MAXIMIZING MUTUAL INFORMATION FOR TACOTRON

A PREPRINT

Peng Liu[†], Xixin Wu[‡], Shiyin Kang[†], Guangzhi Li[†], Dan Su[†], Dong Yu[†] [†] Tencent AI Lab [‡] Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, China {feanorliu, shiyinkang, guangzhilei, dansu, dyu}@tencent.com wuxx@se.cuhk.edu.hk

September 4, 2019

ABSTRACT

End-to-end speech synthesis method such as Tacotron, Tacotron2 and Transformer-TTS already achieves close to human quality performance. However compared to HMM-based method or NN-based frame-to-frame regression method, it is prone to some bad cases, such as missing words, repeating words and incomplete synthesis. More seriously, we cannot know whether such errors exist in a synthesized waveform or not unless we listen to it. We attribute the comparatively high sentence error rate to the local information preference of conditional autoregressive models. Inspired by the success of InfoGAN in learning interpretable representation by a mutual information regularization, in this paper, we propose to maximize the mutual information between the predicted acoustic features and the input text for end-to-end speech synthesis methods to address the local information preference problem and avoid such bad cases. What's more, we provide an indicator to detect errors in the predicted acoustic features as a byproduct. Experiment results show that our method can reduce the rate of bad cases and provide a reliable indicator to detect bad cases automatically.

1 Introduction

Tacotron [35] and Tacotron2 [28] are conditional autoregressive (CAR) models trained with teacher forcing [37]. The condition is summarized form the input text with attention mechanism [1]. Transformer-TTS [24] can be considered as another instance of CAR model, with effective utilization of self-attention mechanism [34]. Such architecture can be trained in an end-to-end way, so it has a much shorter pipeline and needs less expert knowledge and human labor. It is flexible enough to adapt for speaking style [29, 36] and multi-speaker [9, 17]. In addition, it is easy to be combined with neural vocoder [18, 31, 32] to enhance the synthesized waveform quality.

Training with teacher forcing induces a mismatch between the training period and the inference period, usually known as exposure bias [26]. Even worse, it strengthens the local information preference [6] for the CAR model. We explain the local information preference intuitively first. At each time step during training, the CAR model receives a teacher forcing input and a conditional input. The teacher forcing input is one previous time step from the target. The conditional input is the text to be synthesized. If the CAR model learns to copy the teacher forcing input, or to predict the target totally depending on teacher forcing input without using the conditional information, it still gets small training mean square error (MSE). Finally the model, which achieves small MSE, may not learn to depend on the condition at all. So at the inference period, the CAR model generates results that have nothing to do with the condition. Note that local information preference still exists even if teacher forcing is not used. When a random variable x admits autoregressive dependency over a conditional random variable z, i.e. $p(x|z) = \prod_i (x_i|x_{<i}, z)$, an universal function approximator, such as RNNs used in the CAR model, can in theory represent the distribution without condition on z [6].

The local information preference weakens the dependency between the predicted acoustic features and the text condition when training a CAR speech synthesis model. In most cases, the CAR speech synthesis model learns to depend on the text condition to predict the acoustic features. However they are prone to bad cases. We argue that this is caused

by the local information preference of the model. Since the model prefers predicting the acoustic features from the teacher forcing input at training stage, it does not model the dependency between the text condition and the predicted acoustic features sufficiently. If we can strengthen the dependency, we may reduce the bad-case rate. In [5], the authors propose a information-theoretic regularization for generative adversarial networks (GAN) [10] to learn a set of disentangled latent codes. The authors separate GAN's input noise vector into incompressible noise and latent codes with factorized distribution. But the generator of GAN is free to ignore the additional latent codes and predicts observations only conditioning on the incompressible noise. To eliminate such trivial solutions, the authors maximize the mutual information between the laten codes and the observations for GAN. This leads to the InfoGAN model. The idea is straightforward. Since the the mutual dependency between two variables can be measured by mutual information. Maximizing mutual information (MMI) would strengthen the dependency between the laten codes and the observations, and hence eliminate the trivial solutions that the GAN's generator models the observations without depending on the latent codes. Enlightened by InfoGAN, we propose to maximize the mutual information between the text condition and the predicted acoustic features to strengthen the dependency for CAR speech synthesis models. This would alleviate the local information preference problem and reduce the rate of bad cases. Besides we train an auxiliary CTC [12] recognizer to maximize the mutual information. The edit distance between the CTC greedy decoding result and the text condition can be used as an indicator to detect errors in the synthesized acoustic features.

In the following, we begin with reviewing the related work in section 2. Then we explain the local information preference formally for CAR speech synthesis models and review the existing designs in Tacotron which prevent the model from predicting the target totally depending on teacher forcing input in section 3. We also explain why text-based WaveNet [31] and WaveRNN [18] do not suffer from such problems in this section. Finally we explain our method and provide the experiment results.

2 Related work

Many previous works focus on improving Tacotron's reliability. [14] adopts professor forcing to mitigate the exposure bias induced by training with teacher forcing. The authors use diagonal attention penalty to enforce that the alignment between the acoustic features and the text is approximately diagonal in [30]. In [42], the authors propose to use the alignment information form hand-crafted labels or from an HHM-based system to guide the attention for Tacotorn. Since there exists a large body of legacy corpus and HMM-based systems, this is an efficient way to improve Tacotron. But it is not trained in an end-to-end way. The implicit duration model of Tacotron uses alignment information that is not self-contained. Transformer-TTS adopts self-attention structure to improve the training and inference efficiency and to shorten the long range dependency path between any two inputs at different time steps [24].

There is a frequently observed problem in variational autoencoder (VAE) [22] training called optimization challenge [4] or posterior collapse [20, 33]. If the decoder of VAE is expressive enough, especially when an autoregressive decoder is used, a VAE may learn trivial latent representations. Because the decoder could reconstruct the target without using information from the latent representation. This problem is attributed to the local information preference property of VAE in [6]. Various methods are proposed to correct this shortcoming, such as weakening the the decoder [4, 13] and changing the training objective [6, 20, 33, 40]. We draw a lot of inspiration from this line of works. Since Tacotron uses a powerful autoregressive decoder, it may also ignore the latent representation learned from input text.

Maximum mutual information is used to estimate HMMs in speech recognition [2]. In [23], the authors apply maximum mutual information to another sequence-to-sequence task, conversational response generation, to increase the diversity of the generated text.

3 CAR model tends to ignore the condition

In this section we first explain local information preference for CAR model formally. Then we explain why Tacotron still works though it tends to ignore the text condition. Finally we explain why WaveNet and WaveRNN, also integrated with powerful autoregressive decoders, do not suffer from this problem.

3.1 Variational encoder-decoder perspective of CAR model

Usually we perform maximum likelihood estimation (MLE) to train a CAR speech synthesis model. And the model communicates information form the text to the acoustic features through the time-aligned latent variables. Such latent variables exist in various speech recognition and synthesis systems, such as the hidden states in the HMM-based speech synthesis system, the forward-backward search matrix in the CTC recognizer, and the attention variables in Tacotron. We can formalize the CAR speech synthesis model as a variational encoder-decoder (VED) [41]. We use

t and x to represent a text and its corresponding acoustic features in the training set. Since it is a CAR model, the conditional likelihood can be written as $\log p_{\theta}(x|t) = \sum_{i=1}^{N} \log p_{\theta}(x_i|x_{<i}, t)$, where N is the number of acoustic frames in x. For simplicity, we suppose the distribution of the time-aligned latent variables, c, is factorizable, i.e. $\log p_{\theta}(c|x,t) = \sum_{i=1}^{N} \log p_{\theta}(c_i|x_{<i},t)$. At least this is true for the ad hoc treatment of the attention variables in Tacotron [27]. Then $\log p_{\theta}(x|t) = \log \int_{c} p_{\theta}(x,c|t) dc = \sum_{i=1}^{N} \log \int_{c_i} p(x_i,c_i|x_{<i},t) dc_i$. The training objective is to maximize the sum of the conditional likelihood of each t, x pair in the training set. For a training pair at time step i:

$$\log p_{\theta}(\boldsymbol{x}_{i}|\boldsymbol{x}_{< i}, \boldsymbol{t}) = D_{\mathrm{KL}}(q_{\phi}(\boldsymbol{c}_{i}|\boldsymbol{x}_{< i}, \boldsymbol{t})||p_{\theta}(\boldsymbol{c}_{i}|\boldsymbol{x}_{\le i}, \boldsymbol{t})) + \mathcal{L}(\theta, \phi, \boldsymbol{x}, \boldsymbol{t})$$
(1)

The first RHS term is the KL divergence of the encoder approximate from the model posterior (it is the posterior because it has access to the current x_i). The second RHS term is the variational lower bound. Since the KL-divergence term is always non-negative:

$$\log p_{\theta}(\boldsymbol{x}_{i}|\boldsymbol{x}_{

$$= \mathbb{E}_{q_{\phi}(\boldsymbol{c}_{i}|\boldsymbol{x}_{

$$= -D_{\mathrm{KL}}(q_{\phi}(\boldsymbol{c}_{i}|\boldsymbol{x}_{

$$= \log p_{\theta}(\boldsymbol{x}_{i}|\boldsymbol{x}_{(2)
$$(3)$$$$$$$$$$

In the above equations, ϕ is the attention encoder parameters and θ is the autoregressive decoder parameters for a CAR speech synthesis model. Since we suppose the text communicates information to the acoustic features only through the time-aligned latent variables, by Bayes rule, we have $p(x_i, c_i | x_{< i}, t) = p(x_i | x_{< i}, c_i) p(c_i | x_{< i}, t)$ used in Eq 2. Note that the variational encoder-decoder formalization is a bit different from the original VAE. From Eq 1, we can see that $q_{\phi}(c_i|x_{<i},t)$ is used to approximate the model posterior distribution, $p_{\theta}(c_i|x_{<i},t)$. But it does not use information from the current x_i . Because x_i is the acoustic feature frame to predict at inference time step i, we cannot use it as the input to the encoder. We can use $q_{\phi}(c_i | x_{< i}, t)$ as the prior distribution, $p_{\theta}(c_i | x_{< i}, t)$. Then the KL-divergence term becomes 0 in Eq 2. If we use a deterministic function to calculate c_i , Eq 2 becomes the training objective of Tacotron. In Eq 3, the KL-divergence term is 0 only when x_i and c_i is conditional independent. In such case, the time-aligned latent variables, c, are meaningless. If the model learns meaningful time-aligned latent variables, the KL-divergence term is positive. When trained to maximize $\log p_{\theta}(\boldsymbol{x}_i | \boldsymbol{x}_{< i}, \boldsymbol{t})$, the model would not learn meaningful time-aligned latent variables to avoid the extra cost if $x_{< i}$ contains enough information to predict x_i . Since the time aligned latent variables are the bridges that communicate information from the text to the acoustic features, the model cannot exploit the text efficiently without the latent variables. we arrive at a similar conclusion as variational lossy autoencoder (VLAE): information can be modeled locally by the CAR model without using information from the time-aligned latent variables will be modeled locally, only the remainder will be modeled using them [6]. We argue that this is one of the possible reasons why attention mechanism cannot learn alignment under some bad configurations [38].

3.2 Why Tacotron learns to condition on text?

We argue that Tacotron learns to condition on the text mainly because of several designs: the reduction window, the large frame shift and the dropout in the decoder prenet. Reduction window is a frame dropout mechanism like the word dropout used in VAE language model (VAELM) [4] to weaken the connection between autoregressive steps. Setting reduction factor to 5 [35] can be considered as dropping 80% frames at equal intervals. This is a bit different from dropping words randomly to a certain percentage in VAELM. But they work in a similar way.

We can use the MSE between the teacher forcing input and the acoustic target as a metric for information locality. If the MSE is smaller, it is easier for the CAR model to predict the target only based on the teacher forcing input without using information from the text. We list the frame averaged MSE of mean-std normalized log mel wrapped short-time Fourier transform (STFT) magnitude for different configurations for the LJSpeech dataset [16] in Table 1. If a reduction window is used, we repeat the teacher forcing input reduction factor times to make the number of frames consistent with that of the target. From Table1, we can see that using larger reduction factor and frame shift could increase the MSE between the teacher forcing input and the acoustic target, which indicates that the connection between them is weakened. To achieve smaller training MSE, the model has to depend more on the text. In [14], the authors point out that the decoder prenet dropout in Tacotron could make the model condition more on the input text. Intuitively, the dropout makes the teacher forcing input incomplete, so the model has to condition more on the text to reconstruct the target. The reduction factor and large frame shift are vital designs in Tacotron. Not only because they speed up training and inference, but also they weaken the local dependency on autoregressive inputs to make the model depends more on input text. Using 12.5 ms frame shift is also important for Griffin-Lim algorithm to reconstruct acceptable waveforms. The reconstructed waveforms using 5 ms frame shift is too quivering (the quiver can be removed by a WaveRNN vocoder in our experiment). Also dropping teacher forcing input frames randomly to a certain percentage is a cheap trick to make the model more robust, which is not applied to Tacotron in previous works.

	frame shift (ms)	
reduction factor	5	12.5
1	0.08357	0.14002
2	0.13502	0.25805
5	0.26686	0.56535

Table 1: Mean square error (MSE) between the teacher forcing input and the acoustic target for different reduction factor and frame shift configurations for LJSpeech.

3.3 Why text-based WaveNet and WaveRNN do not suffer?

WaveNet and WaveRNN are also CAR models. When trained with text input, they are seldomly reported bad cases such as omitting words or incomplete synthesis. Because the text information is unfolded to acoustic frame level according to a duration model and then upsampled to waveform sample point level by a transposed convolution stack or by repeating [31]. Each waveform sample point has a strict correspondence to a piece of text information. The CAR model could exploit the text information directly without using the time-aligned latent variables. It is easier to model the correlation between the sample point and the corresponding text information. So the model may learn a strong dependency between them.

4 Maximizing mutual information (MMI) for Tacotron

Although the previous mentioned designs in Tacotron alleviate the local information preference, they weaken the autoregressive decoder and decrease the model's performance. A model using reduction factor 2 generates better perceptual results than one using reduction factor 5 [35]. This indicates that the more the autoregressive model is weakened, the more drop in performance is induced. Even worse, Tacotron make mistakes, such as repeating words, omitting words and incomplete sentences, which seldomly appear in HMM-based methods [7] or NN-based frame-to-frame regression methods [8, 19, 39]. The dependency between the predicted acoustic features and the text input in Tacotron is not sufficiently modeled. If the dependency is sufficiently modeled and the model is penalized heavily when it makes mistakes during training, the generated acoustic features should strictly follow the text. So we take the InfoGAN approach that maximize the mutual information between the predicted acoustic features and the input text during training to strengthen the dependency between them.

4.1 MMI with an auxiliary recognizer

The mutual information between the input text, t, and the predicted acoustic features, \tilde{x} , is

$$\begin{split} I(\widetilde{\boldsymbol{x}}; \boldsymbol{t}) &= H(\boldsymbol{t}) - H(\boldsymbol{t}|\widetilde{\boldsymbol{x}}) \\ &= \mathbb{E}_{\widetilde{\boldsymbol{x}} \sim p_{\alpha}(\boldsymbol{x})} [\mathbb{E}_{\boldsymbol{t} \sim p_{\alpha}(\boldsymbol{t}|\widetilde{\boldsymbol{x}})} [\log p_{\alpha}(\boldsymbol{t}|\widetilde{\boldsymbol{x}})]] + H(\boldsymbol{t}) \\ &= \mathbb{E}_{\widetilde{\boldsymbol{x}} \sim p_{\alpha}(\boldsymbol{x})} [D_{\mathrm{KL}}(p_{\alpha}(\boldsymbol{t}|\widetilde{\boldsymbol{x}})) | q_{\beta}(\boldsymbol{t}|\widetilde{\boldsymbol{x}})) + \mathbb{E}_{\boldsymbol{t} \sim p_{\alpha}(\boldsymbol{t}|\widetilde{\boldsymbol{x}})} [\log q_{\beta}(\boldsymbol{t}|\widetilde{\boldsymbol{x}})]] + H(\boldsymbol{t}) \\ &\geq \mathbb{E}_{\widetilde{\boldsymbol{x}} \sim p_{\alpha}(\boldsymbol{x})} [\mathbb{E}_{\boldsymbol{t} \sim p_{\alpha}(\boldsymbol{t}|\widetilde{\boldsymbol{x}})} [\log q_{\beta}(\boldsymbol{t}|\widetilde{\boldsymbol{x}})]] + H(\boldsymbol{t}) \\ &= \mathbb{E}_{\boldsymbol{t} \sim p(\boldsymbol{t}), \widetilde{\boldsymbol{x}} \sim p_{\alpha}(\boldsymbol{x}|\boldsymbol{t})} [\log q_{\beta}(\boldsymbol{t}|\widetilde{\boldsymbol{x}})] + H(\boldsymbol{t}) \end{split}$$
(5)

 $\alpha = \{\theta, \phi\}$ is the CAR model parameters. In Eq 4 we introduce an auxiliary distribution $q_{\beta}(t|\tilde{x})$ to approximate the posterior $p_{\alpha}(t|\tilde{x})$ since it is intractable. The lower bound derivation uses the variational information maximization technique [3, 5]. H(t) is a constant for our problem. From Eq 5, we can see that maximizing the mutual information between the input text and the predicted acoustic features is equivalent to training an auxiliary recognizer which maximizes the probability of recognizing the input text from the predicted acoustic features with respect to the CAR model parameters, α , and the auxiliary recognizer parameters, β . This is intuitively sound. If the predicted acoustic features are consistently recognized as the input text, of course the model gets the correct result. Adding the mutual information term to the training objective in Eq 1 can penalize the model if it ignores the dependency between the predicted acoustic features and the text. When this penalty is stronger than the KL-divergence term in Eq 3, the model learns meaningful time-aligned latent variables to exploit the text.

4.2 CTC recognizer for Tacotron

To keep the end-to-end property, we use a simple CTC recognizer as the auxiliary recognizer. The CTC recognizer uses the same convolution stack + bidirectional LSTM [15] layer structure as the Tacotron2's text encoder for simplicity except that the former has an extra CTC loss layer. Lack of a language model is usually considered as a drawback of the CTC recognizer [11]. But this quite meets our demand. Since we do not want a language model to remedy the detected errors. Minimizing the CTC loss could strengthen the dependency between the predicted acoustic features and the input text during training. At inference stage, we can use greed decoding to get the recognition result and compare it with the input text to find out the errors in the synthesized acoustic features. Naturally we get an automatic error detection method for end-to-end speech synthesis models. Since there is a strong correlation between a text that is hard to be synthesized and the corresponding acoustics that are hard to be recognized, e.g. "dillydally", "namby-pamby" and "hahahahaha", the automatic error detection should be reliable. When the recognizer detects an error, it indicates that the input text is hard to be synthesized. We validate the reliability in our experiment. Besides, an important trick to make the CTC recognizer work is averaging the CTC loss by the number of the acoustic frames. The CTC loss is usually not averaged when training a stand-alone speech recognizer. However the autoregressive loss is already averaged by the number of acoustic frames. If the CTC loss is not averaged, the gradients from it dominates the training and leads to a failure. Also all the tokens in the CTC target should be vocable. Tokens like "HEAD", "TAIL", and "WORD BOUNDARY", which may be used as input tokens, should be skipped.

One concern is that if the CAR model copies the teacher forcing input and the auxiliary recognizer works well, it achieves small recognition error, but does not learn the correlation between the predicted acoustic features and the input text. We argue that the model fitting is an asymptotic progress. Initially the model cannot copy the teacher forcing input correctly, it has to depend more on the input text to reduce the recognition loss. The model should learn the correlation when exploiting the input text. So maximizing the mutual information between the predicted acoustic features and the input text eliminates the copying solution.

The final loss function is:

$$\mathcal{L} = |\boldsymbol{x}_{mel} - \widetilde{\boldsymbol{x}}_{mel}| + |\boldsymbol{x}_{linear} - \widetilde{\boldsymbol{x}}_{linear}| + CELoss(\boldsymbol{x}_{stop}, \widetilde{\boldsymbol{x}}_{stop}) + \lambda CTCLoss(\boldsymbol{t}, \widetilde{\boldsymbol{x}}_{mel})$$
(6)

The first 2 RHS terms are the reconstruction losses for Mel spectrum and linear spectrum. Also the model minimize the cross entropy loss for stop tokens and the CTC loss between the predicted Mel spectrum and the text to be synthesized. λ controls the relative weight for the CTC loss. The linear loss is used, because we use the Griffin-Lim algorithm to reconstruct waveforms to monitor the training progress. There might be a weight for cross entropy loss term too.

5 Experiments

We first verify the correlation between bad-case rate and the configuration of frame shift and reduction factor. Then we show that maximizing the mutual information between the predicted acoustics and the text to be synthesized can reduce the rate of bad case. We also show the reliability of the error indicator.

5.1 Experiment setup

We use LJSpeech for English and Databaker Chinese Standard Mandarin Speech Corpus (db-CSMSC)¹ for Mandarin Chinese in our experiments. LJSpeech contains 13,100 audio clips of a single female speaker. We process the transcriptions with Festival² to get the phoneme sequences. db-CSMSC contains 10,000 standard Mandarin sentences recorded by a single female native speaker and recorded in a professional recording studio. The dataset contains the Chinese character and pinyin transcriptions and hand-crafted time intervals. In our experiments, we only use the pinyin transcription and transfer the pinyin sequence to a pinyin scheme which contains initials and sub-finals. Our pinyin scheme contains much less units than the initial-final pinyin scheme. It can alleviate the out-of-vocabulary and data sparsity problem.

All the waveforms are downsampled to 16k Hz in our experiments. We extract 2048-point STFT magnitudes with Hanning window and wrap the features with Mel filter to 80-band Mel spectrum. We use 5ms/20ms or 12.5ms/50ms window shift for different experiments. Then a log operation is applied to linear spectrum and Mel spectrum. Channel-wise mean-std normalization is applied to the log spectrums. We use dynamic mini-batch size in our experiments. The maximum mini-batch size is 32, if the training samples are too long to be fit in the GPU memory, the mini-batch size is divided by 2. So we can run each experiment on a single Nvidia Telsa M40 24GB GPU. We use repeat padding for the

¹https://www.data-baker.com/open_source.html

²http://www.festvox.org/festival/index.html

corpus	reduction factor	frame shift (ms)	
corpus		5	12.5
	1	100%	100%
LJSpeech	2	100%	16%
	5	22%	10%
	1	100%	55%
db-CSMSC	2	57%	17%
	5	8%	8%

Table 2: Sentence error rate (SER) for different reduction factor and frame shift configuration	e (SER) for different reduction factor and frame shift configurations. ⁶	5
---	---	---

training samples of different lengths in a batch since zero padding would affect the batch normalization statistics. We use the Adam [21] optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon = 10^{-8}$. The initial learning rate is 0.002 and starts to decay by a factor of $\sqrt{4000/step}$ from 4000 step³. The gradient is clipped to maximum global norm of 1.0 [25]. We use Tacotron2 for our experiments. For each dataset, we keep out 1/64 of the dataset as the validation set. For English, the test cases are randomly chosen from the 1132 CMU_ARCTIC⁴ sentences. For Mandarin Chinese, the test cases are chosen from text of different domains. We use a test set of 1000 sentences for error detection. Among these 1000 sentences, we randomly select 100 for listening test. The average numbers of words/characters and phonemes in one utterance is 8.8 and 32.1 for English and 15.6 and 41.4 for Chinese. In the listening test, we use an open-sourced WaveRNN vocoder⁵ to reconstruct waveforms from Mel spectrums.

5.2 Sentence error rate (SER) for different configurations

We train Tracotron2 for the different configurations listed in Table 1. The model is evaluated at 200k step. The SER for LJSpeech and db-CSMSC is recorded in Table 2. The errors are labeled by hand for the test set of 100 sentences. The model performs worse on LJSpeech than on db-CSMSC, mainly because LJSpeech is not a elaborate corpus for speech synthesis. SER negatively correlates with frame shift and reduction factor in Table 2. This result is consistent with the case reported in [28] that when using 5ms frame shift, significantly more pronunciation issues are observed. We also test dropping teacher forcing input frames for 12.5 ms frame shift and reduction factor 2. The dropout frame rate is set to 0.2. Each teacher forcing input frame has a probability of 0.2 to be set to the global mean. The SER is 15% for LJSpeech and 12% for db-CSMSC. It is slightly better than the SER of the same configuration in Table 2. Dropping the teacher forcing input frames increases the information gap between the teacher forcing input and acoustic target. We may conclude that increasing the MSE between the teacher forcing input and the acoustic target makes it hard for the the autoregressive component to predict the latter totally depending on the former. Then the model has to exploit the text more and to learn a strong correlation between the text and the predicted acoustic features. It would make less mistakes with the strong correlation at inference stage. In summary, modeling the correlation sufficiently is crucial for the robustness of Tacotron.

5.3 SER and MOS for Tacotron-MMI

In this part, we use 12.5ms frame shift and reduction window of size 2 for computing efficiency. In Tacotron2, the attention context is concatenated to the LSTM output and projected by a linear transform to predict the Mel spectrum. This means the predicated Mel spectrum contains linear components of the text information. If we use this Mel spectrum as the input to the CTC recognizer, the text information is too easily accessible for the recognizer. This may cause the text information to be encoded in a pathological way in the Mel spectrum and lead to a strict diagonal alignment map (one acoustic frame output for one phoneme input) combined with location-sensitive attention. So before the linear transform operation, we add an extra LSTM layer to mix the text information and acoustic information as depicted in Figure 1 and the output of the extra LSTM layer is linearly projected to predict the Mel spectrum. λ is set to 1.0 at the start of training and begins to increase at 40k step. It increases by 1.0 linearly at every 2k steps and stops increasing

³https://github.com/keithito/tacotron

⁴http://festvox.org/cmu_arctic/

⁵https://github.com/fatchord/WaveRNN

⁶100% SER indicates no intelligible waveform is synthesized.

corpus	error type	MMI + dfr 0.0	MMI + dfr 0.2
LISpeech	1000 detected	6.5%	4.6%
Liopecen	100 undetected	3%	0%
dh CSMSC	1000 detected	2.0%	1.8%

Table 3: SER for detected and undetected errors. dfr is short for dropout frame rate. 1000 detected is the automatically detected SER for the 1000-sized test set. 100 undetected is the SER labeled by a human not including the automatically detected errors for the 100-size test set. Sum of the 1000 detected and 100 undetected SER is the total SER.

Table 4: Mean opinion score (MOS) with 95% confidence intervals for different configurations.

3%

2%

100 undetected

	dfr 0.0	dfr 0.2	MMI + dfr 0.0	MMI + dfr 0.2
MOS	3.84±0.16	$3.92{\pm}0.17$	3.83±0.14	3.87±0.15

when it reaches 10.0. We select the model checkpoint, which achieves smaller average edit distance on the validation set, for evaluation from 100k to 300k training step.

The automatic error detection is experimented on the test set of 1000 sentences. All the waveforms, in which errors are detected, are listened by a human to check the correctness of the error detection. Then the 100-sized test set is also listened by a human to label errors. The results are recorded in Table 3. For db-CSMSC, the detected errors are consistent with the errors in the synthesized waveforms. The undetected errors are all unnatural pauses. Mandarin Chinese does not contain explicit word boundaries, so it is hard to predict the pauses. The recognition targets are all vocable and contain no pause information, hence the indicator can not detect pause errors. For LJSpeech, 34.8% of the waveforms, which are detected as erroneous, are labeled as correct by a human. Inspecting these inconsistent cases, we find that the recognizer is confused by pairs of phonemes that sound very similar, such as "UW", "UH" and "SH", "ZH"⁷. The undetected 3% errors for LJSpeech + dfr 0.0 (dropout frame rate 0.0) are phonemes of unnaturally short durations. Dropping teacher forcing input frames would make the model more robust, so we find no such undetected errors in LJSpeech + dfr 0.2 waveforms. For Mandarin Chinese, the indicator is reliable for mispronunciations, skipping or repeating words and incomplete or unstopped synthesis. For English, the mispronunciation detection is not reliable, but other error detections are still reliable. The indicator for English may be improved if a more powerful recognizer structure is used. We leave this for future study. The sum of the detected and undetected SER is better than the best result in Table 2, specifically, 3.0% vs 10% for LJSpeech (false detection removed) and 3.8% vs 8% for db-CSMSC. We may conclude that MMI can reduce the bad-case rate, since it strengthens the correlation between the text input and the predicted acoustic features during training. Audio and error detection samples are accessible online⁸.

We conduct a mean opinion score (MOS) test to see whether the extra MMI objective would degenerate the synthesized waveform quality or not. Only correctly synthesized waveforms are selected for this test. From Table 4, we can see that Tacotron2 with dfr 0.2 achieves the best perceptual result. MMI has a bit of negative effect on the perceptual performance. Since robustness is also very important for a speech synthesis model, we can accept a little drop in perceptual performance to avoid unacceptable bad cases.

6 Conclusion

In this paper we analyze why Tacotron is prone to bad cases. In short, modeling the correlation between the text and the acoustic features sufficiently is important to avoid the bad cases. To gain this objective, we propose to maximize the mutual information between the text and the predicted acoustic features with an auxiliary CTC recognizer. Experiment results show that our method can reduce the rate of bad cases. The output of CTC recognizer provides a reliable indicator to detect error in synthesized acoustic features. Besides our method can be trained in an end-to-end manner. It keeps the short pipeline of the original method. Moving forward, since we have an automatic error detection method, we can analyze the mistakes the model made and improve the model. For example, some stand-alone finals appear frequently in the detected errors in Mandarin Chinese. We may improve the phone set aiming at making the model more stable on such finals. This work sheds light on how to design a reliable end-to-end speech synthesis model.

⁷CMU Dictionary format http://www.speech.cs.cmu.edu/cgi-bin/cmudict

⁸https://drive.google.com/drive/folders/1f1CKcXNbd82ypUdXFfr_fJSKumVLoEy0



Figure 1: Modified part of Tacotron2 structure for CTC recognizer. The LSTM layer in the dashed line box is added to mix the text and acoustic information. \oplus is the concatenate operation. The text information is the text input processed by the text encoder. The acoustic information is the teacher forcing input processed by the decoder prenet.

Acknowledgments

The authors thank Chengzhu Yu, Heng Lu and Tianxiao Fu form Tencent AI Lab for their helpful discussions and advices.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, 2015.
- [2] L. Bahl, P. Brown, P. de Souza, and R. Mercer. Maximum mutual information estimation of hidden markov model parameters for speech recognition. In *ICASSP '86. IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 11, pages 49–52, April 1986.
- [3] David Barber and Felix V. Agakov. The IM algorithm: A variational approach to information maximization. In Advances in Neural Information Processing Systems 16 [Neural Information Processing Systems, NIPS 2003, December 8-13, 2003, Vancouver and Whistler, British Columbia, Canada], pages 201–208, 2003.
- [4] Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew M. Dai, Rafal Józefowicz, and Samy Bengio. Generating sentences from a continuous space. In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, *CoNLL 2016, Berlin, Germany, August 11-12, 2016*, pages 10–21, 2016.
- [5] Xi Chen, Yan Duan, Rein Houthooft, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain, pages 2172–2180, 2016.
- [6] Xi Chen, Diederik P. Kingma, Tim Salimans, Yan Duan, Prafulla Dhariwal, John Schulman, Ilya Sutskever, and Pieter Abbeel. Variational lossy autoencoder. In 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings, 2017.
- [7] Marvin Coto-Jiménez and John Goddard Close. Speech synthesis based on hidden markov models and deep learning. *Research in Computing Science*, 112:19–28, 2016.
- [8] Yuchen Fan, Yao Qian, Feng-Long Xie, and Frank K. Soong. TTS synthesis with bidirectional LSTM based recurrent neural networks. In INTERSPEECH 2014, 15th Annual Conference of the International Speech Communication Association, Singapore, September 14-18, 2014, pages 1964–1968, 2014.
- [9] Andrew Gibiansky, Sercan Ömer Arik, Gregory Frederick Diamos, John Miller, Kainan Peng, Wei Ping, Jonathan Raiman, and Yanqi Zhou. Deep voice 2: Multi-speaker neural text-to-speech. In Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA, pages 2966–2974, 2017.
- [10] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial networks. *CoRR*, abs/1406.2661, 2014.
- [11] Alex Graves. Sequence transduction with recurrent neural networks. CoRR, abs/1211.3711, 2012.
- [12] Alex Graves, Santiago Fernández, Faustino J. Gomez, and Jürgen Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006), Pittsburgh, Pennsylvania, USA, June 25-29, 2006*, pages 369–376, 2006.
- [13] Ishaan Gulrajani, Kundan Kumar, Faruk Ahmed, Adrien Ali Taïga, Francesco Visin, David Vázquez, and Aaron C. Courville. Pixelvae: A latent variable model for natural images. In 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings, 2017.
- [14] Haohan Guo, Frank K. Soong, Lei He, and Lei Xie. A new gan-based end-to-end tts training algorithm. *CoRR*, abs/1904.04775, 2019.
- [15] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. Neural Computation, 9(8):1735–1780, 1997.
- [16] Keith Ito. The lj speech dataset. https://keithito.com/LJ-Speech-Dataset/, 2017.
- [17] Ye Jia, Yu Zhang, Ron J. Weiss, Quan Wang, Jonathan Shen, Fei Ren, Zhifeng Chen, Patrick Nguyen, Ruoming Pang, Ignacio Lopez-Moreno, and Yonghui Wu. Transfer learning from speaker verification to multispeaker text-to-speech synthesis. In Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada., pages 4485–4495, 2018.
- [18] Nal Kalchbrenner, Erich Elsen, Karen Simonyan, Seb Noury, Norman Casagrande, Edward Lockhart, Florian Stimberg, Aäron van den Oord, Sander Dieleman, and Koray Kavukcuoglu. Efficient neural audio synthesis. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, pages 2415–2424, 2018.
- [19] Shiyin Kang and Helen M. Meng. Statistical parametric speech synthesis using weighted multi-distribution deep belief network. In INTERSPEECH 2014, 15th Annual Conference of the International Speech Communication Association, Singapore, September 14-18, 2014, pages 1959–1963, 2014.
- [20] Yoon Kim, Sam Wiseman, Andrew C. Miller, David Sontag, and Alexander M. Rush. Semi-amortized variational autoencoders. In Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018, pages 2683–2692, 2018.

- [21] Diederik P. Kingma and Jimmy Ba. 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, 2015.
- [22] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings, 2014.
- [23] Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. A diversity-promoting objective function for neural conversation models. In NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016, pages 110–119, 2016.
- [24] Naihan Li, Shujie Liu, Yanqing Liu, Sheng Zhao, Ming Liu, and Ming Zhou. Close to human quality TTS with transformer. CoRR, abs/1809.08895, 2018.
- [25] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, pages 1310–1318, 2013.
- [26] Marc'Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. Sequence level training with recurrent neural networks. In 4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings, 2016.
- [27] Shiv Shankar and Sunita Sarawagi. Posterior attention models for sequence to sequence learning. In *International Conference* on Learning Representations, 2019.
- [28] 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. 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, 2018.
- [29] R. J. Skerry-Ryan, Eric Battenberg, Ying Xiao, Yuxuan Wang, Daisy Stanton, Joel Shor, Ron J. Weiss, Rob Clark, and Rif A. Saurous. Towards end-to-end prosody transfer for expressive speech synthesis with tacotron. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, pages 4700–4709, 2018.
- [30] Hideyuki Tachibana, Katsuya Uenoyama, and Shunsuke Aihara. Efficiently trainable text-to-speech system based on deep convolutional networks with guided attention. In 2018 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2018, Calgary, AB, Canada, April 15-20, 2018, pages 4784–4788, 2018.
- [31] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W. Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. In *The 9th ISCA Speech Synthesis Workshop, Sunnyvale, CA, USA, 13-15 September 2016*, page 125, 2016.
- [32] Aäron van den Oord, Yazhe Li, Igor Babuschkin, Karen Simonyan, Oriol Vinyals, Koray Kavukcuoglu, George van den Driessche, Edward Lockhart, Luis C. Cobo, Florian Stimberg, Norman Casagrande, Dominik Grewe, Seb Noury, Sander Dieleman, Erich Elsen, Nal Kalchbrenner, Heiga Zen, Alex Graves, Helen King, Tom Walters, Dan Belov, and Demis Hassabis. Parallel wavenet: Fast high-fidelity speech synthesis. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, pages 3915–3923, 2018.
- [33] Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA, pages 6309–6318, 2017.
- [34] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA, pages 6000–6010, 2017.
- [35] Yuxuan Wang, R. J. Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J. Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, Quoc V. Le, Yannis Agiomyrgiannakis, Rob Clark, and Rif A. Saurous. Tacotron: A fully end-to-end text-to-speech synthesis model. *CoRR*, abs/1703.10135, 2017.
- [36] Yuxuan Wang, Daisy Stanton, Yu Zhang, R. J. Skerry-Ryan, Eric Battenberg, Joel Shor, Ying Xiao, Ye Jia, Fei Ren, and Rif A. Saurous. Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, pages 5167–5176, 2018.
- [37] Ronald J. Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural Computation*, 1(2):270–280, 1989.
- [38] Xixin Wu, Shiyin Kang, Lifa Sun, Yishuang Ning, Zhiyong Wu, and Helen Meng. Attention-based recurrent generator with gaussian tolerance for statistical parametric speech synthesis. In *The Affective Social Multimedia Computing (ASMMC)*, 2017.
- [39] Heiga Zen, Andrew W. Senior, and Mike Schuster. Statistical parametric speech synthesis using deep neural networks. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2013, Vancouver, BC, Canada, May 26-31, 2013*, pages 7962–7966, 2013.

- [40] Shengjia Zhao, Jiaming Song, and Stefano Ermon. Infovae: Information maximizing variational autoencoders. *CoRR*, abs/1706.02262, 2017.
- [41] Chunting Zhou and Graham Neubig. Multi-space variational encoder-decoders for semi-supervised labeled sequence transduction. In Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers, pages 310–320, 2017.
- [42] Xiaolian Zhu, Yuchao Zhang, Shan Yang, Liumeng Xue, and Lei Xie. Pre-alignment guided attention for improving training efficiency and model stability in end-to-end speech synthesis. *IEEE Access*, pages 1–1, 2019.