Structural Similarity Index (SSIM)

Structural Similarity Index (SSIM) (Wang et al., 2004) is one of the state-of-the-art perceptual metrics to measure image quality, which can capture

¹¹Since the time axis of mel-spectrogram has a variable length, we cannot model the autoregressive along frequency by regarding the time axis as the channel dimension and directly feeding the transposed mel-spectrogram into a causal decoder similar to that in TransformerTTS.

structural information and texture. The value of SSIM is between 0 and 1, where 1 indicates perfect perceptual quality relative to the ground truth. For each data point in predicted and ground-truth mel-spectrograms ($\hat{s} = \hat{y}(t, f)$ and $s = y(t, f)$), the SSIM value is computed as: $SSIM(s, \hat{s}) = \frac{2\mu_s\mu_{\hat{s}}+C_1}{\mu_s^2+\mu_{\hat{s}}^2+C_1} \cdot \frac{2\sigma_{s\hat{s}}+C_2}{\sigma_s^2+\sigma_{\hat{s}}^2+C_2}$, where μ_s and $\mu_{\hat{s}}$ denote the means for two regions, which are centered in s and \hat{s} within a 2D-window with size (W, W) respectively, and we set W to 11; σ_s and $\sigma_{\hat{s}}$ are standard deviation for regions s and \hat{s} ; $\sigma_{s\hat{s}}$ is the covariance of regions s and \hat{s} and $C_1 = 0.0001$ and $C_2 = 0.0009$ are constant values for stabilizing the denominator. The SSIM loss for all mel-spectrogram data points $y(t, f)$ and $\hat{y}(t, f)$ is expressed as: $\mathcal{L}_{SSIM} = \frac{1}{TF} \sum_{t=1}^T \sum_{f=1}^F (1 - SSIM(y(t, f), \hat{y}(t, f)))$. The model architecture of SSIM loss follows the baseline model and we directly replace the MAE loss in the baseline model with SSIM loss as shown in Figure 4b.

A.4 Laplacian Mixture (LM) Loss

Laplacian mixture loss¹² can model samples independently with multimodal distribution. Let $\text{La}(y; \mu, \beta)$ denotes the probability distribution function for a Laplace random variable y , where μ and β are the mean and average absolute deviation (MAD) of the Laplace distribution. The k -component mixture of Laplace distribution is defined as follows:

$$P(y(t, f)) = \sum_{k=1}^K \pi_k(t, f) \times \text{La}(y(t, f); \mu_k(t, f), \beta_k(t, f)),$$

where $\mu_k(t, f)$ and $\beta_k(t, f)$ are the mean and MAD of the distribution of the k -th Laplacian component. The log-likelihood loss under mixture of Laplace distribution is:

$$L_{LM} = -\frac{1}{TF} \sum_{t=1}^T \sum_{f=1}^F \log \left\{ \sum_{k=1}^K \pi_k \text{La}(y(t, f); \mu_k, \beta_k) \right\}.$$

As shown in Figure 4b, the basic architecture of mel-spectrogram decoder follows baseline model and we modify the output layer of the baseline

¹²We choose Laplace distribution as the mixture distribution since the distribution of the magnitude of spectrogram is Laplacian (Tits et al., 2019; Usman et al., 2018; Gazor and Zhang, 2003). We have also tried other mixture distributions (e.g., mixture of logistic and mixture of Gaussian) and have similar findings.

model to predict $\mu_k(t, k)$, $\sigma_k(t, k)$ and $\pi_k(t, k)$ for each component k in each mel-spectrogram bin located in (t, k) . Then the model is optimized with L_{LM} loss only and we set $k = 5$. In inference, for each data point in the mel-spectrogram, we first choose the Laplacian component k according to the probability π_1, \dots, π_K and then sample it from the Laplacian distribution $\text{La}(y(t, f); \mu_k(t, f), \beta_k(t, f))$.

A.5 Generative Adversarial Network (GAN)

We introduce adversarial training to better model the dependent and multimodal distribution. Inspired by Wu and Luan (2020); Binkowski et al. (2020), we adopt multiple random window discriminators where the input mel-spectrogram is randomly clipped into 3 clips with different window lengths, and each clip is fed into one discriminator. The discriminator consists of a 3-layer 2D-convolutional network¹³ with LeakyReLU activation, each followed by the batch normalization and dropout layer, and finally, an extra linear layer is added to project the hidden states to a probability to measure whether the input is a real mel-spectrogram sample. We use the LS-GAN (Mao et al., 2017) loss to train the TTS model G and multi-window discriminators D_1 , D_2 and D_3 : $L_{adv_D} = \sum_{i=1}^3 \mathbb{E}_y (D_i(y) - 1)^2 + \mathbb{E}_{\hat{y}} D_i(\hat{y})^2$, $L_{adv_G} = \frac{1}{3} \sum_{i=1}^3 \mathbb{E}_{\hat{y}} (D_i(\hat{y}) - 1)^2$, where \hat{y} and y are the generated and ground-truth mel-spectrograms.

We use multiple random window discriminators to implement the generative adversarial network. The architecture of multiple random window discriminators is shown in Figure 8.

A.6 Glow

Glow (Kingma and Dhariwal, 2018) is a kind of normalizing flow, which maps data into a known and simple prior (e.g., spherical multivariate Gaussian distribution). Glow is optimized with the exact log-likelihood of the data and as a generative model, it is very good at modeling all dependencies within very high-dimensional input data, which is usually specified in the form of a dependent and multimodal probability distribution. To make a fair comparison with other modeling methods, we only take the output hidden states of the encoder as the condition to each step of flows and use spherical

¹³Through experiments, we find that the 2D-CNN-based network performs better than 1D-CNN in discriminator, especially for the reconstruction of high-frequency details.

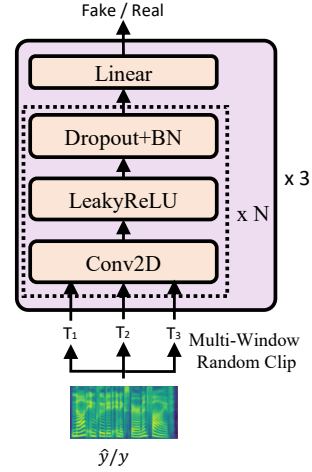


Figure 8: The architecture of multiple random window discriminators. \hat{y} and y denote the predicted and the ground-truth mel-spectrogram respectively.

Gaussian as the prior distribution like in the vanilla Glow (Kingma and Dhariwal, 2018). As shown in Figure 4d, our Glow-based decoder models the distribution of mel-spectrograms conditioned on the encoder output hidden states x . The Glow-based decoder is composed of k flow steps f_1, \dots, f_k , each of which consists of an affine coupling layer, an invertible 1×1 convolution, and an activation normalization layer. Glow-based decoder maps the spherical Gaussian random variables z to the mel-spectrograms as

$$z \sim \mathcal{N}(z; 0, \mathbf{I}), y = f_0 \circ f_1 \circ \dots \circ f_k(z, x). \quad (1)$$

In training (dotted line in Figure 4d), we directly minimize the negative log-likelihood of the data, which can be calculated using a change of variables:

$$z = f_k^{-1} \circ f_{k-1}^{-1} \circ \dots \circ f_0^{-1}(y), \quad (2)$$

$$\log p_\theta(y|x) = \log p_\theta(z) + \sum_{i=1}^k \log |\det(\mathbf{J}(f_i^{-1}(y, x)))|, \quad (3)$$

where the first term in Eq. (3) is the log-likelihood of the spherical Gaussian, \mathbf{J} is the Jacobian and θ is the parameters of the Glow-based mel-spectrogram decoder. In inference (solid line in Figure 4d), we sample z from the spherical multivariate Gaussian distribution and generate mel-spectrogram using Eq. (1).

The architecture of Glow-based modeling method is shown in Figure 9. We choose non-causal WaveNet (Van Den Oord et al., 2016) as the network in the affine coupling layers following Kim et al. (2020). We set the number of flow steps K to 24 and the layers of WaveNet to 4^{14} . We also set the dilation of WaveNet to 1 since we do not need very large receptive fields as that used in vocoder.




Figure 9: The architecture for Glow-based modeling method. The training and inference directions are represented with dotted and solid lines respectively.

B Experimental Settings

In this section, we describe detailed experimental settings including datasets, training and inference details and evaluation criterion.

Datasets We conduct all experiments on LJSpeech dataset (Ito, 2017), which contains 13,100 English audio clips (about 24 hours) of a single speaker and corresponding text transcripts. We split the dataset into three sets: 12,228 samples for training, 349 samples (with document title LJ003) for validation and 523 samples (with document title LJ001 and LJ002) for testing following Ren et al. (2020). We convert the text sequence into the phoneme sequence (Arik et al., 2017; Wang et al., 2017; Shen et al., 2018) with an open-source grapheme-to-phoneme tool¹⁵ to alleviate the mispronunciation problem. The raw waveform is transformed into mel-spectrograms following Shen et al. (2018) and we set frame size

¹⁴The total number of parameters of Glow-based method (43M) is about twice that of other modeling methods (26M) and we find that the performance degrades when using the smaller model, indicating that the Glow-based modeling method requires a large model footprint to keep the bijection.

¹⁵<https://github.com/Kyubyong/g2p>

and hop size to 1024 and 256 with respect to the sample rate 22050.

Training and Inference We train the model on 1 NVIDIA 2080Ti GPU, with batch size of 48 sentences. We use the Adam optimizer (Kingma and Ba, 2015) with $\beta_1 = 0.5$, $\beta_2 = 0.998$ and $\varepsilon = 10^{-9}$ which can stabilize the adversarial training. We follow the same learning rate schedule in Vaswani et al. (2017). It takes 160k steps for training until convergence, except Glow-based model, which needs 480k steps. In the inference process, the output mel-spectrograms are transformed into audio samples using pre-trained Parallel WaveGAN (Yamamoto et al., 2020)¹⁶.

Subjective Evaluation To evaluate the perceptual quality, we conduct the MOS and CMOS (Loizou, 2011) tests. We randomly choose 30 samples from the test set for subjective evaluation. Twenty native English speakers are asked to make quality judgments about the synthesized speech samples. The text content keeps consistent among different systems so that all testers only examine the audio quality without other interference factors. For MOS, each evaluator is asked to mark the subjective naturalness of a sentence on a 1-5 Likert scale. For CMOS, listeners are asked to compare pairs of audio generated by systems A and B and indicate which of the two audio they prefer and choose one of the following scores: 0 indicating no difference, 1 indicating small difference, 2 indicating a large difference and 3 indicating a very large difference. The estimated hourly wage paid to participants is about \$10 and we totally spent about \$900 on participant compensation.

Objective Evaluation Pech-Pacheco et al. (2000) propose the variation of the Laplacian as a “blurriness metric” for images since the Laplacian filter can define edges well and blurry images barely have any edges while a well-focused (clear) image is expected to have a high variation of the Laplacian in grey levels. Inspired by their work, we introduce the variation of the Laplacian to mel-spectrograms and analyze its correlation with the smoothness. Specifically, the variation of the Laplacian Var_L is given by:

$$\text{Var}_L(\bar{y}) = \sum_t \sum_f \left(|L(t, f)| - \frac{1}{NM} \sum_m \sum_n |L(m, n)| \right)^2$$

¹⁶<https://github.com/kan-bayashi/ParallelWaveGAN>

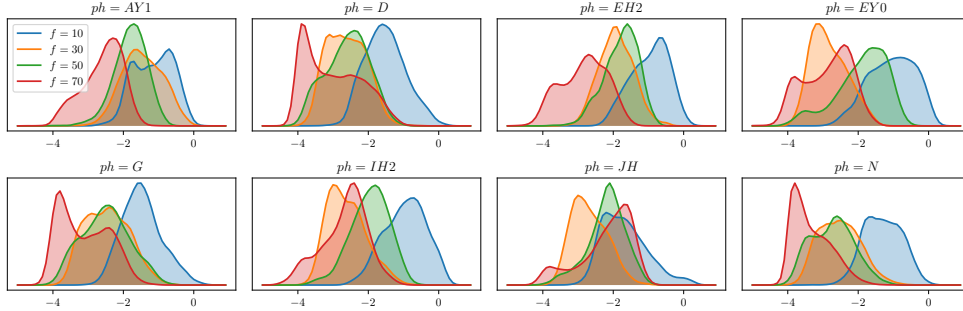


Figure 10: More marginal distributions mel-spectrogram $P(y(t, f)|x = ph)$ for several different phonemes ph .




Figure 11: More joint distributions $P(y(t, f_1), y(t, f_2)|x)$ of two data points in mel-spectrogram along the frequency axis ($f_1 = 10$ and $f_2 = 11$).




Figure 12: More joint distributions $P(y(t, f_1), y(t, f_2)|x)$ of two data points in mel-spectrogram along the frequency axis ($f_1 = 10$ and $f_2 = 50$).




Figure 13: More joint distributions $P(y(t, f), y(t + 1, f)|x)$ of two data points in mel-spectrogram along the time axis.

where $L(t, f)$ is the convolution of the (predicted or ground-truth) mel-spectrogram $\bar{y}(t, f)$ with the

Laplacian operation mask L defined as:

$$8211 \quad L = \frac{1}{6} \begin{pmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{pmatrix}.$$

The variation of the Laplacian increases with decreased smoothness of mel-spectrograms. We calculate the variation of the Laplacian Var_L of the predicted mel-spectrogram and that of the ground-truth mel-spectrogram to see how close they are. We use an open-source tool¹⁷ to compute Var_L .

C Experimental Settings for Multi-Speaker TTS

In this section, we describe the experimental settings for multi-speaker TTS (Section 5.2). We first introduce the dataset, and then describe the model details.

Dataset We conduct our experiments on the *train-clean-100* subset in LibriTTS (Zen et al., 2019), which contains about 54 hours speech audio samples and their corresponding text transcriptions. We choose the *train-clean-100* subset since the data in this subset is clean and the total time of speech audio in this subset is comparable with that in LJSpeech, which help us analyze the influence of multi-speaker setting while excluding the influence of training data size and noisy audio. These audio samples are recorded by 123 female speakers and 124 male speakers. We randomly choose 200 audio samples as the validation set, 200 of them as the test set and the rest of them as the training set.

Model Details Based on the single-speaker baseline model described in Section 3.2, to introduce the speaker identity information into our model, we add an extra speaker embedding module. We look up the speaker embedding from this module and add it to the encoder outputs. The hidden size of the speaker embedding module is the same as the encoder hidden size.

D More Analyses on Mel-Spectrogram Distributions

D.1 More Visualizations

We plot more marginal and joint distributions (along time and frequency) of mel-spectrogram in Figure 10, 11, 12 and 13.

D.2 Multimodality Evaluation on the Marginal Distributions

To further demonstrate that combination of the methods from two categories (FastSpeech 2 and

¹⁷https://github.com/petronav/Blur_Detection/blob/master/Variance_of_Laplacian/blur_check_vol.py

other enhanced modeling methods) can alleviate the over-smoothing problem, we conduct Hartigan’s dip test (Hartigan et al., 1985) on the marginal distributions $P(y(t, f)|x = ph)$ given f and ph to measure the degree of multimodality of each method. Specifically, we denote the dip test value given the distribution P as $\mathcal{D}(P)$ and a lower value means more multimodal. we calculate the $\mathcal{D}(P)$ under the marginal distributions $P(\bar{y}(t, f)|x = ph)$ and average $\mathcal{D}(P)$ under different ph and f to obtain the averaged dip test value $\bar{\mathcal{D}}$. We compute $\bar{\mathcal{D}}$ on the following systems: 1) *GT*, the ground-truth mel-spectrogram; 2) *MAE*, which is the baseline model as described in Section 3.2; 3) *FastSpeech 2* as described in Section 3.1; 4) *FastSpeech 2 + SSIM*, which replaces MAE loss with SSIM loss; 5) *FastSpeech 2 + LM*, which predicts the k-component mixture of Laplace distribution and is trained with LM loss; 6) *FastSpeech 2 + GAN*, which adds the adversarial loss to FastSpeech 2; and 7) *FastSpeech 2 + Glow*, which replaces the mel-spectrogram decoder with Glow.

Table 7: Results of different models combining the basic ideas of two categories. The best scores are in bold.

Methods	$\bar{\mathcal{D}}$
<i>GT</i>	0.049
<i>MAE</i>	0.064
<i>FastSpeech 2</i>	0.060
<i>FastSpeech 2 + SSIM</i>	0.058
<i>FastSpeech 2 + LM</i>	0.060
<i>FastSpeech 2 + GAN</i>	0.054
<i>FastSpeech 2 + Glow</i>	0.056

The results are shown in Table 7. We can see that all of these combined methods can increase the degree of multimodality of the marginal distributions $P(y(t, f)|x = ph)$ and GAN achieves the best $\bar{\mathcal{D}}$ among these methods, which is consistent with our findings in Section 5 that all the combined methods can alleviate the over-smoothing problem and GAN performs the best.

E Potential Negative Societal Impacts and Limitations

Although our analyses can inspire the community and industry to develop more powerful TTS models, it may result in unemployment for people with related occupations such as broadcaster and radio host. Besides, powerful TTS systems may be used in non-consensual voice cloning and fake media

generation, which might be harmful to society. As for the limitation, we combine existing methods to build a better TTS system rather than proposing a brand new text-to-speech model.

F Ethics Statement

F.1 For every submission

F.1.1 Did you discuss the *limitations* of your work?

Yes. See Appendix E.

F.1.2 Did you discuss any potential *risks* of your work?

Yes. See Appendix E.

F.1.3 Do the abstract and introduction summarize the paper’s main claims?

Yes. See abstract and Section 1.

F.2 Did you use or create *scientific artifacts*?

No.

F.3 Did you run *computational experiments*?

Yes.

If yes:

F.3.1 Did you report the *number of parameters* in the models used, the *total computational budget* (e.g., GPU hours), and *computing infrastructure* used?

Yes. See Section 3.2 and Appendix B.

F.3.2 Did you discuss the experimental setup, including *hyperparameter search* and *best-found hyperparameter values*?

Yes. See Section 3.2 and Appendix B.

F.3.3 Did you report *descriptive statistics* about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Yes. See Section 3.2, 4.2 and 5.1.

F.3.4 If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Yes. See Appendix B.

F.4 Did you use *human annotators* (e.g., crowdworkers) or *research with human subjects*?

Yes.

If yes:

F.4.1 Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

Yes. See Appendix B.

F.4.2 Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such *payment is adequate* given the participants’ demographic (e.g., country of residence)?

Yes. See Appendix B.

F.4.3 Did you discuss whether and how *consent* was obtained from people whose data you’re using/curating (e.g., did your instructions explain how the data would be used)?

Yes. See Appendix B.

F.4.4 Was the data collection protocol *approved* (or *determined exempt*) by an ethics review board?

N/A.

F.4.5 Did you report the basic demographic and geographic characteristics of the *annotator* population that is the source of the data?

Yes, they are all native English speakers. See Appendix B.