

On the Problem of Text-To-Speech Model Selection for Synthetic Data Generation in Automatic Speech Recognition

Nick Rossenbach^{1,2,3,*}, Ralf Schlueter^{2,3}, Sakriani Sakti^{1,4}

¹Japan Advanced Institute of Science and Technology, Japan

²RWTH Aachen University, Germany; ³AppTek GmbH, Germany

⁴Nara Institute of Science and Technology, Japan

rossenbach@cs.rwth-aachen.de, ssakti@is.naist.jp, schlueter@cs.rwth-aachen.de

Abstract

The rapid development of neural text-to-speech (TTS) systems enabled its usage in other areas of natural language processing such as automatic speech recognition (ASR) or spoken language translation (SLT). Due to the large number of different TTS architectures and their extensions, selecting which TTS systems to use for synthetic data creation is not an easy task. We use the comparison of five different TTS decoder architectures in the scope of synthetic data generation to show the impact on CTC-based speech recognition training. We compare the recognition results to computable metrics like NISQA MOS and intelligibility, finding that there are no clear relations to the ASR performance. We also observe that for data generation autoregressive decoding performs better than non-autoregressive decoding, and propose an approach to quantify TTS generalization capabilities.

Index Terms: speech recognition, synthetic data generation, text-to-speech

1. Introduction

The usage of synthetic data from text-to-speech (TTS) systems in the context of automatic speech recognition (ASR) has been explored in many different ways. This ranges from simply generating new data [1, 2] to jointly training ASR and TTS systems [3]. Prior work improved the integration of the synthetic data into the ASR systems, for example by different weighting of synthetic inputs [4], or the usage of internal ASR model statistics to reject synthetic samples not matching the real data [5]. As most publications aim to introduce new methodologies, in each case they usually use one specific TTS architecture and one specific ASR architecture. The number of fundamentally different ASR architectures is rather limited and roughly categorized into: Hybrid deep neural network hidden Markov model (DNN-HMM) [6], connectionist-temporal-classification (CTC) [7], Transducer [8] and Attention-Encoder-Decoder (AED) [9]. For TTS there is a vast amount of different architectures and models, especially when including different approaches of multi-speaker modeling [10]. From the perspective of working on ASR this makes it very difficult to choose which TTS system to use for synthetic data generation. In addition, human evaluations are not suited to evaluate data for training purposes, as conducting the evaluation usually takes more time than running the training itself. Then, evaluation for naturalness and human subjective evaluations alone can not be expected to be a sufficient criterion for utility estimation of synthetic data. Instead, aspects like speaker diversity, temporal diversity or expressiveness in general are assumed to be relevant

to yield good results by integrating synthetic data into the ASR process [11, 12]. Unfortunately, those aspects are difficult to quantify. Furthermore, LJSpeech [13] is often the only common dataset to compare new TTS architectures, which can not be expected to give representative results for the multi-speaker case [14]. This is especially relevant as one might not have training data suited for TTS training at hand for a particular domain or language, but instead has to use the ASR data.

Motivated by those problems, we present an experiment pipeline to systematically analyze the relation of TTS generated data to ASR training. For this we compare different decoder architectures, which are based on: Non-autoregressive (NAR) Transformer similar to FastSpeech-2 [15], autoregressive (AR) long-short term memory (LSTM) similar to Non-attentive Tacotron [16], Glow-TTS [17] as flow-based decoder and Grad-TTS [18] as diffusion based decoder. We keep the rest of the TTS systems as simple as possible, using a Transformer encoder and convolution based duration prediction network, as used in [15, 17, 19]. Given the large amount of possible TTS architectures, we choose these systems for multiple reasons: First, they have a widespread recognition in the scientific community. Then, they already share similar encoder and up-sampling concepts, sometimes even using the same code, which makes comparison easier. Finally, they do not require training with an integrated vocoder model (as e.g. in VITS) which would make a direct comparison more difficult. While comparisons of different TTS architectures with respect to naturalness and expressiveness exist [20], this is to our knowledge the first work covering a broader range of different TTS output characteristics and quality levels targeted for ASR training. Also, we perform a direct comparison of NAR vs AR decoding for data generation with a consistent sequence mapping method. Lastly, we present a way of measuring generalization of the TTS systems with respect to ASR training data generation by using different conditions for synthesizing the training data. The Sisyphus [21] based setup and additional materials are available on Github¹.

2. Text-to-speech systems

2.1. General architecture

In the design of the TTS systems we follow what we found to be most prevalent in recent literature. After the introduction of Transformer blocks with convolution for TTS [19] and the establishing of duration based state up-sampling [22], many systems continued using these components. Thus, for each system we use the same layout of a Transformer based-encoder, a convolution based duration predictor and a discrete encoder state

*Work done while being an internship student at JAIST

¹https://github.com/rwth-i6/i6_experiments/tree/main/users/rossenbach/experiments/jaist_project

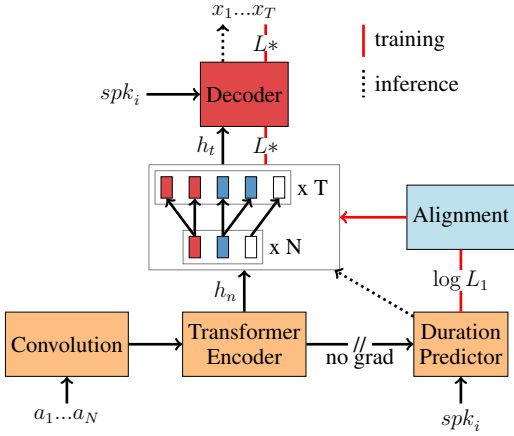


Figure 1: The TTS architecture. We change the decoder marked in red for each of the different systems.

up-sampling method. Only the decoder, which converts the up-sampled encoder states h_1^T into the final log-Mel spectrograms x_1^T , is changed. Figure 1 shows the TTS architecture used in this work. We do not use explicit f_0 or energy prediction. The overall structure follows exactly Glow-TTS, in which many parts were taken from TransformerTTS [19] or FastSpeech-2. We also used the code of the official Glow-TTS repository as reference for our implementation.

2.2. Decoder variants

A large amount of proposed TTS decoder architectures can be categorized into three groups: Direct prediction, flow and diffusion. Direct prediction means that the output of the TTS neural network is a log-Mel spectrogram which is trained to directly match the target x_1^T . Common losses are mean absolute error or mean squared error. The generation is a single forward pass through the decoder network g_θ , which can be written as:

$$\hat{x}_1^T = g_\theta(h_1^T) \quad (1)$$

The generation can also be done in an auto-regressive manner:

$$\hat{x}_t = g_\theta(h_1^t, \hat{x}_1^{t-1}) \quad (2)$$

For flow-based TTS systems, the decoder is an invertible neural network that can be applied in both directions from the spectrograms x_1^T into latent space variables z_1^T and vice-versa. During training, the target features are transformed into the latent space and matched against a normal distribution with a mean $\mu_1^T = \mu_\theta(h_1^T)$ and unit variance. The target objective is to maximize the likelihood of the latent space samples given the distributions. For inference, the latent space values are sampled from the distribution with a mean function μ_θ , temperature τ , and passed backwards through the flow f_θ :

$$\hat{z}_1^T \sim \mathcal{N}(\mu_\theta(h_1^T), \tau \mathbf{1}) \quad \text{and} \quad \hat{x}_1^T = f_\theta^{-1}(\hat{z}_1^T) \quad (3)$$

For the diffusion-based TTS system we follow the Grad-TTS architecture. A normal distribution as in the flow-based approach is used to sample a noisy output spectrogram \hat{X}_I with $X = x_1^T$. The noise spectrogram is cleaned in I steps using a differential equation containing the solver network s_θ and pre-defined noise schedule β_i to reach the final output \hat{X}_0 :

$$\hat{X}_I \sim \mathcal{N}(\mu_\theta(h_1^T), \tau \mathbf{1}) \quad (4)$$

$$\hat{X}_{i-1} = \hat{X}_i - \frac{1}{2I} \left(\mu_\theta(h_1^T) - \hat{X}_i - s_\theta(h_1^T, \hat{X}_i) \right) \beta_i \quad (5)$$

For each of the decoder principles we implemented one model for which the details will be presented later in Section 4.3.

Table 1: TTS model parameter count, training time on Nvidia A40 GPU and inference real-time-factor on AMD Epyc 7H12 CPU using 2 physical cores and including Griffin & Lim.

TTS Decoder	# params	training time	CPU RTF
Transformer	24.4M	70h	0.069
NAR LSTM	29.3M	27h	0.025
AR LSTM	33.6M	94h	0.028
Glow-TTS	57.3M	53h	0.036
Grad-TTS	25.5M	61h	0.866

3. Evaluation pipeline and metrics

To evaluate the TTS systems in the context of ASR, we train a CTC-based ASR system purely on data generated by TTS. We generate the same amount of data as used for the baseline training, and measure the performance of the ASR systems in word error rate (WER) for recognizing real human speech. In the optimal case, the performance of the ASR system trained on the synthetic data would match the performance of the system trained on real data. In order to evaluate the generalization capabilities of TTS towards new conditions, we generate the synthetic data in 3 ways: a) using exactly the same text and speaker assignment as in training, b) using the same text but shuffling the speakers and finally c) using the same amount of text from a different corpus in the same domain. That way we can measure how good the TTS systems are at generalizing to new conditions different from the training. The main evaluation criterion to measure the utility of the TTS-generated data is the WER of the trained ASR systems calculated on the official LibriSpeech [23] test sets. In addition, we use metrics directly on the synthetic data to evaluate their relation to the final WER. We use NISQA [24] as a tool to generate automatic mean-opinion-score (MOS) values. While computer generated MOS values are not a replacement for subjective evaluations, they can be used as tool to determine naturalness as possible indicator for the ASR performance. Subjective evaluation itself can not be useful to evaluate synthetic data in a development process, as doing the evaluation often takes more effort than simply training the ASR system. We also use a larger ASR model trained on the full LibriSpeech corpus to recognize the synthesized cross-validation data to check for completeness and pronunciation correctness. We refer to this metric as synthetic WER (sWER). In the literature this is often called *intelligibility* metric, and sometimes character error rate is used instead of WER, e.g. in [25].

4. Experiments

4.1. Data

We use the LibriSpeech train-clean-100 as baseline data to train the TTS and ASR models, and the transcriptions of train-clean-360 as additional text data. This is a common setup that has been used extensively in previous literature [1, 2, 4]. We subsample the exact same amount of utterances, roughly 28k, from train-clean-360 to have the same amount of unseen text as text in the training data. We take 4 utterances of each of the 251 speakers in train-clean-100 out of the TTS training data to have a cross-validation set to report MOS and sWER. We perform no further pre-processing of the data, unlike other work [2, 11]. I.e. no excessive silence is removed, and no speakers are discarded. We use the official LibriSpeech lexicon for word-to-phoneme conversion and use Sequitur [26] to generate phonemes for words that are not in the lexicon. For ASR, we augment the phoneme set with end-of-word markers as presented in [27].

4.2. Automatic speech recognition

As ASR system we use a CTC model with a Conformer-based [28] encoder with 12 blocks, a base dimension of 384 and a convolutional frontend with a sub-sampling factor of 4. For recognition we use the lexicon-constrained Flashlight [29] CTC decoder via the available Torchaudio interface. The search process includes the official LibriSpeech 4-gram LM model. Since it has been shown that synthetic data generation is successful mostly for AED-ASR [2], we chose the context-free phoneme-based CTC-ASR specifically for the reason that it should be influenced less by changing or adding new text.

4.3. Training and model settings

All models are trained using 80-dimensional log-Mel spectrogram features, for TTS with 12.5ms shift and for ASR with 10ms shift and SpecAugment [30] applied. The TTS models are trained with Adam [31] for 400 epochs. For all models except Grad-TTS we use linear learning rate (LR) scheduling, which we found to perform better than inverse square reduction. We increase the LR from $5e-5$ to $5e-4$ for the first 100 epochs, and reduce linearly from $5e-4$ to $5e-7$. For Grad-TTS we use a constant LR of $1e-4$. The ASR models are trained using AdamW with weight-decay $1e-3$ for roughly 80 epochs and a maximum learning rate of $7e-4$. We generate an alignment by training the Glow-TTS system once using implicit alignment search. In each TTS training the fixed alignment is used for up-sampling and as target for the duration predictor. For vocoding, we convert the log-Mel to linear features using a separate BLSTM-based network and apply the Librosa Griffin & Lim (G&L) [32] function with 32 iterations and momentum of 0.99.

4.3.1. TTS parameter settings

For the different decoder architectures the number of parameters, training time and inference time differ, as shown in Table 1. Making a fair comparison is difficult, as for each system we would need to do extensive tuning experiments until we reach the optimal performance. Thus, we stick mostly to the sizes presented in previous publications. Due to the high cost of subjectively evaluating TTS systems, ablation studies on the model sizes are usually not presented in scientific publications.

Base model: For the TTS Transformer encoder we use 6 layers with an internal dimension of 256, preceded by 3 convolution layers with 256 channels and kernel size 5. The duration prediction network uses 2 convolutions with kernel size 3 and 384 channels.

Transformer decoder: For the Transformer-based decoder we simply re-use the same structure as for the encoder Transformer, having 6 layers with an internal base dimension of 256. The system is trained using L1 loss on the output spectrograms.

Autoregressive LSTM decoder: For the AR-LSTM decoder we are using 2 layers of 1024-dimensional Zoneout-regularized LSTM layers. We use a reduction factor of 2 to predict two feature frames per decoder step. The pre-net for the AR feedback loop and the convolutional NAR post-net are derived from [16]. We apply L1 loss for both the decoder output and the combined prediction with the post-net.

Non-autoregressive LSTM decoder: This decoder is similar to the AR decoder, but we remove the AR feedback loop network, and replace the uni-directional Zoneout LSTM cell by a normal bi-directional LSTM with dropout. Each LSTM direction has 512 dimensions. For Transformer, AR- and NAR-LSTM, a 256-dimensional speaker embedding vector is con-

Table 2: Evaluation of the different TTS decoder architectures. We use 28k sequences from train-clean-360 for the synthetic training data generation and the cross-validation subset to evaluate NISQA MOS and the sWER. The reference values are the training / evaluation on the real audio. For MOS we include the 95% confidence interval using the bootstrap method from [33].

TTS Decoder	Syn. data ASR		NISQA MOS \uparrow	sWER [%] \downarrow
	test [WER %] \downarrow			
	clean	other		
Transformer	11.8	33.7	2.74 ± 0.05	1.6
NAR LSTM	10.2	30.5	2.87 ± 0.05	1.7
AR LSTM	9.4	28.4	3.08 ± 0.05	2.5
Glow-TTS	13.8	34.9	3.33 ± 0.04	10.8
Grad-TTS	11.5	30.0	2.47 ± 0.05	15.1
Reference	5.1	14.9	4.06 ± 0.07	1.6

catenated to each upsampled encoder state.

Glow-TTS decoder: For a flow-based decoder, including speaker conditioning, we follow exactly the architecture of Glow-TTS [17]. The decoder consists of 12 invertible coupling blocks with an internal dimension of 256. We merge each 2 consecutive frames for a latent-space dimension of 160.

Grad-TTS decoder: For a diffusion-based model we follow the architecture and settings of Grad-TTS [18]. As the original architecture is not designed for multi-speaker TTS, modifications were necessary for convergence. We added two additional Transformer layers on top of the 6 encoder layers which are conditioned on the speaker embedding. That way, we have a speaker-conditioned mean prediction network. We also used the same speaker embeddings to condition the u-net, which is used for the differential equation solving. We found $\tau = 0.7$ to be the optimal noise scale for both Glow-TTS and Grad-TTS.

5. Results

5.1. Synthetic data, automatic MOS and recognizability

In a first comparison we evaluated the systems via the synthetic data only ASR training, the NISQA MOS and the sWER. Table 2 shows the results of each of the five decoders compared to the reference values of the real data. For all metrics we used text data unseen during training. With respect to the final ASR training evaluation, all TTS systems have a reasonable performance and there is no strong outlier. In general, using the synthetic data for training results in a system with double the error rate when training TTS only on the available ASR data, which is a consistent result to previous literature [5, 11]. [34] showed that even when using vast amounts of training data to create strong foundation-model sized TTS models, the generated synthetic data seems to be not yet on par with real data. The AR-LSTM decoder produces the best synthetic training data, which is unexpected as the Glow-TTS and Grad-TTS should produce more expressive and variable audio due to their internal sampling process. Also, Glow-TTS is the only system that exactly follows the given literature, while the other systems have small adaptations or are simplifications of the original model architecture. The third and fourth column in Table 2 show the NISQA MOS and the sWER on the TTS cross validation set. It is visible that there is no direct relation of the two metrics compared to the final ASR performance. For example, when comparing Grad-TTS to the Glow-TTS and Transformer decoder, there is a lower MOS and a much higher sWER, but still the final system performance is better. This means both metrics are not reliable to determine which system to pick for data generation.




Figure 2: Example spectrograms of a selected sequence from the cross validation set. Used TTS decoder from top to bottom: a) Transformer b) AR-LSTM c) Glow-TTS d) Grad-TTS

5.2. Decoder characteristics

In Figure 2 we show a selected sequence synthesized by four different decoders. The Transformer decoder spectrograms are smooth, lacking a lot of detail in the high-frequency area. The AR-LSTM spectrogram in comparison has much more details. The spectrogram from Glow-TTS shows visible noise resulting from the latent space sampling. It also has less distinct boundaries between the phonemes, but rather fuzzy transitions, which might be an explanation for the worse recognizability. The Grad-TTS has less distorted spectrograms than Glow-TTS, but we found that for very short sequences the output is sometimes just a noisy grid pattern. Also the voice lines become very blurred together for lower frequencies, causing a very unnatural sound, which matches with the low automatic MOS score. Thus we assume, that the NISQA MOS and the sWER are very sensitive to certain artifacts that have less impact in on the actual ASR training. One can see that all four TTS decoders produce spectrograms with different characteristics, even with other components being identical. Thus, it is important to not only consider a single TTS architecture when trying to show a general relation between the ASR performance and metrics calculated on the level of the TTS system or the synthetic data.

5.3. Generalization

Table 3 shows how good the TTS systems generalize with respect to the synthetic data generation task. Recall that condition a) means to exactly reproduce the training data, so with perfect memorization during training it would be possible to produce identically performing data. It is visible that across all models the values for a) and c) are mostly not too far from each other, so the generalization abilities are high. This gives a hint that current models are under-fitting and much larger models might yield better results, even on very constrained datasets such as LibriSpeech 100h. For synthesizing new text, so comparing b) to c), most models show no or only a small change. Only for the AR LSTM and Glow-TTS decoder there is meaningful degradation in the WER, mostly for the "other" test. For the Glow-TTS

Table 3: Comparison of the different TTS decoder architectures regarding overfitting conditions by synthesizing data with a) same text and same speaker as in training b) same text and shuffled speakers c) new text of same amount.

TTS Decoder	test [WER %]↓					
	clean			other		
	a)	b)	c)	a)	b)	c)
Transformer	11.1	11.8	11.8	32.5	33.5	33.7
NAR LSTM	9.7	10.1	10.2	30.0	30.2	30.5
AR LSTM	8.8	9.1	9.4	26.3	27.1	28.4
Glow-TTS	12.6	13.3	13.8	31.0	32.9	34.9
Grad-TTS	11.0	11.5	11.5	28.0	30.1	30.0
Reference	5.1			14.9		

a possible reason could be that it suffers from pronunciation issues, which get even stronger when synthesizing new text. Testing for the generalization capabilities of TTS this way might not be used only for the comparison of different models like in this work, but allow for easier hyperparameter tuning which is usually not done extensively in TTS literature.

5.4. Discussion

When training with synthetic data only, we achieved a WER performance which is less than factor 2 worse compared to training on the real data. In relation, [5] reported reaching ratios of around 2 on the LibriSpeech-960 task, and [11] reported ratios of more than 3. In this work we specifically investigated decoder architectures. In future work this gap can be reduced by extending the results of this work with extensive hyperparameter tuning as well as prosody [11] and more sophisticated duration modeling [12]. We have seen that NISQA MOS and the sWER are not reliable to determine the usefulness of the TTS system for ASR data generation. This means there are more and/or different metrics needed to give accurate estimates about how good the synthetic data can be utilized for ASR training. While in our work synthesizing new data and training the ASR systems was done in less than a day, this becomes much more of a problem when training large models on magnitudes more of data. Thus, we think it is important to find ways to estimate the capabilities of TTS models before starting large-scale data generation.

6. Conclusion

In this work we evaluated different TTS decoder architectures w.r.t. their ability to generate synthetic data for ASR training. We compared five different TTS decoder architectures and found that for our setup the AR-LSTM-based system with L1 loss performed best in creating training data for ASR. We were able to create synthetic data that performs less than a factor of 2 worse than real data while training the TTS only on the same data as available for the ASR training. We observed that the systems' outputs differ from each other in many aspects such as naturalness or pronunciation accuracy. Still, the metrics we used to capture them showed a surprisingly low correlation to the performance of the ASR systems trained on such outputs. This indicates that better metrics have to be found in order to do a meaningful evaluation of TTS systems before using it to create synthetic training data. We also introduced a way to measure the generalization capabilities of TTS systems for data generation, which can be used to optimize TTS towards being able to generate data that reaches parity with real data.

7. Acknowledgements

This work was partially supported by NeuroSys, which as part of the initiative “Clusters4Future” is funded by the Federal Ministry of Education and Research BMBF (03ZU1106DA), and by the project RESCALE within the program *AI Lighthouse Projects for the Environment, Climate, Nature and Resources* funded by the Federal Ministry for the Environment, Nature Conservation, Nuclear Safety and Consumer Protection (BMUV), funding ID: 67KI32006A. Part of this work is supported by JSPS KAKENHI Grant Numbers JP21H05054 and JP23K21681. The authors thank Lukas Rilling, Peter Vieting and Benedikt Hilmes for their additional input and feedback.

8. References

- [1] A. Laptev, R. Korostik, A. Svischev, A. Andrusenko, I. Medenikov, and S. Rybin, “You do not need more data: Improving end-to-end speech recognition by text-to-speech data augmentation,” *13th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*, October 2020.
- [2] N. Rossenbach, M. Zeineldeen, B. Hilmes, R. Schlüter, and H. Ney, “Comparing the benefit of synthetic training data for various automatic speech recognition architectures,” in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2021, pp. 788–795.
- [3] A. Tjandra, S. Sakti, and S. Nakamura, “End-to-end feedback loss in speech chain framework via straight-through estimator,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 6281–6285.
- [4] M. K. Baskar, L. Burget, S. Watanabe, R. F. Astudillo, and J. H. Cernocky, “Eat: Enhanced ASR-TTS for self-supervised speech recognition,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, June 2021.
- [5] T.-Y. Hu, M. Armandpour, A. Shrivastava, J.-H. R. Chang, H. Koppula, and O. Tuzel, “SYNT++: Utilizing Imperfect Synthetic Data to Improve Speech Recognition,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2022.
- [6] H. A. Bourlard and N. Morgan, *Connectionist speech recognition: a hybrid approach*. Springer Science & Business Media, 2012, vol. 247.
- [7] A. Graves, S. Fernández, and F. Gomez, “Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks,” in *In Proceedings of the International Conference on Machine Learning, ICML*, 2006, pp. 369–376.
- [8] A. Graves, “Sequence transduction with recurrent neural networks,” *ArXiv*, vol. abs/1211.3711, 2012.
- [9] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 2016, pp. 4960–4964.
- [10] X. Tan, T. Qin, F. K. Soong, and T.-Y. Liu, “A survey on neural speech synthesis,” *ArXiv*, vol. abs/2106.15561, 2021.
- [11] C. Minixhofer, O. Klejch, and P. Bell, “Evaluating and reducing the distance between synthetic and real speech distributions,” in *Proc. Interspeech*, 2023, pp. 2078–2082.
- [12] N. Rossenbach, B. Hilmes, and R. Schlüter, “On the relevance of phoneme duration variability of synthesized training data for automatic speech recognition,” in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2023, pp. 1–8.
- [13] K. Ito, “The ljspeech dataset.” [Online]. Available: <https://keithito.com/LJ-Speech-Dataset/>
- [14] Y. Wu, X. Tan, B. Li, L. He, S. Zhao, R. Song, T. Qin, and T.-Y. Liu, “AdaSpeech 4: Adaptive Text to Speech in Zero-Shot Scenarios,” in *Proc. Interspeech*, 2022, pp. 2568–2572.
- [15] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, “FastSpeech 2: Fast and High-Quality End-to-End Text to Speech,” in *International Conference on Learning Representations (ICLR)*, December 2021.
- [16] J. Shen, Y. Jia, M. Chrzanowski, Y. Zhang, I. Elias, H. Zen, and Y. Wu, “Non-attentive tacotron: Robust and controllable neural TTS synthesis including unsupervised duration modeling,” *ArXiv*, vol. abs/2010.04301v3, October 2020.
- [17] J. Kim, S. Kim, J. Kong, and S. Yoon, “Glow-TTS: A Generative Flow for Text-to-Speech via Monotonic Alignment Search,” in *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc., 2020, pp. 8067–8077.
- [18] V. Popov, I. Vovk, V. Gogoryan, T. Sadekova, and M. Kudinov, “Grad-TTS: A Diffusion Probabilistic Model for Text-to-Speech,” in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 139. PMLR, Jul. 2021, pp. 8599–8608.
- [19] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu, “Neural Speech Synthesis with Transformer Network,” in *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, ser. AAAI’19/IAAI’19/EAAI’19. AAAI Press, 2019.
- [20] G. Zhang, T. Merritt, S. Ribeiro, B. Tura-Vecino, K. Yanagisawa, K. Pokora, A. Ezzer, S. Cygert, A. Abbas, P. Bilinski, R. Barra-Chicote, D. Korzekwa, and J. Lorenzo-Trueba, “Comparing normalizing flows and diffusion models for prosody and acoustic modelling in text-to-speech,” in *Proc. Interspeech*, 2023, pp. 27–31.
- [21] J. Peter, E. Beck, and H. Ney, “Sisyphus, a workflow manager designed for machine translation and automatic speech recognition,” in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP*, pp. 84–89.
- [22] Y. Ren, Y. Ruan, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, “FastSpeech: Fast, Robust and Controllable Text to Speech,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [23] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An asr corpus based on public domain audio books.” in *ICASSP*. IEEE, 2015, pp. 5206–5210.
- [24] G. Mittag and S. Möller, “Deep Learning Based Assessment of Synthetic Speech Naturalness,” in *Proc. Interspeech*, 2020, pp. 1748–1752.
- [25] T. Hayashi, R. Yamamoto, K. Inoue, T. Yoshimura, S. Watanabe, T. Toda, K. Takeda, Y. Zhang, and X. Tan, “Espnet-tts: Unified, reproducible, and integratable open source end-to-end text-to-speech toolkit,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7654–7658.
- [26] M. Bisani and H. Ney, “Joint-sequence models for grapheme-to-phoneme conversion,” *Speech Communication*, vol. 50, no. 5, pp. 434–451, May 2008.
- [27] W. Zhou, S. Berger, R. Schlüter, and H. Ney, “Phoneme based neural transducer for large vocabulary speech recognition,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 5644–5648.
- [28] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, “Conformer: Convolution-augmented Transformer for Speech Recognition,” in *Proc. Interspeech*, 2020, pp. 5036–5040.
- [29] J. D. Kahn, V. Pratap, T. Likhomanenko, Q. Xu, A. Hannun, J. Cai, P. Tomasello, A. Lee, E. Grave, G. Avidov, B. Steiner, V. Liptchinsky, G. Synnaeve, and R. Collobert, “Flashlight: Enabling innovation in tools for machine learning,” in *Proceedings of the 39th International Conference on Machine Learning*, vol. 162, 17–23 Jul 2022, pp. 10557–10574.
- [30] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition,” in *Proc. Interspeech*, 2019, pp. 2613–2617.
- [31] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *3rd International Conference on Learning Representations, ICLR*, 2015.
- [32] D. W. Griffin, D. S. Deadrick, and J. S. Lim, “Speech synthesis from short-time fourier transform magnitude and its application to speech processing,” in *ICASSP*, 1984, pp. 61–64.
- [33] L. Ferrer and P. Riera, “Confidence intervals for evaluation in machine learning.” [Online]. Available: <https://github.com/luferrer/ConfidenceIntervals>
- [34] Seed-Team and ByteDance, “Seed-TTS: A family of high-quality versatile speech generation models,” *ArXiv*, vol. abs/2406.02430, 2024.