



The Effects of Input Type and Pronunciation Dictionary Usage in Transfer Learning for Low-Resource Text-to-Speech

Phat Do¹, Matt Coler¹, Jelske Dijkstra², Esther Klabbbers³

¹Campus Fryslân, University of Groningen, the Netherlands

²Fryske Akademy/Mercator Research Centre, the Netherlands

³ReadSpeaker, the Netherlands

{t.p.do, m.coler}@rug.nl, jdijkstra@fryske-akademy.nl, esther.judd@readspeaker.com

Abstract

We compare phone labels and articulatory features as input for cross-lingual transfer learning in text-to-speech (TTS) for low-resource languages (LRLs). Experiments with FastSpeech 2 and the LRL West Frisian show that using articulatory features outperformed using phone labels in both intelligibility and naturalness. For LRLs without pronunciation dictionaries, we propose two novel approaches: a) using a massively multilingual model to convert grapheme-to-phone (G2P) in both training and synthesizing, and b) using a universal phone recognizer to create a makeshift dictionary. Results show that the G2P approach performs largely on par with using a ground-truth dictionary and the phone recognition approach, while performing generally worse, remains a viable option for LRLs less suitable for the G2P approach. Within each approach, using articulatory features as input outperforms using phone labels.

Index Terms: neural text-to-speech synthesis, low-resource languages, articulatory features, pronunciation dictionary

1. Introduction

1.1. Multilingual data for TTS in low-resource languages

Neural TTS produces speech that is both more intelligible and more natural than its preceding paradigms [1]. However, it requires large amounts of training data: LJSpeech [2] is one of the most commonly used data sets in neural TTS research and it has nearly 24 hours of professionally recorded single-speaker American English speech. This makes it hard to directly use neural TTS for low-resource languages (LRLs), which make up most of the languages in the world. One approach to deal with this is to make use of data from high-resource languages. This is done by taking advantage of the latent space information shared by languages, even if they are apparently very different. A meta-analysis [3] found this to be effective, and even more so with neural TTS than with earlier TTS paradigms.

In this multilingual data approach, one particular method is cross-lingual transfer learning: pre-training the acoustic model with ample data from a high-resource language (the “*source language*”) and fine-tuning it with limited data in the LRL (the “*target language*”). However, this has an inherent challenge of input mismatch. Two separate languages, however close to each other, almost always have different phone sets. While using grapheme input may avoid the issue of different phone sets, it may also lead to mispronunciation issues. Another challenge is with unseen phones: phones of the target language that are not in the source language. For these, the model has to initialize their associated weights from scratch and learn from the limited data. This makes transfer learning less effective. To avoid this, [4] and our previous work [5] explored solving both challenges

simultaneously by applying different methods to map the target languages’ phones (“*target phones*”) to their closest counterparts in the source languages (“*source phones*”). By using the pre-trained weights of the source phones, this approach was able to benefit the transfer learning process.

Such phone mapping gives promising results, but it also has issues. First, it may introduce “*accented*” speech. This is when the target phones still (partially) sound like the source phones, especially with very little fine-tuning data. Second, the use of phone labels (e.g., IPA or X-SAMPA symbols) as model input means that the input embeddings are treated in an all-or-nothing manner. In other words, two phones that are represented by two different labels will be considered completely different, even if their pronunciations are close to each other. Consequently, there are likely source phones that are not mapped to any phones and thus simply unused. This leads to inefficient use of data, which is not helpful for LRLs. Both issues may be solved by replacing phone labels (and thus avoiding phone mapping) with universal articulatory features. These are features associated with how the phones are pronounced and can be systematically looked up from the phone labels. The universality of these features also pre-emptively avoids the input mismatch issue, making cross-lingual transfer learning more extensively applicable.

Thus, articulatory features were used as input in transfer learning by [6]. However, they did not find significant improvements in speech quality compared to using phone labels. Since they used Tacotron 2 [7], an autoregressive (AR) TTS architecture, their result could be partly due to the architecture’s higher data requirement (compared to non-AR TTS [8]) and its association with unstable attention training [9]. Therefore, this study investigates whether using a non-AR TTS architecture can lead to better results in transfer learning with articulatory features.

1.2. TTS for LRLs without pronunciation dictionaries

Another issue for LRLs is they often lack a pronunciation dictionary. This limits TTS for such LRLs to two options. The first is using grapheme input, which may lead to mispronunciation issues, which are then compounded in transfer learning. The second is building a pronunciation dictionary or a grapheme-to-phone (G2P) system from scratch. This requires deep linguistic expertise and heavy time investment, both not viable for LRLs.

One approach is to circumvent the need for an explicit dictionary and use G2P conversion during both training and synthesizing. To this end, G2P performance in LRLs can be aided by using multilingual data (similar to TTS acoustic models). One approach that has been used is to train a massively multilingual (269-language) G2P model using found data from Wiktionary, as demonstrated by [10]. To enable parameter sharing across all training languages, they combined data from all lan-

guages into a single training set and used special language tokens in the input sequences to distinguish between languages. For predicting each test utterance in an unseen test language, they 1) searched for the test language’s 10 “nearest” training languages using a phylogenetic tree, 2) generated one prediction from each of these 10 languages (conditioned by the corresponding language token) and 3) ensembled the 10 predictions to form the final prediction. They reached an average phone error rate (PER) of 35.7% across 605 test languages. Though far from ideal, we argue that this is useful enough for TTS in LRLs, even more so if articulatory features are used for input, since these can facilitate efficient (transfer) learning even if the phone labels are not identical, as mentioned in Section 1.1.

However, this G2P approach may not work equally well for all LRLs. For any LRL, if its phone inventory and G2P rules are too different from those of its 10 nearest training languages, the approach’s performance will likely be lower. It will also degrade if these nearest languages themselves have limited training data and thus the G2P predictions conditioned on them are already not adequately accurate. Therefore, we propose another approach that is theoretically more language-independent: using a universal phone recognizer to predict phone sequences in the audio segments and creating a makeshift pronunciation dictionary using the accompanying texts. This dictionary is then used in the TTS pipeline. This approach is inspired by a recent work in phone recognition by [11]. First, they cascaded a language-independent narrow phone layer and a language-specific broad phone layer in the recognition model. Second, they used phone sets from PHOIBLE [12] (a database of phone inventories of 2,186 languages) to filter the phone output. They reached a PER of 64.2% for the test language Tusom. While even further from satisfactory than the G2P approach, this approach is more likely to work language-independently. Thus, we posit that this approach can be considered for TTS in LRLs, especially with the use of articulatory features as input. As mentioned in Section 1.1, this input type helps with transfer learning even if the phone labels are not identical and thus is likely more tolerant of inaccuracies in phone prediction.

1.3. Contributions

Accordingly, we aim to make the following contributions:

- 1) We use a non-autoregressive architecture (FastSpeech 2 [13]) in cross-lingual transfer learning for TTS in the low-resource language (LRL) West Frisian to compare the effectiveness of phone labels and articulatory features as model input.
- 2) We explore two options in the absence of a pronunciation dictionary for the target LRL: *a)* predicting phone sequences in both training and synthesizing with a massively multilingual grapheme-to-phone model, and *b)* predicting phone sequences in audio with a universal phone recognizer and building a makeshift pronunciation dictionary.

2. Data sets & proposed pipelines

2.1. Languages & data used

For pre-training, we used the LJSpeech data set to facilitate comparison in future research thanks to its popularity. For the target LRL, we chose West Frisian (“Frysk”, hereafter Frisian), the second official language of the Netherlands with roughly 350,000 native speakers [14]. We created a small single-speaker data set from a Frisian audiobook to be used as training data. We split the recordings into utterances by silence and manu-

ally checked to obtain their corresponding texts, normalized and expanded the texts where relevant, and trimmed all preceding and trailing silence. We randomly selected 150 utterances (~15 minutes) for the training and validation set, and another 100 utterances for the test set. Both sets have similar distributions of duration d : $\bar{d} = 6.0$, $s_d = 2.2$, $1 \leq d \leq 10$ (seconds). For the baseline models, we used a pronunciation dictionary. This dictionary has roughly 73,000 entries and was derived from the Frisian Audio Mining Enterprise (FAME) project [15]. For the other models, to simulate the lack of pronunciation dictionaries in LRLs, we used this dictionary only to analyze the results.

2.2. Multilingual grapheme-to-phone (G2P) model

Following Section 1.2, we used the pre-trained model from [10] to do G2P conversion for the Frisian data in both training and synthesizing. Using the dictionary as baseline, this model had a mean phone error rate of $33\% \pm 6\%$ (SE). It should be noted that the model’s training data does include Frisian among its 269 training languages, so it is not exactly an unseen language. It is thus reasonable to expect the model to perform better on Frisian than on truly unseen languages. However, the Frisian training data was still rather limited at 991 entries. This places it among more than 95% of the 874 languages in [10] with fewer than 5,000 entries. We thus assumed that the results from experimenting with Frisian are still relevant for many LRLs, despite technically not being representative for all languages.

2.3. Universal phone recognition model

We used the recognizer checkpoint provided by the authors of [11] to perform phone recognition on the Frisian training and validation sets, limiting the model’s output phone pool to the 40-phone Frisian phone set provided by PHOIBLE v2.0. Initial tests showed that the result from this was insufficiently accurate and thus not suitable for TTS training. As a result, we used ground truth phone sequences of 30 utterances (~3 minutes of speech) to fine-tune the phone recognizer model. Admittedly, such ground truth information is likely not available in practical use with many LRLs, so this partially reduces the approach’s applicability for LRLs. We discuss this further in Section 4.3.

The fine-tuned model was used to obtain 150 corresponding pairs of texts and predicted phone sequences. Although these were sufficient to train the TTS model, we also needed to generate phone sequences for the test utterances. This is because in a practical scenario, there is no ground truth audio and there is also a need for G2P capability for out-of-vocabulary (OOV) words. There was no straightforward method to do this since we lacked word boundaries information in the audio. Therefore, we first tried using this 150-pair mini data set with OpenNMT v2.0 [16] to train a G2P model, but could not reach convergence due to the small data size. Consequently, we built a mini makeshift pronunciation dictionary using the predicted phone sequences and their corresponding texts. We manually decided the word boundaries. For words that had more than one predicted phone sequence, we simply chose the most common prediction. We then used this dictionary to train a G2P model so that it could cover all words (including OOVs) in the test set.

3. Experiment details

3.1. TTS model architecture & training

For the acoustic model, we used the FastSpeech 2 implementation (~35M parameters) of [17]. This used phone duration

information extracted by the Montreal Forced Aligner (MFA) [18]. MFA (v2.0.5) was also used for G2P when needed. For the models with articulatory features as input, we used the implementation in the IMS-Toucan Toolkit v2.3 from [19]. Phone labels were converted into one-hot encoded vectors of corresponding values from a set of 62 binary articulatory features. This followed the convention in [20] and complemented their set of 49 features with features such as phone length and more fine-grained values for place and manner of consonant articulation. For the vocoder, we used pre-trained HiFi-GAN V1 (~14M parameters) [21] for all models.

We first pre-trained two models in English, one with phone labels and one with articulatory features as input. Each model was trained for 200,000 parameter updates with a batch size of 32 (~500 epochs) using the Adam optimizer ([22], $\beta_1 = 0.9$, $\beta_2 = 0.98$, $\epsilon = 10^{-9}$). We then fine-tuned these two models using the ground-truth Frisian pronunciation dictionary and call the resulting models *ph-gt* and *ft-gt* following their input types. Each fine-tuning was run for 5,000 parameter updates with a batch size of 12 (~400 epochs). All training was done on a single NVIDIA V100 32GB GPU, taking roughly 45 hours for pre-training and roughly 50 minutes for each fine-tuning.

From the pre-trained models, we also fine-tuned models that did not have access to the dictionary following the pipelines in Section 2.2 and 2.3. We call models from the G2P approach *-g2p* and those from the phone recognition approach *-rec*. To facilitate comparisons between input types, we did this with both phone labels (*ph-*) and articulatory features (*ft-*) input, resulting in four models: *ph-g2p*, *ph-rec*, *ft-g2p*, and *ft-rec*.

3.2. Evaluation

We used the fine-tuned models to synthesize the same set of 100 unseen utterances and evaluated them in intelligibility and naturalness. For intelligibility, we used an automatic speech recognition (ASR) model¹ fine-tuned from the “large” multilingual checkpoint of the self-supervised learning model *wav2vec 2.0* [23] on the 50-hour Frisian data set from Common Voice (v8) [24]. It reports a word error rate of 16.25% on Common Voice’s test set. We used the ASR model directly (without a language model) on the synthesized utterances and calculated character error rates (CER) for evaluation.

For naturalness, we used an automatic MOS (Mean Opinion Score) prediction model that is based on and fine-tuned from *wav2vec 2.0 Base* [23]. This follows the approach in [25]. Following the result of our work on efficient fine-tuning strategies for MOS prediction in LRLs [26], we first pre-trained the model on the BVCC data set [25] and then further trained it on the neural-TTS-only SOMOS data set [27]. Both of these data sets are in English. We then fine-tuned the model further on the MOS data² from our previous work [5]. This data contains 2,024 MOS ratings for 220 synthetic utterances (from 11 systems: 10 TTS and 1 resynthesis) rated by 46 participants. The original 100-scale MUSHRA scores were linearly converted to the 5-point MOS scale. We split this data into 80% training, 10% validation, and 10% test sets. Table 1 shows the prediction accuracy of the fine-tuned model on the test set, measured in MSE (Mean Squared Error) and LCC (Linear Correlation Coefficient, i.e., Pearson’s r). For LCC, corresponding p -values were checked for statistical significance. We also include the best-performing measures from the out-of-domain track (OOD)

¹huggingface.co/wietsedv/wav2vec2-large-xlsr-53-frisian

²phat-do.github.io/sigul22

of the recent VoiceMOS Challenge [28]. This OOD track is similar to our settings in terms of cross-lingual MOS prediction with limited ground-truth data. Table 1 shows that the fine-tuned model has an adequate level of prediction accuracy, with an MSE of 0.190 and an LCC of 0.821 for utterance-level predictions. Besides, the test set in this study (described in Section 2.1) has many similarities with the fine-tuning data from [5]: both are trained with data from audiobooks in Frisian, both share the FastSpeech 2 architecture, etc. In short, we posit that this fine-tuned model was suitable to be used for this study’s evaluation and analysis.

Table 1: *Fine-tuned MOS prediction accuracy*

Level	Metric	Current	VoiceMOS’s best
Utterance	MSE	0.190	0.162
	LCC	0.821	0.921
System	MSE	0.074	0.030
	LCC	0.946	0.988

4. Results & discussion

We used 100 test utterances for analysis, but for practical reasons, we randomly picked out 20 utterances (140 audio samples) and shared them online for reference³. Figures 1 and 2 show the boxplots of CER and MOS for all systems described in Section 3. For reference, we also included results from resynthesized audio (generated from ground-truth spectrograms).




Figure 1: *CER on the test set (lower is better)*




Figure 2: *Predicted MOS on the test set (higher is better)*

We used a linear mixed effects model and treated input type (*ph* and *ft*) and “dictionary type” (*gt*, *g2p*, and *rec*), plus their interaction, as fixed effects. For random effects, we used random intercepts for utterances and by-utterance random slopes for both input type and dictionary type. Residual plots were

³phat-do.github.io/nodict-IS23

used to check for assumptions of homoscedasticity and normality. Statistical significance was checked using p -values from the likelihood ratio tests between models with and without each effect in question. The results from this show that input type, dictionary type, and their interaction all had significant effects on both CER and MOS. Therefore, we conducted a post-hoc analysis using Tukey’s multiple comparison of means [29] to check if the mean CER and MOS of all scenarios (combinations of input and dictionary type) were significantly different from each other. To avoid cluttering, only relevant scenarios are included in Table 2. Significant differences are in bold.

Table 2: Results from Tukey’s multiple comparison of means

Group 1		Group 2		$\mu_{G2} - \mu_{G1}$ & p -value			
Input	Dict	Input	Dict	CER		MOS (predicted)	
<i>ph</i>	<i>gt</i>	<i>ft</i>	<i>gt</i>	-0.045	.003	0.982	< .001
<i>ph</i>	<i>gt</i>	<i>g2p</i>		-0.026	.267	0.101	.386
	<i>gt</i>	<i>ph</i>	<i>rec</i>	0.065	< .001	0.114	.341
	<i>g2p</i>		<i>rec</i>	0.092	< .001	0.004	.999
<i>ft</i>	<i>gt</i>	<i>g2p</i>		0.055	< .001	-0.021	.999
	<i>gt</i>	<i>ft</i>	<i>rec</i>	0.125	< .001	-0.543	< .001
	<i>g2p</i>		<i>rec</i>	0.070	< .001	-0.522	< .001
<i>ph</i>	<i>g2p</i>	<i>ft</i>	<i>g2p</i>	0.035	.042	0.851	< .001
	<i>rec</i>		<i>rec</i>	0.014	.857	0.324	< .001

4.1. Input type comparison with ground-truth dictionary

With *gt*, *ft* decreased CER by 0.045 (± 0.01) and increased MOS by 0.982 (± 0.05) compared to *ph*. In other words, with all things equal (the same TTS architecture, training data, training procedure, evaluation method, etc.), using articulatory features as input outperformed using phone labels for cross-lingual transfer learning. This likely comes from the better learning efficiency hypothesized in Section 1.1. As a result, TTS for LRLs is expected to benefit from this. This is at least with a non-autoregressive (non-AR) architecture as used in this study, as opposed to the AR architecture in [6].

4.2. Multilingual G2P & phone recognition

Phone labels input: Compared to *gt*, *g2p* did not significantly change either CER or MOS. Meanwhile, *rec* increased CER by 0.065 (± 0.01) compared to *gt* and 0.092 (± 0.01) compared to *g2p*. This means that with phone labels, using multilingual G2P was a viable approach, giving output speech quality comparable to that from the ground-truth dictionary. The phone recognition approach was less effective, producing a similar level of naturalness but worse intelligibility.

Articulatory features input: Compared to *gt*, *g2p* led to an increase of 0.055 (± 0.01) in CER, but no significant change in MOS. However, *rec* increased CER by 0.125 (± 0.01) and decreased MOS by 0.543 (± 0.04) compared to *gt*. This means that with articulatory features, multilingual G2P was only comparable to the ground-truth dictionary in naturalness, and phone recognition led to worse quality in both criteria.

Phone labels vs. articulatory features: For *g2p*, *ft* led to both an increase of 0.035 (± 0.006) in CER and an increase of 0.851 (± 0.03) in MOS. For *rec*, *ft* had no effect on CER, but increased MOS by 0.324 (± 0.03). In other words, in the absence of an actual pronunciation dictionary, articulatory features outperformed phone labels in naturalness while performing either worse or similarly in intelligibility. This is partially in line with

our hypothesis in Section 1.2.

4.3. Assumptions & limitations

Assumptions: We assumed that the use of CER calculated from ASR predictions was relevant in evaluating the intelligibility of synthetic speech, i.e., it is comparable to human evaluation. This is especially notable since we used an ASR model without a language model. Meanwhile, humans cannot really turn off their language models [30]. For MOS evaluation, besides the usual assumption that MOS can be reliably used to judge “naturalness”, we assumed that our model was fine-tuned enough to give scores representative of those made by human listeners.

Limitations: As mentioned in Section 2.2, Frisian was among the G2P model’s training data (albeit with limited data). Therefore, we could not reliably infer that the approach would work equally well for truly unseen LRLs. For the phone recognition model, since it relies on phone sets from PHOIBLE, its performance would degrade for an LRL not among the 2,186 languages covered by this database. Even for a covered language, the phone recognition model required some fine-tuning before being directly integrated into the TTS pipeline due to its base performance. This fine-tuning may require some expertise in phonetics and/or the LRL, which may limit its applicability.

5. Conclusions & future work

In this study, we explored cross-lingual transfer learning from English (LJSpeech) to West Frisian, using 15 minutes of audio-book data. The output speech was evaluated with CER obtained from an ASR model and MOS from an MOS prediction model, both based on *wav2vec 2.0*. The results confirmed that using articulatory features as input for the TTS acoustic model significantly outperformed using phone labels in both intelligibility and naturalness. Accordingly, we believe that this approach is beneficial for similar work in low-resource languages (LRLs).

For LRLs without available pronunciation dictionaries, we propose two approaches. The first is doing grapheme-to-phone (G2P) conversion directly in both training and synthesizing with a massively multilingual G2P model. The second is creating a makeshift dictionary using a universal phone recognition model. The results of our experiments indicated that the G2P approach was generally comparable to a ground-truth dictionary, while the phone recognition approach only showed comparable results in one case. Thus, G2P is a viable option in almost all cases while phone recognition is more limited but remains an option for LRLs that cannot benefit much from G2P. The two approaches’ effects also differed between the input types: articulatory features generally led to better speech quality than phone labels. These findings are expected to enable and provide strategies toward developing TTS for the thousands of LRLs that have no available pronunciation dictionaries.

Future research is planned to explore bypassing the intermediate step of phonetic transcriptions (to skip a potential source of inaccuracy) and converting audio directly to articulatory features to be used in creating a makeshift dictionary.

6. Acknowledgements

We thank the Center for Information Technology of the University of Groningen for providing access to the Hábrók high performance computing cluster. We are thankful for the permission to use the recordings by Geartsje de Vries of books written by Koos Tiemersma and published by the publisher Audiofrysk.

7. References

- [1] X. Tan, T. Qin, F. Soong, and T.-Y. Liu, “A Survey on Neural Speech Synthesis,” *arXiv:2106.15561 [cs, eess]*, Jun. 2021.
- [2] K. Ito and L. Johnson, “The LJ Speech Dataset,” <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [3] P. Do, M. Coler, J. Dijkstra, and E. Klabbbers, “A Systematic Review and Analysis of Multilingual Data Strategies in Text-to-Speech for Low-Resource Languages,” in *Interspeech 2021*. ISCA, Aug. 2021, pp. 16–20.
- [4] Y.-J. Chen, T. Tu, C.-C. Yeh, and H.-Y. Lee, “End-to-End Text-to-Speech for Low-Resource Languages by Cross-Lingual Transfer Learning,” in *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, vol. 2019-September, 2019, pp. 2075–2079.
- [5] P. Do, M. Coler, J. Dijkstra, and E. Klabbbers, “Text-to-Speech for Under-Resourced Languages: Phoneme Mapping and Source Language Selection in Transfer Learning,” in *Proceedings of the 1st Annual Meeting of the ELRA/ISCA Special Interest Group on Under-Resourced Languages*. Marseille, France: European Language Resources Association, Jun. 2022, pp. 16–22.
- [6] D. Wells and K. Richmond, “Cross-lingual Transfer of Phonological Features for Low-resource Speech Synthesis,” in *11th ISCA Speech Synthesis Workshop (SSW 11)*. ISCA, Aug. 2021, pp. 160–165.
- [7] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan, R. A. Saurous, Y. Agiomvrgiannakis, and Y. Wu, “Natural TTS Synthesis by Conditioning Wavenet on MEL Spectrogram Predictions,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2018, pp. 4779–4783.
- [8] A. Pine, D. Wells, N. Brinklow, P. Littell, and K. Richmond, “Requirements and Motivations of Low-Resource Speech Synthesis for Language Revitalization,” in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Dublin, Ireland: Association for Computational Linguistics, May 2022, pp. 7346–7359.
- [9] C. Valentini-Botinhao and S. King, “Detection and Analysis of Attention Errors in Sequence-to-Sequence Text-to-Speech,” in *Interspeech 2021*. ISCA, Aug. 2021, pp. 2746–2750.
- [10] X. Li, F. Metzger, D. Mortensen, S. Watanabe, and A. Black, “Zero-shot Learning for Grapheme to Phoneme Conversion with Language Ensemble,” in *Findings of the Association for Computational Linguistics: ACL 2022*. Dublin, Ireland: Association for Computational Linguistics, May 2022, pp. 2106–2115.
- [11] X. Li, S. Dalmia, J. Li, M. Lee, P. Littell, J. Yao, A. Anastasopoulos, D. R. Mortensen, G. Neubig, and A. W. Black, “Universal Phone Recognition with a Multilingual Allophone System,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 8249–8253.
- [12] S. Moran and D. McCloy, Eds., *PHOIBLE 2.0*. Jena: Max Planck Institute for the Science of Human History, 2019. [Online]. Available: <https://phoible.org/>
- [13] 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*, Feb. 2023.
- [14] D. Gorter, C. van der Meer, and A. Riemersma, “Frisian in the Netherlands,” *Multilingual Europe: facts and policies*, pp. 185–206, 2008.
- [15] E. Yilmaz, M. Andringa, S. Kingma, J. Dijkstra, F. van der Kuip, H. Van de Velde, F. Kampstra, J. Algra, H. van den Heuvel, and D. van Leeuwen, “A Longitudinal Bilingual Frisian-Dutch Radio Broadcast Database Designed for Code-Switching Research,” in *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*. Portorož, Slovenia: European Language Resources Association (ELRA), May 2016, pp. 4666–4669.
- [16] G. Klein, Y. Kim, Y. Deng, J. Senellart, and A. M. Rush, “Open-NMT: Open-Source Toolkit for Neural Machine Translation,” in *Proceedings of ACL 2017, System Demonstrations*, 2017, pp. 67–72.
- [17] C.-M. Chien, J.-H. Lin, C.-y. Huang, P.-c. Hsu, and H.-y. Lee, “Investigating on Incorporating Pretrained and Learnable Speaker Representations for Multi-Speaker Multi-Style Text-to-Speech,” in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Jun. 2021, pp. 8588–8592.
- [18] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, “Montreal Forced Aligner: Trainable Text-Speech Alignment Using Kaldi,” in *Interspeech 2017*. ISCA, Aug. 2017, pp. 498–502.
- [19] F. Lux and T. Vu, “Language-Agnostic Meta-Learning for Low-Resource Text-to-Speech with Articulatory Features,” in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Dublin, Ireland: Association for Computational Linguistics, May 2022, pp. 6858–6868.
- [20] M. Staib, T. H. Teh, A. Torresquintero, D. S. R. Mohan, L. Foglianti, R. Lenain, and J. Gao, “Phonological Features for 0-shot Multilingual Speech Synthesis,” *Interspeech 2020*, pp. 2942–2946, Oct. 2020.
- [21] J. Kong, J. Kim, and J. Bae, “HiFi-GAN: Generative Adversarial Networks for Efficient and High Fidelity Speech Synthesis,” in *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc., 2020, pp. 17 022–17 033.
- [22] D. P. Kingma and J. 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*, Y. Bengio and Y. LeCun, Eds., 2015.
- [23] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “Wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations,” in *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc., 2020, pp. 12 449–12 460.
- [24] R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. Tyers, and G. Weber, “Common Voice: A Massively-Multilingual Speech Corpus,” in *Proceedings of the 12th Language Resources and Evaluation Conference*, 2020, pp. 4218–4222.
- [25] E. Cooper, W.-C. Huang, T. Toda, and J. Yamagishi, “Generalization Ability of MOS Prediction Networks,” in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2022, pp. 8442–8446.
- [26] P. Do, M. Coler, J. Dijkstra, and E. Klabbbers, “Resource-Efficient Fine-Tuning Strategies for Automatic MOS Prediction in Text-to-Speech for Low-Resource Languages,” *arXiv:2305.19396 [cs, eess]*, May 2023.
- [27] G. Maniati, A. Vioni, N. Ellinas, K. Nikitaras, K. Klapsas, J. S. Sung, G. Jho, A. Chalamandaris, and P. Tsiakoulis, “SOMOS: The Samsung Open MOS Dataset for the Evaluation of Neural Text-to-Speech Synthesis,” in *Interspeech 2022*. ISCA, Sep. 2022, pp. 2388–2392.
- [28] W. C. Huang, E. Cooper, Y. Tsao, H.-M. Wang, T. Toda, and J. Yamagishi, “The VoiceMOS Challenge 2022,” in *Interspeech 2022*. ISCA, Sep. 2022, pp. 4536–4540.
- [29] W. Haynes, “Tukey’s Test,” in *Encyclopedia of Systems Biology*, W. Dubitzky, O. Wolkenhauer, K.-H. Cho, and H. Yokota, Eds. New York, NY: Springer, 2013, pp. 2303–2304.
- [30] K. Lakhotia, E. Kharitonov, W.-N. Hsu, Y. Adi, A. Polyak, B. Bolte, T.-A. Nguyen, J. Copet, A. Baevski, A. Mohamed, and E. Dupoux, “On Generative Spoken Language Modeling from Raw Audio,” *Transactions of the Association for Computational Linguistics*, vol. 9, pp. 1336–1354, 2021.