

Mix and Match: An Empirical Study on Training Corpus Composition for Polyglot Text-To-Speech (TTS)

Ziyao Zhang¹, Alessio Falai¹, Ariadna Sanchez¹, Orazio Angelini¹, Kayoko Yanagisawa¹

¹Alexa AI, Amazon

{zhaziyao, falai, cariadn, aorazio, yakayoko}@amazon.com

Abstract

Training multilingual Neural Text-To-Speech (NTTS) models using only monolingual corpora has emerged as a popular way for building voice cloning based Polyglot NTTS systems. In order to train these models, it is essential to understand how the composition of the training corpora affects the quality of multilingual speech synthesis. In this context, it is common to hear questions such as “Would including more Spanish data help my Italian synthesis, given the closeness of both languages?”. Unfortunately, we found existing literature on the topic lacking in completeness in this regard. In the present work, we conduct an extensive ablation study aimed at understanding how various factors of the training corpora, such as language family affiliation, gender composition, and the number of speakers, contribute to the quality of Polyglot synthesis. Our findings include the observation that female speaker data are preferred in most scenarios, and that it is not always beneficial to have more speakers from the target language variant in the training corpus. The findings herein are informative for the process of data procurement and corpora building.

Index Terms: neural text-to-speech (NTTS), multilingual speech synthesis, polyglot TTS

1. Introduction

Attempts to go beyond building single-speaker, monolingual Text-To-Speech (TTS) systems date far back in time [1]. The advancements in Neural TTS (NTTS) have seen the emergence of TTS models which are able to deliver high-quality multilingual and multi-speaker speech synthesis [2, 3, 4]. These multilingual TTS models differ with each other significantly in terms of the level of flexibility of speaker and language variant (LV) hereafter, which is defined as local variants of a language) combinations that they are able to deliver. In this paper, we define a multi-speaker multilingual TTS system, which enables any speaker in the training corpus to speak any LV present in the same corpus, as Polyglot TTS.

Multilingual TTS systems can generally be categorised into two realms, depending on whether cross-lingual voice cloning [2], defined as converting a certain speaker’s voice into speaking a new language, is attempted. In this paper, we focus on the paradigm where multilingual TTS is realised via voice cloning. In such a scenario, model training can be very challenging, because of the requirement of speaker and language disentanglement [5]. Although it has been shown [6] that having “polyglot” training corpora, which include recordings of the same speaker in multiple languages, improves multilingual synthesis, it is rarely practical to procure such data and the approach is unscalable. In light of this, the most realistic solution for building a voice cloning-based multilingual NTTS system is to train models using training corpora consisting of multiple monolingual speakers, which is the scenario we consider in this paper.

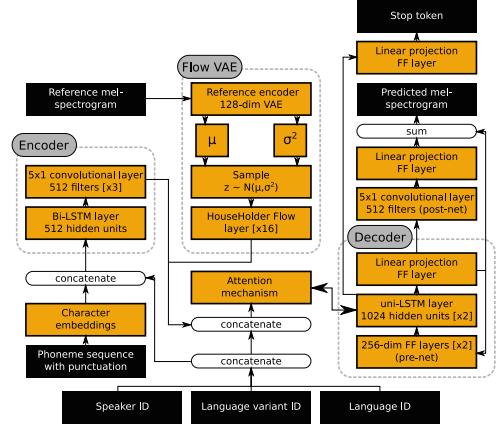


Figure 1: The model architecture used in all our experiments.

There have been multiple studies on various aspects of multilingual NTTS model training using monolingual corpora. For example, [2, 7, 8, 9] introduce several representative NTTS architectures for building voice cloning-based TTS systems. To enable cross-lingual voice conversion, authors in [10, 11] focus on achieving better disentanglement of several elements of speech, such as speaker characteristic and linguistic information. On addressing the data scarcity problems in some LVs, [12, 13] build multilingual TTS models for delivering TTS services in low-resource languages. Amongst these studies, one important topic that is largely missing is how the language and speaker composition of the training corpora affects the quality of multilingual speech from speakers present in the dataset. This paper’s contribution fills the gap with large scale experiments comparing various training data compositions in terms of gender, language family mixture of speakers and the number of speakers. Our aim is to understand how these factors affect the target speaker’s ability to speak the target language. We believe these results are valuable for constructing training corpora used for training Polyglot NTTS models.

The rest of the paper is organised as follows. Section 2 describes the architecture of the NTTS model we use for all our experiments. Section 3 details the experiment designs and the aim of each individual experiment. Section 4 presents experiment results accompanied by analyses. Finally, Section 5 concludes the paper.

2. NTTS architecture for Polyglot TTS

2.1. Model architecture

Our model is based on Tacotron 2 [14], i.e., it is sequence-to-sequence and attention-based [15], and consists of 3 main parts: an encoder, a decoder and an attention mechanism (Figure 1).

The encoder takes as input a sequence of phonemes and embeds them into 256-dimensional latent vectors. These phoneme

embeddings are then passed through a stack of 3 convolutional blocks, composed of 1D convolutions, ReLU activation, batch normalisation and dropout. Each 1D convolution in the encoder has 512 filters and a kernel size of 5, while the dropout rate for each block is set to 0.3. After that, hidden representations go through a single bi-LSTM unit, having an output dimension of 512. We employ a commonly used Tacotron 2 style decoder, similarly to [16, 17]. The attention mechanism feeds into the LSTM decoder. The output is passed through two linear projections. One projection leads to the stop token prediction. The other projection feeds into both the 512 CNN filters of the post-net and the 256-d feed-forward pre-net, which in turn feeds into the LSTM layer of the decoder. As for the attention mechanism, we rely on the default location-sensitive cell [18], with the same settings reported in [17], and we accumulate the attention weights at each decoder time-step to speed up the training procedure.

In addition, we employ specialized lookup tables to enable multi-speaker cross-lingual synthesis. In particular, we allocate 256-dimensional embeddings to condition the encoder on speaker, language and LV information, by concatenating them at each time-step. Finally, a custom reference encoder [19] is used to model speaking styles, where we consider “style” as anything that cannot be inferred solely from text. The reference encoder is composed of a VAE [20] that takes as input an 80-d vector representing a mel-spectrogram frame and outputs a 128-d latent vector. The VAE is then followed by 16 *householder* flow layers [21]. The reference encoder’s output is then concatenated to the encoder’s hidden states. Finally, we use a universal vocoder based on [22].

2.2. Data preparation and training strategy

Our Polyglot model takes as input a sequence of phonemes that are generated by passing input text through a TTS frontend. Each LV has its own phone set with punctuation, although all LVs share the same service tokens, such as word boundary and sentence start tokens. The training targets are mel-scaled spectrograms extracted from audio recordings.

We use a common set of training strategies, such as the optimiser, learning rate, gradient clipping, training loss, etc., for all models trained in this paper. As this is not the focus of our study, we refer readers to [16] for more details on model training settings.

3. Experiment and evaluation designs

3.1. Terminology

To facilitate the descriptions of experiment designs, we define the following terms and acronyms.

- **Target Speaker (TS):** The speaker whose voice is used for synthesising in multiple LVs.
- **Target Language variant (TL):** The LV(s) which the Polyglot model is intended to speak.
- **Speaker of Target Language variant (STL):** Speaker(s) whose data is included in Polyglot model training, and their native LV is a target LV.
- **Assisting Speaker (AS):** Speaker(s) whose data is included in model training, but their native LV is not one of the target LVs.

Note that we use the ISO 639 standard for denoting codes of language variants, i.e., they are noted in the format of {language code}-{LV CODE}, except for Arabic (arb).

3.2. Core dataset

Our experiments found that in order to ensure that audio samples are free from conspicuous quality artefacts such as distorted speaker identity, which could bias the evaluation of overall quality of Polyglot synthesis, a core subset of the available training corpus is needed. Note that finding an optimised strategy for maintaining consistent speaker identity during voice cloning is beyond the scope of this paper. We use the same core data composition, selected empirically, for all experiments conducted in this study. In particular, the core dataset we use includes 40 speakers whose native languages span 11 LVs across 3 languages (English, French, Spanish). In addition, we include 2K-4.5K utterances per speaker for AS and STL. For TSEs, we include around 15K-20K utterances by default, for better voice quality. Each utterance on average contains 55-75 phonemes and lasts about 5 seconds. All audio data are recorded in studio-grade settings. The designs of our experiments, described in Section 3.4, ensure that all evaluations are controlled, i.e., any set of models compared in an evaluation have a single changing variable amongst them. Therefore, in Section 3.4, instead of repeating the description of the part of training dataset shared by all models, we will implicitly intend that the core dataset described above is always included, and emphasise on the changing variable in data composition when describing the models.

3.3. Evaluation method

We use a mixture of preference test and Multiple Stimuli with Hidden Reference and Anchor (MUSHRA) test [23] as the means of subjective evaluation. As most evaluations we are conducting only involve comparing a pair of controlled models, we mainly rely on preference tests in which we ask test participants (native speakers of the LV being tested) to listen to samples from both models and indicate which audio sample they prefer overall. The preference test is essentially a binomial test. We use participants’ votes to establish whether to accept our null hypothesis that one model has higher probability of being perceived as better than the other. If the preference test does not show any statistically significant preference for a particular model, we then conduct more granular MUSHRA tests in which we ask test participants to rate audio samples on a scale between 0 and 100, in terms of naturalness and accent. For the naturalness MUSHRA test, we use the audio recordings of the TS speaking their native language as the upper anchor. For accent evaluation, we use recordings from a native speaker of the synthesised language as upper anchor, and synthesis from a model that has been trained with monolingual data of the original speaker only as lower anchor. For the lower anchor, we map each phone of the language we want to synthesise with to a similar phone that exists in the TS’s monolingual model, purposely producing accented speech. For both tests, we use the threshold of 5% to indicate whether results are statistically significant. We synthesise 100 testing samples for all LVs being tested, and each test sample is evaluated 50/30 times for the preference/MUSHRA tests, respectively.

3.4. Experiment design

In this section, we describe the design of the experiments for understanding how various factors in data composition affect the TS’s Polyglot synthesis in TLs. It should be noted that our experiments here are by no means exhaustive searches on all possible data compositions, which is infeasible in practice. Instead, we carefully design several controlled experiments to reveal the

Table 1: *Models trained & evaluated in Experiment 1.*

Exp.	Model	TL	Difference in data composition
1.1	M1	it-IT	4 Romance language speakers
	M2	de-DE	4 Germanic language speakers
1.2	M3	arb de-DE	Core dataset and data from STL
	M4	it-IT	12 ASes from multiple language families

Evaluation Results				
LV	M1 vote	No preference	M2 vote	p-value
de-DE	34.65%	26.94%	38.41%	0.03%
it-IT	50.41%	21.78%	27.82%	0.00%

LV	M3 vote	No pref	M4 vote	p-value
de-DE	20.10%	14.28%	65.62%	0.00%
it-IT	26.82%	13.36%	59.82%	0.00%
arb	37.45%	25.82%	36.73%	-

trend. The TSs, TLs and ASes for our experiments are selected to be representative, e.g. of a language family, while also considering data availability. All subjective evaluations conducted are controlled in terms of the number of training epochs.

3.4.1. Experiment 1: Language family affiliation of AS

In this experiment, the goal is to understand how the language family affiliations of ASes affect the quality of the TSes speaking the TLs. To this end, we design 2 sets of experiments as follows.

Experiment 1.1: For the 2 models (M1 and M2) trained in this experiment, we select it-IT (a Romance language) and de-DE (a Germanic language) as TLs. The only difference between the 2 models is that in the training dataset of M1 there are 4 ASes from Romance languages (1 fr-FR, 1 es-ES, 1 pt-BR, 1 es-US) in addition to the core dataset, whereas M2 has 4 additional ASes from 4 Germanic languages (1 en-US, 1 nb-NO, 1 sv-SE, 1 is-IS). For evaluation, we expect to observe whether the addition of ASes from different language families affects the TLs belonging to different language families differently.

Experiment 1.2: In Experiment 1.1 above, we only added ASes from the same language family in each model. We are also interested in knowing the effect of adding ASes belonging to multiple language families. Therefore, in Experiment 1.2, we add a basket of ASes from different language families, and evaluate how this would affect different TLs. Specifically, in model M3, the training data composition only includes the core dataset (described in Section 3.2) and data from STLs. By contrast, for training M4, we add 12 additional ASes in cy-GB, da-DA, is-IS, fr-FR, pl-PL, ro-RO, pt-BR, ru-RU, tr-TR, sv-SE, ko-KR, and nb-NO. There are 4 Germanic and 3 Romance language ASes amongst them, together with 5 ASes from other language families. For both M3 and M4, the TLs are selected to be arb, de-DE and it-IT. The addition of arb as TL in this experiment is meant to test whether adding ASes in distant language families from the TL has any bearing on the TL.

Table 1 summarises the experiment configurations. Note that all models evaluated in the same evaluation here have identical gender composition for speakers present in the datasets.

3.4.2. Experiment 2: Gender mixture of AS

In this experiment, we investigate the influence of gender mixture of ASes on male and female TS’s Polyglot synthesis, to

Table 2: *Models trained & evaluated in Experiment 2.*

Model	TL	TS	Difference in data composition
M5	arb it-IT	female male	8 female AS
M6			4 female AS + 4 male AS
M7			8 male AS

Evaluation Results				
Female TS				
LV	M5 vote	No pref	M6 vote	p-value
it-IT	42.86%	22.30%	34.84%	0.00%
arb	39.70%	25.90%	34.40%	0.00%
LV	M6 vote	No pref	M7 vote	p-value
it-IT	32.49%	33.90%	33.60%	-
arb	33.19%	33.54%	33.28%	-
Male TS				
LV	M5 vote	No pref	M6 vote	p-value
it-IT	43.96%	24.92%	31.12%	0.00%
arb	42.43%	21.97%	35.60%	0.00%
LV	M6 vote	No pref	M7 vote	p-value
it-IT	37.71%	31.38%	30.91%	0.00%
arb	32.87%	34.14%	32.99%	-

answer questions pertaining to the gender mixture of the speakers in the training corpora. For this purpose, all models we train for this experiment have a male and a female TS. The TLs for this experiment are arb and it-IT. Then, we consider 3 groups of gender mixtures of ASes, i.e., all female, all male, and gender-balanced. The models trained for this experiment (M5, M6, and M7) are summarised in Table 2. Note that all models trained for Experiment 2 and 3 have identical language composition.

3.4.3. Experiment 3: Gender mixture of STL

STLs differ from ASes in that in the end we expect the TSes to speak in the native LVs of STLs, but not of ASes. Therefore, we explore the gender mixture of STL separately in this experiment. Similar to Experiment 2 above, we train 3 models for Experiment 3, i.e., models with all female (M8), gender-balanced (M9), and all male (M10) STLs in the training dataset. Based on the availability of our voice data, we select de-DE, en-GB, and fr-CA as TLs. The comparison of models for this experiment can be found in Table 3.

3.4.4. Experiment 4: The number of STL

For a voice cloning-based multilingual TTS system, the quality of TS speaking the TL is directly affected by the amount of data from STL in the training dataset. Therefore, in this section we carry out experiments to examine the effect of the number of STL present in the training dataset. We aim to understand, for example, whether it is always better to have more STL in the training dataset. Intuitively, when there is a diverse training corpus in TL, we would expect the model to benefit from better disentanglement between language and speaker identity in the TL. In particular, for this experiment we use 4, 10, 16 STLs per TL for training 3 models M11, M12, and M13, respectively. The TLs here are en-GB, fr-CA, and de-DE.

Table 3: *Models trained & evaluated in Experiment 3.*

Model	TL	TS	Difference in data composition
M8	de-DE	female male	8 female STL/TL
M9	en-GB		4 female + 4 male STL/TL
M10	fr-CA		8 male STL/TL

Evaluation Results

Female TS				
LV	M8 vote	No pref	M9 vote	p-value
de-DE	58.25%	16.13%	25.63%	0.00%
en-GB	52.85%	10.38%	36.77%	0.00%
fr-CA	52.29%	18.14%	29.58%	0.00%
LV	M9 vote	No pref	M10 vote	p-value
de-DE	40.45%	18.59%	40.95%	-
en-GB	47.48%	11.12%	41.40%	0.00%
fr-CA	38.57%	27.87%	33.56%	0.00%
Male TS				
LV	M8 vote	No pref	M9 vote	p-value
de-DE	43.28%	23.89%	32.83%	0.00%
en-GB	49.74%	14.85%	35.41%	0.00%
fr-CA	47.34%	20.02%	32.65%	0.00%
LV	M9 vote	No pref	M10 vote	p-value
de-DE	52.66%	18.52%	28.82%	0.00%
en-GB	49.63%	10.46%	39.91%	0.00%
fr-CA	40.43%	23.75%	35.82%	0.01%

4. Evaluation result and discussion

4.1. Effects of language family affiliation of AS

The evaluation results for Experiment 1 are summarised in Table 1, which clearly show that adding ASes speaking languages in the same language family as the TL improves TS’s polyglot synthesis in that TL. Furthermore, evaluation results from comparing M3 and M4 show that adding together ASes from several language families benefits the TLs which share the language family affiliation of ASes (see Section 3.4.1 for detailed composition). The result also suggests that the extent of improvement for a TL is related to the amount of data from the ASes in the same language family as that of the TL, as de-DE receives higher percentage of votes than it-IT (the number of Germanic utterances is 80% more than that of Romance utterances). For the other TL, arb, M3 and M4 are equally preferred. This can be explained by the absence of additional languages that are close to Arabic from M3’s to M4’s data compositions. We conclude that adding AS whose language is close to the language of TL can help improve the TL while adding ASes far from the TL makes little difference for the TL.

4.2. Effects of gender mixture of AS and STL

Tables 2 and 3 summarise the evaluation results for Experiments 2 and 3, which test the influence of gender composition of AS and STL, respectively. The differences in AS/STL data composition can be also be found there. As the results from both experiments suggest the same trend in terms of preferred AS/STL gender mixture, we report them together in this section. A brief summary for this is that *all female* > *gender-balanced* ≥ *all male* for AS and STL data composition. This is true for both female and male TSes. In addition, the difference between

Table 4: *Models trained & evaluated in Experiment 4.*

Model	TL	Difference in data composition
M11	en-GB	4 STL/TL
M12	fr-CA	10 STL/TL
M13	de-DE	16 STL/TL

Evaluation Results

Item	Model	en-GB	fr-CA	de-DE	Significance
Nat.	M11	59.25	67.72	63.03	all difference significant
	M12	60.25	69.47	64.38	
	M13	57.68	65.33	61.68	
Acc.	M11	63.12	74.40	74.38	only diff. M11&M13 insignificant
	M12	64.37	75.95	76.28	
	M13	63.34	74.03	74.5	

gender-balanced and all male in AS/STL is more pronounced for male TS, as there are more cases where there is statistically significant preference for gender-balanced AS/STL over all male AS/STL. Informal subjective evaluations found that for male TS, models trained with more female data tend to reduce the creakiness of TS’s voice speaking the TLs. The benefits of having more female AS/STL is less noticeable for female TS’s, but we nonetheless observed improved articulation clarity in some cases. This result is surprising as we did not observe any speaker identity issues, i.e., the TSes still keep their speaker identities even when all training data in the TLs is from STLs of the opposite gender. This is a strong data point which indicates that by including the core dataset described in Section 3.2, the model achieved good speaker-language disentanglement.

4.3. Effects of the amount of STL data

In Table 4, we present evaluation results from MUSHRA tests for comparing M11, M12, and M13. As mentioned in Section 3.3, here we use MUSHRA tests for evaluating naturalness and accent separately, as preference tests failed to produce statistically significant distinctions between the models. From the MUSHRA scores in the table, we can see that having 10 STL produces the best scores for both naturalness and accent, for all 3 LVs tested. For naturalness, there is $10\text{ STL} > 4\text{ STL} > 16\text{ STL}$; whereas we have $10\text{ STL} > 4\text{ STL} = 16\text{ STL}$ for accent. This result refutes our original hypothesis that having more STLs always improves the model’s TL performance. Instead, we observe that there is a sweet spot, and more data is not always better. This might be explained by the fact that the training dataset becomes more unbalanced towards the TLs, as we keep adding STL data, which in turn hinders the model’s ability to generalise and disentangle different LVs.

5. Conclusions

In this paper, we studied how the data composition of training corpora affects multilingual NTTS synthesis in terms of the language family affiliation, gender balance and the number of speakers present. Our most surprising result is that including female speakers is better than having a balanced mixture of genders. Also, supporting speakers whose native LVs are closer to the target language are beneficial. Moreover, we conclude that adding more speakers in the target LV improves the model only up to a point, after which adding more incurs regressions. While these results are useful guidelines for building training corpora for multilingual NTTS systems, we are not yet able to fully explain some of them, which will be investigated as future work.

6. References

- [1] H. Javkin, K. Hata, L. Mendes *et al.*, “A multilingual text-to-speech system,” in *International Conference on Acoustics, Speech, and Signal Processing*, IEEE, 1989, pp. 242–245.
- [2] Y. Zhang, R. J. Weiss, H. Zen *et al.*, “Learning to speak fluently in a foreign language: Multilingual speech synthesis and cross-language voice cloning,” *Proc. Interspeech 2019*, pp. 2080–2084, 2019.
- [3] Q. Yu, P. Liu, Z. Wu *et al.*, “Learning cross-lingual information with multilingual blstm for speech synthesis of low-resource languages,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 5545–5549.
- [4] P. Do, M. Coler, J. Dijkstra *et al.*, “A systematic review and analysis of multilingual data strategies in text-to-speech for low-resource languages,” in *Interspeech 2021*. ISCA, 2021, pp. 16–20.
- [5] J. Fong, J. Wu, P. Agrawal *et al.*, “Improving polyglot speech synthesis through multi-task and adversarial learning,” in *Proc. 11th ISCA Speech Synthesis Workshop (SSW 11)*, pp. 172–176.
- [6] S. Zhao, T. H. Nguyen, H. Wang *et al.*, “Towards natural bilingual and code-switched speech synthesis based on mix of monolingual recordings and cross-lingual voice conversion,” *Proc. Interspeech 2020*, pp. 2927–2931, 2020.
- [7] E. Nachmani and L. Wolf, “Unsupervised polyglot text-to-speech,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 7055–7059.
- [8] I. Himawan, S. Aryal, I. Ouyang *et al.*, “Speaker adaptation of a multilingual acoustic model for cross-language synthesis,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7629–7633.
- [9] J. Yang and L. He, “Cross-lingual text-to-speech using multi-task learning and speaker classifier joint training,” *arXiv preprint arXiv:2201.08124*, 2022.
- [10] J. Williams, J. Fong, E. Cooper *et al.*, “Exploring disentanglement with multilingual and monolingual vq-vae,” *arXiv preprint arXiv:2105.01573*, 2021.
- [11] Z. Wang, W. Ge, X. Wang *et al.*, “Accent and speaker disentanglement in many-to-many voice conversion,” in *2021 12th International Symposium on Chinese Spoken Language Processing (ISCSLP)*. IEEE, 2021, pp. 1–5.
- [12] G. Maniati, N. Ellinas, K. Markopoulos *et al.*, “Cross-lingual low resource speaker adaptation using phonological features,” *arXiv preprint arXiv:2111.09075*, 2021.
- [13] I. Demirsahin, M. Jansche, and A. Gutkin, “A unified phonological representation of south asian languages for multilingual text-to-speech,” in *Proc. The 6th Intl. Workshop on Spoken Language Technologies for Under-Resourced Languages*, pp. 80–84.
- [14] J. Shen, R. Pang, R. J. Weiss *et al.*, “Natural tts synthesis by conditioning wavenet on mel spectrogram predictions,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 4779–4783.
- [15] I. Sutskever, O. Vinyals, and Q. V. Le, “Sequence to sequence learning with neural networks,” in *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, ser. NIPS’14. Cambridge, MA, USA: MIT Press, 2014, p. 3104–3112.
- [16] T. Merritt, B. Putrycz, A. Nadolski *et al.*, “Comprehensive evaluation of statistical speech waveform synthesis,” in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 325–331.
- [17] J. Latorre, J. Lachowicz, J. Lorenzo-Trueba *et al.*, “Effect of data reduction on sequence-to-sequence neural TTS,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 7075–7079.
- [18] J. K. Chorowski, D. Bahdanau, D. Serdyuk *et al.*, “Attention-based models for speech recognition,” in *Advances in Neural Information Processing Systems*, C. Cortes, N. Lawrence, D. Lee *et al.*, Eds., vol. 28. Curran Associates, Inc., 2015.
- [19] R. J. Skerry-Ryan, E. Battenberg, Y. Xiao *et al.*, “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*, J. G. Dy and A. Krause, Eds.
- [20] D. P. Kingma and M. 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.
- [21] J. M. Tomczak and M. Welling, “Improving variational auto-encoders using householder flow,” *CoRR*, vol. abs/1611.09630, 2016. [Online]. Available: <http://arxiv.org/abs/1611.09630>
- [22] Y. Jiao, A. Gabryś, G. Tinchev *et al.*, “Universal neural vocoding with parallel wavenet,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6044–6048.
- [23] ITUR Recommendation, “Bs. 1534-1. method for the subjective assessment of intermediate sound quality (mushra),” *International Telecommunications Union, Geneva*, 2001.