

# Low-data? No problem: low-resource, language-agnostic conversational text-to-speech via F0-conditioned data augmentation

Giulia Comini, Goeric Huybrechts, Manuel Sam Ribeiro, Adam Gabryś, Jaime Lorenzo-Trueba

Amazon Alexa, TTS Research

{gcomini, huybrech, manuerib, gabrysa, truebaj}@amazon.com

## Abstract

The availability of data in expressive styles across languages is limited, and recording sessions are costly and time consuming. To overcome these issues, we demonstrate how to build low-resource, neural text-to-speech (TTS) voices with only 1 hour of conversational speech, when no other conversational data are available in the same language. Assuming the availability of non-expressive speech data in that language, we propose a 3-step technology: 1) we train an F0-conditioned voice conversion (VC) model as data augmentation technique; 2) we train an F0 predictor to control the conversational flavour of the voice-converted synthetic data; 3) we train a TTS system that consumes the augmented data. We prove that our technology enables F0 controllability, is scalable across speakers and languages and is competitive in terms of naturalness over a state-of-the-art baseline model, another augmented method which does not make use of F0 information.

**Index Terms:** text-to-speech, expressive speaking styles, F0 prediction, voice conversion, data augmentation

## 1. Introduction

Neural text-to-speech (TTS) has been shown to produce high-quality synthesised speech from written text. State-of-the-art TTS models usually require a large amount of high-quality speech data [1–8]. Traditionally, these data were recorded in a flat, neutral style, however, with the advancement in TTS research, a new interest in various speech styles has arisen [9–13]. Thus, in order to generate a varied portfolio of TTS voices which covers voices in multiple styles and languages, a large quantity of data is needed. Consequently, in an effort to mitigate data collection, which is costly and time consuming, interest in low-resource TTS research has grown significantly. The two main research trends of low-resource TTS focus on: 1) transfer learning to a low-resource voice, leveraging a large multi-speaker dataset [14–24]; 2) data augmentation, leveraging multi-speaker data [25], voice conversion (VC) capabilities [12, 13], and TTS [26]. The main limitation of these approaches is the assumption that a large amount of data are available for any language and style. In order to make the technology scalable, we propose a language-agnostic method which relies on a minimal amount of such data, regardless of its style. Our work is inspired by [12], which shows that it is possible to build high-quality voices with only a small amount of speech data (*target data*). [12] leverages an unlimited amount of data from other voices (*supporting data*), which are converted to the target voice via VC. A multi-speaker TTS model is then trained with the original and augmented data. We remove the main constraint of [12], i.e. the assumption that a large amount of expressive data are available for any language; we assume instead to have flat, inexpressive, *neutral-style* supporting data. We also reduce the amount of supporting data from an “unlim-

ited amount” to 8-10 hours. Further, our low-resource target data are recorded in an expressive, friendly and spontaneous style, which we call *conversational style*. We show that our technology is language-agnostic, building 9 voices in 5 locales, while using the same configuration of hyper-parameters. As in [12] we perform data augmentation via VC, using a model based on Copycat [27]. Since we have supporting neutral data, but we would like the augmented data to retain the conversational style of the target voice, we enable the possibility of controlling the output speaker style during VC, leveraging log-F0 (F0) information. Controlling prosody with F0 information has been extensively explored in both VC [28] (especially for emotion conversion [29–32]) and TTS literature [4, 8, 14, 33, 34]. We learn how to model F0 through a separate model, which is used to produce F0 trajectories in the expressive, personalised style of the target speaker. The output of such model is then used to feed the VC model at inference time. Our 3-steps pipeline consists of: 1) an F0-conditioned VC model to do data augmentation; 2) an F0 predictor that controls the conversational style of the voice-converted synthetic data; 3) a TTS system that consumes the augmented data. The main contributions of this paper are: 1) an F0 predictor model that produces reliable F0 trajectories; 2) a method to control the style of the augmented samples consumed by the TTS model; 3) a scalable, language-agnostic technology for building high-quality, low-resource TTS voices.

## 2. Methodology

Our goal is to build low-resource, neural TTS voices with only 1 hour of conversational speech, when no other conversational data are available in the same language. In order to do so, we additionally require 8-10 hours of neutral supporting data from another voice in the same gender and locale as the target voice. Given these conditions, we propose a 3-step technology. Firstly, a modified version of Copycat [27] (VC) which makes use of F0 trajectories is trained in order to augment the target speaker data. Secondly, a neural F0 predictor learns to predict F0 trajectories, which are in turn injected into the VC system at inference time. In this way we leverage the supporting data, which are converted to the target speaker identity, but also the target data, from which we learn the conversational style that we want to maintain during VC. Finally, a single speaker TTS model is trained with the original and synthetic data.

### 2.1. F0-conditioned Voice Conversion

Our proposed VC approach (Figure 1) is a modification of CopyCat [27], with two additions: the concatenation of the speaker embeddings to the upsampled phonemes (before passing them to the phoneme encoder), as in [12]; frame-level F0 as input to the decoder, similarly to [28], except for the fact that we do not apply normalisation. The idea behind this architec-




Figure 1: Voice Conversion + TTS pipeline

ture is to learn to disentangle between linguistic information, prosody content and speaker identity, so that VC is performed effectively. [27] uses a variational auto-encoder (VAE) [35] to preserve the supporting data prosody, downsampling and then upsampling the latent spectrogram representation along the time dimension. If the supporting speaker is recorded in a neutral, inexpressive style, her prosody will be passed onto the target speaker during inference. However, in our use-case we would like to maintain the conversational, spontaneous style of the target speaker. To be able to control the style of the converted samples at inference time, we condition the decoder with  $F_0$ , interpolating it over the unvoiced phonemes. Any other prosodic information which is not  $F_0$  and duration is still going to be carried over by the reference encoder. The model is trained on the target and supporting speakers’ data for 100k steps with batch size of 32, and then fine-tuned on the target speaker data for an additional 4k steps. Fine-tuning has been shown to remove the potential remaining speaker leakage [12], without degrading the quality of the synthetic samples.  $F_0$  is extracted by the WORLD vocoder [36]. We use a Kullback–Leibler divergence (KLD) loss for training the VAE reference encoder and an L1 loss is computed between the oracle and the predicted mel-spectrogram.

## 2.2. $F_0$ predictor

As mentioned in the previous section, our aim is to generate synthetic data which mimic the conversational style of the target speaker voice. The architecture of the VC system (Section 2.1) allows us to control prosody at inference time, conditioning the decoder with an  $F_0$  trajectory. Since we do not want to maintain the neutral-style  $F_0$  trajectories of the supporting speaker, we learn to reproduce the target speaker  $F_0$  by training a separate  $F_0$  predictor. The  $F_0$  predictor output is then used in the VC system at inference time. The model architecture (Figure 2) consists of: a Tacotron 2 text encoder [5] which, similarly to the VC system, receives as input state-level phonetic information upsampled at frame level; a parallel decoder that, given the encoded phonemes and a one-hot speaker ID flag, is trained to predict the  $F_0$  trajectories. The  $F_0$  predictor is trained for 10k steps on the neutral supporting and conversational target speaker data, and fine-tuned for another 10k steps on the target




Figure 2:  $F_0$  predictor model architecture

speaker data only. Interpolation of  $F_0$  is used for assigning an  $F_0$  value to the unvoiced phonemes. An L1 loss is computed between the oracle and the predicted  $F_0$  trajectories.

## 2.3. Text-to-Speech

Our TTS model is based on Tacotron 2 [5]. As highlighted in Figure 1, it consists of a phonetic encoder, an additional VAE [35] that captures prosodic information [37, 38], and an auto-regressive decoder. A single-head location-attention mechanism [39] is used between the encoder and the decoder. Our TTS systems are trained on 1 hour of target speaker real data and 8-10 hours of synthetic data, generated by the  $F_0$ -conditioned VC system (where the  $F_0$  trajectory is the one predicted by the  $F_0$  predictor of Section 2.2). As in [12], we additionally fine-tune the TTS models for 4k steps on the target speaker data. Fine-tuning guides the model towards the space of the real data, which results in a refinement in signal and segmental quality, as well as expressiveness. At inference time we condition the decoder with a VAE z-vector computed averaging the target speaker real data. Finally, the predicted mel-spectrograms are converted to audio waveforms by a Parallel Wavenet universal neural vocoder [40].

## 3. Experiments

In this section we assess the strengths of our approach, performing objective analyses of the components of the data augmentation pipeline, as well as perceptual evaluations on the final TTS voices. More specifically, we aim to demonstrate that: 1) the  $F_0$ -predictor model produces conversational  $F_0$  trajectories; 2) we are able to achieve a good level of  $F_0$  controllability in VC; 3) the voice-converted samples resemble the target speaker  $F_0$  distribution; 4) our proposed technology is able to build low-resource TTS voices and either improves or is on par with the previous best augmented system [12] across different languages.

**Experimental setting.** To demonstrate that our technology is scalable in terms of languages and voices, we build 9 conversational TTS voices<sup>1</sup>, of which 5 are females and 4 are males, in 5 different locales (Canadian French, French, Italian,

<sup>1</sup>Samples available on [www.amazon.science/low-data-no-problem](http://www.amazon.science/low-data-no-problem)

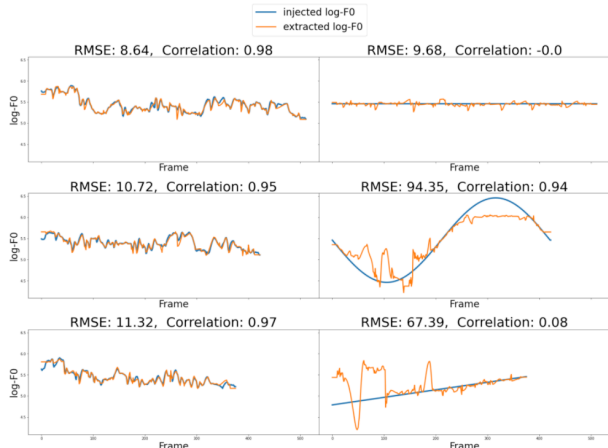


Figure 3: *F0 controllability*. The blue lines correspond to the *F0* trajectories used to condition the decoder in VC, while the orange ones are extracted from the synthetic samples.

German and Spanish), using only 1 hour of data for each target voice. In each experiment we additionally use 8-10 hours of supporting speaker data, i.e. neutral recordings from a voice of the same gender as the target speaker. The data used in these experiments have been professionally recorded and are part of our internal dataset. Every model has been trained in parallel on 4 GPUs (Tesla V100-SXM2-32GB) with batch size of 32, and the same hyper-parameters have been used across all the experiments. The test set consists of 200 samples per voice. For each target voice we also build a baseline system based on [12], which is trained in the same fashion as the proposed one, with the only difference that the synthetic data have been generated with a VC conversion system that does not make use of *F0* (therefore the *F0*-predictor is not needed). Everything else besides the *F0*-conditioning component is as described in Sections 2.1 and 2.3.

### 3.1. Objective analysis

***F0* predictor quality assessment.** Following the procedure described in Section 2.2, we train an *F0* predictor for each one of the 9 target voices. Every model is trained on target and supporting speaker pairs. We use the Root Mean Square Error (RMSE) and Pearson’s Correlation Coefficient (Correlation) as metrics to compare the predicted *F0* trajectories and the oracle ones in the log domain. The objective metrics are computed on an interpolated *F0*. Across all the voices, we get an average RMSE of 49.6 and 0.65 for Correlation. We will explain the implication of the results later in this section

***F0* controllability in VC.** The main advantage of an *F0*-conditioned model is the ability to gain control over *F0* at inference time, which is crucial for handling speaking styles. In this analysis we show how the proposed VC system (Section 2.1) is able to follow the injected *F0* trajectory. The left column of Figure 3 presents three examples of injected trajectories (in blue) together with their counterpart extracted from synthetic audio files (in orange). The injected trajectories come from the *F0* predictor model from Section 2.2. In the second column of Figure 3, we show how, for the same source samples of the left column, the extracted *F0* changes when we inject artificially constructed trajectories. Our model does well in following realistic *F0*s, while it struggles with the artificial trajectories. It can be argued that the decoder is trained to reconstruct somewhat

“intelligible” speech, while the artificial trajectories might not follow the short-term *F0* variation that depends on the phonetic sequence. The result is essentially a compromise between an attempt to both reconstruct speech and follow the *F0* trajectory. Over 200 converted samples, the average RMSE between the injected and extracted *F0* curves is 12.01 and the average Correlation is 0.94. As a reference, we compute the same values when injecting the flat line, the sinusoidal curve and the linear ascendent line, obtaining, as average over the 200 samples, 13.56, 106.42 and 76.87 for RMSE, and 0, 0.86 and 0 for average Correlation, respectively. Note that from a mathematical perspective, a correlation of 0 between the flat linear trajectories and the extracted ones is expected.

***F0* distribution in the synthetic data.** In this experiment we show that our VC technology generates synthetic samples that reflect the speaking style of the target speaker. We analyse the distributions of *F0* and delta-*F0* values of both *F0*-conditioned and non *F0*-conditioned VC models. We obtain the delta-*F0* values by computing the first derivative on the *F0* trajectories. Delta-*F0* encodes the speed of change which we believe is an important feature for characterising styles. To perform this analysis, we generate synthetic data with three data augmentation techniques: 1) *No-F0 VC* is Copycat as in the baseline [12], without *F0* conditioning; 2) *Orig-F0 VC* is the proposed VC model (Section 2.1), where at inference time the decoder is conditioned with the original *F0* of the supporting speaker, rescaled by the mean of the target speaker; 3) *Pred-F0 VC* is *Orig-F0 VC*, where at inference time the decoder is conditioned with the *F0* trajectory produced by the *F0* predictor of Section 2.2.

We compare the *F0* distributions of the synthetic samples: firstly, we convert the *F0* values into probability distributions; then we compute the KLD of each system with respect to the target speaker *F0* (and delta-*F0*) distributions. We perform this analysis over 4 speakers from 4 locales (Canadian French, French, Italian, Spanish). We chose these speakers since they have a wide range of absolute distance in terms of *F0* mean and variance with respect to their respective supporting speakers. Per-speaker *F0* mean and variance have been computed across the overall speaker data in the linear frequency scaling. We observed that if a target speaker has a much higher *F0* variance with respect to the supporting speaker, i.e. the target speaker is more expressive than the supporting one, it is usually harder to produce synthetic data as expressive as the target voice. Having the control over *F0* can help to mitigate this effect.

In Table 1 we observe that, for Speaker 1, *Pred-F0 VC* has the smallest KLD value with respect to the real data, while for Speaker 4, the speaker which is the closest to her supporting counter-part, the best system is *Orig-F0 VC*. Generally, *Pred-F0 VC* gets small KLD values for the delta-*F0* distributions. Our intuition is that when the supporting and the target speakers are close in terms of *F0* behaviour, using the supporting speaker original *F0* (rescaled by the target speaker *F0* mean) is enough to resemble the *F0* distribution of the target speaker (see Speaker 4). When the speakers are further apart, as in the case of Speaker 1, having an *F0*-predictor can have a high impact. On the other hand, *Orig-F0 VC* is also the model that leads to higher fluctuations in terms of KLD scores, which indicates that its performances are less controllable, since we do not put any restrictions on the supporting speaker except for being of the same locale and gender of the target voice. *No-F0 VC* can also lead to low KLD values, however, as in the case of *Orig-F0 VC*, it can be unpredictable. Ultimately, *Pred-F0 VC* gives us the advantage of controllability and shows consistently low

Table 1: F0 distribution analysis. We highlight the differences in terms of F0 mean and variance between the target speakers and their respective supporting speaker, since it has an impact on the models’ performances. The KLD values are computed between the F0/delta-F0 distributions of synthetic samples with respect to the target data F0/delta-F0 distributions. Values closer to 0 are preferred.

	Difference in F0 mean (wrt Target)	Difference in F0 variance (wrt Target)	KLD (wrt Target F0   delta-F0 distributions)					
			No-F0 VC		Orig-F0 VC		Pred-F0 VC	
<b>Speaker 1 (it)</b>	+47.6	+1840.1	0.27	<b>0.02</b>	0.64	<b>0.02</b>	<b>0.10</b>	<b>0.02</b>
<b>Speaker 2 (fr-ca)</b>	+31.4	+1637.8	<b>0.05</b>	<b>0.01</b>	0.10	0.02	0.08	<b>0.01</b>
<b>Speaker 3 (fr)</b>	-9.27	+980.39	<b>0.06</b>	0.04	0.32	0.08	<b>0.06</b>	<b>0.03</b>
<b>Speaker 4 (es)</b>	+29.7	+484.6	0.09	0.02	<b>0.01</b>	<b>0.01</b>	0.06	0.02

Table 2: Results for the Naturalness MUSHRA evaluation. TTS from No-F0 VC is the baseline system; TTS from F0 VC is the proposed system; the p-value comes from the Student’s t-test and it is computed between the baseline and the proposed system; the % of improvement is the relative improvement of the proposed system with respect to the baseline and the recordings.

Locale	Gender	Recordings	TTS from No-F0 VC	TTS from F0 VC	p-value	% improvement
<b>Canadian French</b>	<b>Female</b>	72.09±1.39	68.27±1.35	67.28±1.38	0.1801	/
	<b>French</b>	<b>Female</b>	68.2±0.94	61.27±0.94	62.03±0.94	0.0953
<b>Italian</b>	<b>Male</b>	68.67±0.89	56.51±0.96	59.14±0.93	<b>0.0</b>	21.70%
	<b>Female</b>	79.84±0.74	61.96±0.85	63.35±0.86	<b>0.0001</b>	7.80%
<b>German</b>	<b>Male</b>	79.01±0.73	61.97±0.84	62.16±0.84	0.5791	/
	<b>Female</b>	75.47±0.82	63.50±0.83	62.82±0.84	0.1118	/
<b>Spanish</b>	<b>Male</b>	75.32±0.81	57.32±0.96	62.23±0.83	<b>0.0</b>	27.30%
	<b>Female</b>	75.49±0.79	67.22±0.87	68.45±0.86	<b>0.0028</b>	17.70%
	<b>Male</b>	79.62±0.71	62.25±0.9	63.73±0.88	<b>0.0001</b>	8.50%

KLD values across speakers. We believe that controlling the F0 generation with an high-quality F0 predictor is a key factor in order to consistently produce synthetic data that resemble the target voice’s typical prosody and speaking patterns.

### 3.2. Perceptual evaluations

We design the experiments in order to highlight how our technology: 1) scales across languages; 2) scales across voices; 3) outperforms our baseline, the state-of-the-art data augmentation approach which does not make use of F0 [12]. A few additional observations: for space reasons, we did not provide results for the TTS systems augmented with data from *Orig-F0 VC* (the VC system which uses the supporting speaker rescaled F0 to condition the decoder), since they were generally worse than the proposed approach; we also did not use a multi-speaker model as a baseline since [12] already showed that data augmentation leads to significant improvements. Our experiments confirmed the finding.

**Evaluation protocol.** We evaluate our models running MUSHRA (MUltiple Stimuli with Hidden Reference and Anchor)-like [41] tests on Naturalness. In the evaluations we omit the reference system, while we use the original recordings, as upper anchor, and the baseline system based on [12]. Each one of the 200 test-cases is evaluated by 15 different users. We apply a mechanism to get the fairest possible results: we select the test utterances to be mainly long utterances, i.e. more than 95% of the utterances contain 10 or more words. This gives the listeners a higher chance of detecting bad TTS samples. The issue with short utterances is that there is less chance of detecting incorrect intonation, rhythm, etc.

**Results.** In Table 2 we show the results for the MUSHRA Naturalness evaluations for 9 synthetic voices. We can observe that our proposed data augmentation approach, *TTS from F0 VC*, statistically improves the baseline, *TTS from no-F0 VC*, for

5 out of 9 voices. In the remaining 4 cases, our approach is on par with the baseline. On average, when there is an improvement, our models close the gap between the baselines and the recordings by 16.6%. For two of the voices where we are on par with the baseline (the Canadian French and French female voices), the proposed approach is very close to recordings, being respectively at 94.5% and 90.5% of Naturalness relative to recordings, making it hard to beat the already high performing baseline system. These results indicate that, generally, having the control on F0 during the data augmentation phase leads to more natural TTS voices. Additionally, we demonstrate that we are able to get high-quality voices across female and male speakers belonging to 5 different locales. On average, the proposed method scores 85.92% in terms of Naturalness with respect to natural recordings.

## 4. Conclusion and future work

In this work we present a language-agnostic methodology to build low-resource TTS conversational voices. The proposed technology augments the low-resource target data via a VC system that enables F0 controllability. An external F0 predictor ensures that the augmented samples resemble the conversational style of the target speaker. Finally, a TTS model consumes the original and augmented data. The ability of learning and controlling a speaker expressiveness means that we do not need to rely on any other conversational supporting data to build a TTS voice. Therefore our technology is largely scalable to voices and languages. Our experiments demonstrate that: 1) the F0-predictor model produces reliable F0 trajectories; 2) we are able to achieve a high level of F0 controllability in VC; 3) the voice-converted samples resemble the target speaker F0 distribution; 4) we produce high-quality TTS voices from low-resource data, being on par or improving the previous best augmented system; 5) our technology is language-agnostic, hence largely scalable.

## 5. References

- [1] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," *CoRR*, 2016.
- [2] A. van den Oord, Y. Li, I. Babuschkin, K. Simonyan, O. Vinyals, K. Kavukcuoglu, G. van den Driessche, E. Lockhart, L. C. Cobo, F. Stimberg, N. Casagrande, D. Grewe, S. Noury, S. Dieleman, E. Elsen, N. Kalchbrenner, H. Zen, A. Graves, H. King, T. Walters, D. Belov, and D. Hassabis, "Parallel wavenet: Fast High-Fidelity Speech Synthesis," in *ICML*, 2018.
- [3] J. Sotelo, S. Mehri, K. Kumar, J. F. Santos, K. Kastner, A. C. Courville, and Y. Bengio, "Char2Wav: End-to-End Speech Synthesis," in *ICLR*, 2017.
- [4] S. Ö. Arik, M. Chrzanowski, A. Coates, G. F. Diamos, A. Gibiansky, Y. Kang, X. Li, J. Miller, A. Y. Ng, J. Raiman, S. Sengupta, and M. Shoenybi, "Deep Voice: Real-time Neural Text-to-Speech," in *ICML*, 2017.
- [5] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. J. S. Ryan, R. A. Saurous, Y. Agiomyriannakis, and Y. Wu, "Natural TTS synthesis by conditioning wavenet on MEL spectrogram predictions," in *ICASSP*, 2018.
- [6] W. Ping, K. Peng, and J. Chen, "ClariNet: Parallel Wave Generation in End-to-End Text-to-Speech," in *ICLR*, 2019.
- [7] Y. Zheng, X. Li, F. Xie, and L. Lu, "Improving End-to-End Speech Synthesis with Local Recurrent Neural Network Enhanced Transformer," in *ICASSP*, 2020.
- [8] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T. Liu, "FastSpeech 2: Fast and High-Quality End-to-End Text to Speech," in *ICLR*, 2021.
- [9] Y. Wang, D. Stanton, Y. Zhang, R. J. Skerry-Ryan, E. Battenberg, J. Shor, Y. Xiao, Y. Jia, F. R., and R. A. Saurous, "Style Tokens: Unsupervised Style Modeling, Control and Transfer in End-to-End Speech Synthesis," in *ICML*, 2018.
- [10] R. J. Skerry-Ryan, E. Battenberg, Y. Xiao, Y. Wang, D. Stanton, J. Shor, R. J. Weiss, R. Clark, and R. A. Saurous, "Towards End-to-End Prosody Transfer for Expressive Speech Synthesis with Tacotron," in *ICML*, 2018.
- [11] K. Lee, K. Park, and D. Kim, "STYLER: Style Factor Modeling with Rapidity and Robustness via Speech Decomposition for Expressive and Controllable neural text-to-speech," in *Interspeech*, 2021.
- [12] G. Huybrechts, T. Merritt, G. Comini, B. Perz, R. Shah, and J. Lorenzo-Trueba, "Low-Resource Expressive Text-To-Speech Using Data Augmentation," in *ICASSP*, 2021.
- [13] R. Shah, K. Pokora, A. Ezger, V. Klimkov, G. Huybrechts, B. Putrycz, D. Korzekwa, and T. Merritt, "Non-Autoregressive TTS with Explicit Duration Modelling for Low-Resource Highly Expressive Speech," *CoRR*, 2021.
- [14] S. Ö. Arik, G. F. Diamos, A. Gibiansky, J. Miller, K. Peng, W. Ping, J. Raiman, and Y. Zhou, "Deep Voice 2: Multi-Speaker Neural Text-to-Speech," *CoRR*, 2017.
- [15] W. Ping, K. Peng, A. Gibiansky, S. Ö. Arik, A. Kannan, S. Narang, J. Raiman, and J. Miller, "Deep Voice 3: Scaling Text-to-Speech with Convolutional Sequence Learning," in *ICLR*, 2018.
- [16] Y. Taigman, L. Wolf, A. Polyak, and E. Nachmani, "VoiceLoop: Voice Fitting and Synthesis via a Phonological Loop," in *ICLR*, 2018.
- [17] S. Ö. Arik, J. Chen, K. Peng, W. Ping, and Y. Zhou, "Neural voice cloning with a few samples," in *NeurIPS*, 2018.
- [18] Y. Lee, T. Kim, and S. Lee, "Voice Imitating Text-to-Speech Neural Networks," *CoRR*, 2018.
- [19] Y. Jia, Y. Zhang, R. J. Weiss, Q. Wang, J. Shen, F. R., Z. Chen, P. Nguyen, R. Pang, I. Lopez-Moreno, and Y. Wu, "Transfer Learning from Speaker Verification to Multispeaker Text-To-Speech Synthesis," in *NeurIPS*, 2018.
- [20] Y. Chen, T. Tu, C. Yeh, and H. Lee, "End-to-End Text-to-Speech for Low-Resource Languages by Cross-Lingual Transfer Learning," in *Interspeech*, 2019.
- [21] N. Tits, K. El Haddad, and T. Dutoit, "Exploring Transfer Learning for Low Resource Emotional TTS," in *IntelliSys*, 2019.
- [22] H. Zhang and Y. Lin, "Unsupervised Learning for Sequence-to-Sequence Text-to-Speech for Low-Resource Languages," in *Interspeech*, 2020.
- [23] E. Cooper, C. Lai, Y. Yasuda, F. Fang, X. Wang, N. Chen, and J. Yamagishi, "Zero-Shot Multi-Speaker Text-To-Speech with State-Of-The-Art Neural Embeddings," in *ICASSP*, 2020.
- [24] E. Casanova, C. Shulby, E. Gölge, N. M. Müller, F. S. de Oliveira, A. Cândido Júnior, A. da Silva Soares, S. M. Aluísio, and M. A. Ponti, "SC-GlowTTS: an Efficient Zero-Shot Multi-Speaker Text-To-Speech Model," *CoRR*, 2021.
- [25] Y. Jia, Y. Zhang, R. J. Weiss, Q. Wang, J. Shen, F. R., Z. Chen, P. Nguyen, R. Pang, I. Lopez-Moreno, and Y. Wu, "Transfer Learning from Speaker Verification to Multispeaker Text-To-Speech Synthesis," in *NeurIPS*, 2018.
- [26] M. Hwang, R. Yamamoto, E. Song, and J. Kim, "TTS-by-TTS: TTS-Driven Data Augmentation for Fast and High-Quality Speech Synthesis," in *ICASSP*, 2021.
- [27] S. Karlapati, A. Moinet, A. Joly, V. Klimkov, D. Sáez-Trigueros, and T. Drugman, "CopyCat: Many-to-Many Fine-Grained Prosody Transfer for Neural Text-to-Speech," in *Interspeech*, 2020.
- [28] K. Qian, Z. Jin, M. Hasegawa-Johnson, and G. J. Mysore, "F0-Consistent Many-To-Many Non-Parallel Voice Conversion Via Conditional Autoencoder," in *ICASSP*, 2020.
- [29] Z. Luo, J. Chen, T. Takiguchi, and Y. Arik, "Emotional Voice Conversion Using Dual Supervised Adversarial Networks With Continuous Wavelet Transform F0 Features," *IEEE ACM Trans. Audio Speech Lang. Process.*, vol. 27, no. 10, 2021.
- [30] K. Zhou, B. Sisman, and H. Li, "Transforming Spectrum and Prosody for Emotional Voice Conversion with Non-Parallel Training Data," in *Odyssey*, 2020.
- [31] K. Zhou, B. Sisman, R. Liu, and H. Li, "Seen and Unseen Emotional Style Transfer for Voice Conversion with A New Emotional Speech Dataset," in *ICASSP*, 2021.
- [32] C. Le Moine, N. Obin, and A. Roebel, "Towards end-to-end F0 voice conversion based on Dual-GAN with convolutional wavelet kernels," in *EUSIPCO*, 2021.
- [33] A. Lancucki, "Fastpitch: Parallel Text-to-Speech with Pitch Prediction," in *ICASSP*, 2021.
- [34] R. Valle, J. Li, R. Prenger, and B. Catanzaro, "Mellotron: Multi-speaker Expressive Voice Synthesis by Conditioning on Rhythm, Pitch and Global Style Tokens," in *ICASSP*, 2020.
- [35] D. P. Kingma and J. Welling, "Auto-Encoding Variational Bayes," in *ICLR*, 2014.
- [36] M. Morise, F. Yokomori, and K. Ozawa, "WORLD: A vocoder-based high-quality speech synthesis system for real-time applications," *IEICE Trans. Inf. Syst.*, vol. 99-D, no. 7, 2016.
- [37] Y. Zhang, S. Pan, L. He, and Z. Ling, "Learning Latent Representations for Style Control and Transfer in End-to-end Speech Synthesis," in *ICASSP*, 2019.
- [38] S. Tyagi, M. Nicolis, J. Rohnke, T. Drugman, and J. Lorenzo-Trueba, "Dynamic prosody generation for speech synthesis using linguistics-driven acoustic embedding selection," in *Interspeech*, 2020.
- [39] D. Bahdanau, K. Cho, and Y. Bengio, "Neural Machine Translation by Jointly Learning to Align and Translate," in *ICLR*, 2015.
- [40] Y. Jiao, A. Gabrys, G. Tinchev, B. Putrycz, D. Korzekwa, and V. Klimkov, "Universal neural vocoding with parallel wavenet," in *ICASSP*, 2021.
- [41] B. Series, "Method for the subjective assessment of intermediate quality level of audio systems," *International Telecommunication Union Radiocommunication Assembly*, 2014.