

CIArTTS: An Open-Source Classical Arabic Text-to-Speech Corpus

Ajinkya Kulkarni✉, Atharva Kulkarni, Sara Abedalmon'em Mohammad Shatnawi✉, Hanan Aldarmaki✉

MBZUAI UAE✉, Erisha Labs India

Ajinkya.Kulkarni@mbzuai.ac.ae, atharva7kulkarni@gmail.com, sara.shatnawi@mbzuai.ac.ae, Hanan.Aldarmaki@mbzuai.ac.ae

Abstract

At present, Text-to-speech (TTS) systems that are trained with high-quality transcribed speech data using end-to-end neural models can generate speech that is intelligible, natural, and closely resembles human speech. These models are trained with relatively large single-speaker professionally recorded audio, typically extracted from audiobooks. Meanwhile, due to the scarcity of freely available speech corpora of this kind, a larger gap exists in Arabic TTS research and development. Most of the existing freely available Arabic speech corpora are not suitable for TTS training as they contain multi-speaker casual speech with variations in recording conditions and quality, whereas the corpus curated for speech synthesis are generally small in size and not suitable for training state-of-the-art end-to-end models. In a move towards filling this gap in resources, we present a speech corpus for Classical Arabic Text-to-Speech (CIArTTS) to support the development of end-to-end TTS systems for Arabic. The speech is extracted from a LibriVox audiobook, which is then processed, segmented, and manually transcribed and annotated. The final CIArTTS corpus contains about 12 hours of speech from a single male speaker sampled at 40100 kHz. In this paper, we describe the process of corpus creation and provide details of corpus statistics and a comparison with existing resources. Furthermore, we develop two TTS systems based on Grad-TTS and Glow-TTS and illustrate the performance of the resulting systems via subjective and objective evaluations. The corpus will be made publicly available at www.clartts.com for research purposes, along with the baseline TTS systems demo.

Index Terms: arabic speech corpus, text-to-speech

1. Introduction

Neural text-to-speech (TTS) models are becoming mainstream due to their superior performance in synthesizing intelligible and natural-sounding speech. Compared to older concatenative (e.g. [1]) or HMM-based [2] TTS models, neural models can generate raw waveform directly from text inputs without complex pre-processing and phonetic feature extraction. Neural TTS models commonly have two main components: an acoustic model that generates acoustic features (e.g. mel-spectrograms) directly from text, and a vocoder to generate a waveform from the acoustic features (see for example [3]). Fully end-to-end TTS models that combine both stages have also been explored [4]. While these neural architectures can be complex, end-to-end training alleviates the need for feature engineering and other design choices that are prone to be suboptimal. One of the bottlenecks in TTS system design, however, is the availability and quality of the corpus used for training. Unlike ASR datasets, where it is desirable to have a variety of speakers and recording conditions to achieve robust performance, it is far more avan-

tageous to have consistent single-speaker corpora for TTS to achieve intelligible and natural sounding synthesis. Therefore, speech data used for training TTS models need to have more consistent acoustic features that ideally only vary along phonetic and prosodic dimensions.

Most existing corpora for Arabic TTS are carefully designed and reduced datasets that are optimized for phonetic coverage while maintaining a relatively small number of units [5][6]. This choice is partially a remnant of early concatenative models that have a real-time computational cost proportional to the size of the dataset. Another reason for this choice is the relative difficulty of constructing consistent datasets that are suitable for TTS training, especially if they need to be annotated at the phonetic level for traditional TTS systems, so a reduced dataset that maintains phonetic coverage is more manageable to construct. For example, one of the most commonly used public TTS datasets for Arabic is the Arabic Speech Corpus (ASC) [5], which has around 3.4 hours of speech. The ASC was designed to maximize phonetic coverage using a greedy optimization strategy. While such optimization technique is the most commonly used in most TTS data construction projects, there is some evidence that a random subset of the same size could potentially lead to similar or even more natural-sounding speech synthesis [7]. In addition, for neural TTS models, quantity is more beneficial to the overall quality of the synthesized speech as they are more robust to small variations in input conditions. Moreover, neural TTS models can work directly with text utterances as input without the need for phonetic annotations, which makes the construction of larger datasets more feasible.

In this work, we construct a relatively large single-speaker corpus for the purpose of developing neural TTS systems for Arabic. In particular, the corpus consists of audio recordings by a male speaker of a book written in Classical Arabic. The audiobook is publicly available in the LibriVox project. To create a corpus for text-to-speech synthesis, we segmented the corpus into short utterances, checked for quality and consistency of recording conditions, then manually annotated the audio segments with fully diacritized transcriptions. Samples can be found at clartts.com. We will make the corpus available publicly for research use. As text transcripts were not available for Arabic audiobooks, we had to perform a manual annotation process to create the CIArTTS corpus. This corpus comprises 12 hours and 10 minutes of speech, consisting of 10,334 utterances from a single male speaker, and was sampled at 40,100 Hz. We also build several neural TTS systems using this corpus and demonstrate the quality of the synthesized speech using subjective and objective evaluations. We show the synthesis performance for both Classical and Modern Standard Arabic. Furthermore, we show the performance of the models using raw character inputs vs. phonetic inputs using a rule-based grapheme-to-phoneme algorithm.

This paper is organized as follows: the first section gives a brief overview of the related works. Then, corpus construction provides the details of building CIArTTS corpus from audiobooks and the annotation process used for it. In section 4, we present the corpus statistics and comparison of the CIArTTS corpus with existing Arabic speech synthesis corpora. We created baseline TTS systems on two Arabic speech synthesis corpus using Glow-TTS and Grad-TTS as described in Section 5. Furthermore, we also explained the experimentation setup along with the evaluation approach to estimate the performance of TTS systems in Section 6, followed by the conclusion in Section 7.

2. Related Work

Currently, Arabic speech synthesis systems are of lower quality compared to their English counterparts, largely due to the limited availability of Arabic speech synthesis data in corpora [8]. The most commonly used approaches for Arabic speech synthesizers are either based on unit selection or parametric speech synthesis [9, 10, 11]. However, many speech synthesis corpora for English have been developed using audiobooks for which text transcripts are readily available, whereas Arabic audiobooks lack these transcripts, making it difficult to develop speech synthesis corpora [8].

In this study, we present the CIArTTS corpus, which is based on an audiobook with manually annotated text transcripts. The Arabic Speech Corpus (ASC) contains around 3.4 hours of south Levantine Arabic speech recorded at 48KHz using fully diacritized text collected from Aljazeera Learn, a language learning website [12]. Diphone-based greedy optimization strategies were used to reduce the size of the transcripts, and non-sense or dummy utterances were recorded to cover the gaps of underrepresented phonemes.

Another approach proposed a fully unsupervised framework to build a TTS system using broadcast news recordings [13]. They used both manual and automatic dataset selection and transfer learning by using high-resource languages in the TTS model from the LJSpeech dataset and fine-tuned it with one hour of Arabic speech. In [14] NatiQ, a Tacotron 2 [15] based Arabic TTS system, high-quality speech data was recorded at a sampling rate of 44kHz from two speakers. In another study, a pre-recorded Audiobook from the Masmoo3 Audiobooks website was used to create a 4-hour Arabic speech synthesis corpus for TTS applications [16]. The balanced Arabic speech corpus was explicitly designed to ensure phonetically balanced Arabic speech (BAC), which was specifically designed for the unit selection and rule-based speech synthesis approach [17]. The main objective of the BAC corpus was to ensure that all potential phonemes and some impossible phoneme combinations between words were included.

3. Corpus Construction

In this section, we describe the steps involved in building our Classical Arabic speech synthesis corpus: audio pre-processing, the annotation process, final corpus creation, and corpus statistics.

3.1. Audio Pre-processing

For the creation of a classical Arabic text-to-speech (CIArTTS) corpus, we selected an audiobook recorded by a single speaker

from the LibriVox project¹. The classical book is titled *Kitab Adab al-Dunya w'al-Din* by Abu al-Hasan al-Mawardi (972-1058 AD). The audiobook is recorded by a single speaker and consists of approximately 16 hours of audio without accompanying text. While scanned copies of the book exist, we opted for manual annotation of the audio data to create text transcripts that truly match the audio recording using the Praat annotation tool².

The audiobook consists of 20 long audio files, each representing a chapter of the book in MP3 format. We converted this audio to WAV format using *ffmpeg* command-line tool to ensure compatibility with the Praat program. We kept the original sampling rate of 40100 Hz. We ran a rule-based Praat script to mark pauses and speech segments in the long audio files. This script created a TextGrid object for a LongSound object and set boundaries at pauses based on intensity analysis. We validated the marking of pauses and speech segments provided by the Praat tool using energy-based VAD from the Kaldi toolkit³.

3.2. Annotation Process

The process of annotating an audiobook involved transcribing audio content into written text, along with additional tags for speech pauses, background noise, inaudible speech segments, and stuttering. The Praat tool was used for the annotation, and the annotators were given TextGrid Praat files that contained the audio recording and a framework for marking speech and pause segments. This helped the annotators efficiently and accurately transcribe the speech segments into written text.

A team of three Arabic annotators was involved in the transcription process to ensure a reliable and accurate final transcript that considered multiple perspectives. To enhance the quality of the transcripts, two rounds of validation were conducted. The first validation was done by the annotators themselves, followed by a check by two other annotators for accuracy and consistency. The text transcripts were marked with Arabic diacritical marks to increase the accuracy of the transcripts for speech analysis and pronunciation.

In addition to the TextGrid Praat files, the annotators were also given a text image of the original book for reference. This made it easier for the annotators to transcribe the speech segments accurately by referring to the original text. Guidelines were provided to the annotators during the annotation process, including instructions for using abbreviations, numbers, special characters, and punctuation according to Arabic language rules. Specific speech segments were marked with tags, including [B] for background noise, [H] for stuttering or hesitation, [*] for unclear speech, and [O] for human noise. The combination of the Praat tool, three annotators, two levels of validation, text transcripts with Arabic diacritization markers, and reference materials helped ensure the accuracy and reliability of the final transcripts.

3.3. Final Corpus Creation

The total amount of original audio is around 16hrs, spanning 20 chapters, so it was recorded in multiple sessions. We observed slight variations in speaking style between the chapters, even though it was neutral (non-emotional) overall. Therefore, we conducted subjective listening tests by listening to random parts of each chapter and removed three chapters that diverge in

¹www.LibriVox.org

²<https://www.fon.hum.uva.nl/praat/>

³<https://kaldi-asr.org>

Table 1: *Corpus statistics comparison between Arabic speech corpus (ASC), Balanced Arabic corpus (BAC) and CIArTTS.*

Count	BAC	ASC	CIArTTS
Sentences	202	1,913	10,334
Words	1,254	17,275	82,970
Words/sentence (Avg)	6	9	8
Unique words	975	12,144	27,870
Phonemes	6,174	135,232	518,682
Diphones	3,614	72,797	282,487
Unique diphones	-	682	520

Table 2: *Percentage of a subset of frequent (Top) and infrequent (bottom) diphones in the CIArTTS corpus vs. a larger text corpus (Tashkeela)*

Diphone	CIArTTS	Tashkeela
w-a	3.62	3.21%
l-a	3.09	3.00%
<-a	2.89	3.53%
l-aa	2.82	1.6%
a-l	2.53	1.39%
E-a	2.32	2.34%
m-a	2.24	1.95%
n-aa	2.19	1.03%
u1-S	.00035	.00065%
u1-T	.00035	.00175%
u1-^	.00035	.00034%
i1-T	.00035	.00163%
u1-E	.00035	.00009%
A-j	.00070	.00011%
A-x	.00070	.00082%
u1-g	.00070	.00154%

speaking style compared to the rest. We split each long-audios using the textgrid obtained through the Praat tool and manual annotation process with speech and silence segments. For ensuring high audio quality, we used signal-to-noise ratio (SNR) to guide the selection process. We estimated the waveform amplitude distribution analysis SNR [] by taking into account the noise power in silence (non-speech) segments adjacent to the given speech segment. We used a threshold value of 20dB SNR for the first level of speech segment selection.

We concentrated adjacent speech segments to create a minimum speech segment duration of 2 seconds. Furthermore, during the concatenation process, we kept only 20% of silence segments between two speech segments if the silence segment duration was exceeding the average silence duration computed across the given long audio. We also removed the preamble speech segments, during which the reader briefly talked about the LibriVox project, stated their name and book information, and may have mentioned copyright descriptions or LibriVox project-related content.

During the segmentation process, we ensured that each segmented speech utterance had a duration of at least 2 seconds and a maximum duration of 10 seconds. Furthermore, we also observed that the Praat pause marking script was unable to tag the last silence segments. Therefore, we manually removed the silence frame in the last audio segments marked by Praat tools. We also removed the speech segments consisting of text transcripts with non-Arabic characters.

We used 3.34% of the corpus as the test set and 96.66% as




Figure 1: *Percentage coverage of phonemes for ASC, CIArTTS, and Taksheela corpus*

the training set. All text files were saved in UTF-16 encoding and non-Arabic characters were removed. The number of training speech utterances was 10000, and test speech utterances were 334. The total duration for training data is 11 hours 45 mins and for the test 25 mins.

4. Corpus statistics

The corpora that are recorded specifically for the purpose of speech synthesis typically follow a specific procedure to maximize phonetic coverage while minimizing total corpus size [6]. However, since we do not record the corpus and instead use a pre-existing audiobook, we are constrained only by the size of the audiobook. As a result, CIArTTS may not include all possible phonetic combinations, but instead follows the phonetic distribution of the language. In Figure 1, we stated the comparison of monophone coverage across the three corpora namely Arabic speech corpus (ASC), Arabic diacritizer text corpus, Tashkeela, and [18] presented CIArTTS corpus. Figure 1 indicates similar monophone distribution from text information for all the corpus.

We compare our corpus statistic with the Balanced Arabic Corpus (BAC) described in [6] and the Arabic Speech Corpus (ASC) in Table 1. CIArTTS is the largest corpus in terms of the number of sentences, words, unique words, phonemes, and diphones, indicating that it is a more extensive and diverse corpus than the other two. ASC has the second-largest number of sentences and words, but its unique words, phonemes, and diphones are lower than CIArTTS. BAC is the smallest corpus in terms of all the measures listed in the table, suggesting that it may not be as comprehensive or representative of Arabic speech as the other two corpora. The only statistic where we observe a shortage is the number of unique diphones. In the ASC, dummy utterances are recorded to artificially maximize the total number of diphones, even though these diphones are rare or impossible in the language. Therefore, this shortage in diphone coverage is unlikely to degrade TTS performance for most utterances. In Table 2, we present the percentage of diphone coverage in the CIArTTS corpus and a large text corpus, the Taksheela Arabic diacritization corpus, where diphone symbols are represented using Buckwalter transcription format. It clearly indicates that CIArTTS have diphone coverage similar to the Arabic text corpus for both the most frequent and most infrequent diphone combinations.

Arabic speech corpus displays better coverage for a few phonemes than CIArTTS possibly due to the presence of

Table 3: Evaluation metrics computed to measure the performance of baseline end-to-end TTS systems on two Arabic speech synthesis corpora, namely Arabic speech corpus (ASC) and Classical ArabicTTS corpus (CIArTTS).

System	Corpus	MOS	PESQ	MCD	Lf0RMSE	BAP	Speaker similarity
GroundTruth	ASC	4.01	—	—	—	—	—
GroundTruth	CIArTTS	4.39	—	—	—	—	—
Grad-TTS	ASC	3.02	1.48	6.38	12.25	1.14	0.51
Glow-TTS	ASC	3.19	1.41	6.27	10.03	1.12	0.56
Grad-TTS	CIArTTS	3.63	2.25	4.94	9.03	0.85	0.71
Glow-TTS	CIArTTS	3.84	2.23	4.83	8.04	0.93	0.78

dummy utterances. CIArTTS corpus still has better coverage for the majority of the phonemes naturally.

5. Baseline TTS systems

The goal of our research was to compare the performance of two baseline text-to-speech (TTS) systems, Grad-TTS [19] and Glow-TTS [20], on the CIArTTS corpus and Arabic speech corpus. We used the default network parameters as mentioned in the papers [19] and [20] respectively for these TTS systems without using any explicit Arabic grapheme to phoneme module on text transcripts. We used the train set and test set as discussed in section 3.4 for training baseline TTS systems on ASC and CIArTTS. We trained the Grad-TTS and Glow-TTS systems individually on both corpus for 1000 epochs.

To synthesize the speech from the predicted Mel spectrograms, we opted for a Hi-Fi GAN-based neural vocoder [21]. The ASC and CIArTTS corpora have speech utterances with different sampling rate that is 48000 Hz and 40100 Hz. Therefore, we trained two Hi-Fi GAN neural vocoders to create compatibility with the different sampling rates of both corpus. We used ASC for training Hi-Fi GAN neural vocoder with 48000 Hz, while for 40100 Hz, we used CIArTTS corpus. We used the V1 configuration of the Hi-Fi GAN neural vocoder for training both neural vocoders as detailed in [21]. We applied the short-time Fourier transform (STFT) with an FFT length of 1024, a hop length of 256, and a window size of 1024, and extracted Mel spectrograms using 80 Mel filters.

6. Evaluation and Results

In Table 1, we present the performance of baseline TTS systems and subjective evaluation of ASC and CIArTTS corpus. We evaluated E2E TTS systems using a Mean Opinion Score (MOS) [22] based listening test. Each listener had to assign a score for synthesized speech utterance on a scale between 1 to 5 considering the intelligibility, naturalness, and quality of speech utterance. A total of 30 Arabic listeners participated in this MOS test and results are displayed in Table 1 with an associated 95% confidence interval. Furthermore, To validate the coherence of subjective listening test with objective evaluation, we opted for Perceptual Evaluation of Speech Quality (PESQ) [23] as an automated assessment of audio quality which takes into account various factors such as Audio sharpness, volume, background noise, lag in audio, clipping and audio interference. PESQ is computed on a scale from -0.5 to 4.5, where 4.5 represents the best similarity.

We used MCD (Mel Cepstral Distortion), an objective evaluation metric that measures the spectral distortion between the synthesized speech and the original speech signal. Lf0 RMSE (Root Mean Square Error of Log F0): an objective evaluation

metric that measures the pitch accuracy of synthesized speech. BAP (Band Aperiodicity): an objective evaluation metric that measures the spectral envelope accuracy of synthesized speech. These evaluations are conducted by computing errors between reference speech utterances and synthesized speech utterances aligned using the dynamic time-warping algorithm.

We selected a cosine distance-based speaker similarity score [24] to measure the consistency of the speaker’s voice quality in synthesized speech. We utilized the pre-trained ECAPA-TDNN-based speaker embedding extractor to measure the similarity scores from synthesized speech and reference speech from the original speech synthesis corpus [25].

Table 3 shows that the ground truth samples of both corpora have higher MOS scores than the synthesized speech generated by the two TTS systems. The Glow-TTS system outperforms the Grad-TTS system in terms of MOS and PESQ scores for both corpora. The CIArTTS corpus has higher MOS scores and lower MCD, Lf0 RMSE, and BAP scores than the ASC corpus, indicating that the CIArTTS corpus is easier to synthesize. Finally, the speaker similarity scores of the synthesized speech are relatively low for ASC-based TTS systems, compared to the CIArTTS corpus counterpart. Thus, it shows that CIArTTS-based systems are better at retaining the speaker’s voice characteristics in synthesized speech.

7. Conclusion

In this work, we presented a single-speaker classical Arabic TTS corpus named CIArTTS corpus based on an audiobook in a Male speaker’s voice. The CIArTTS is developed with aiming to facilitate the research in Arabic end-to-end TTS system with a large-scale speech synthesis dataset consisting of a total of 12 hours and 10 mins of annotated speech. Furthermore, we have shown the comparative study on corpus statistics with two Arabic speech synthesis corpora namely Arabic speech corpus (ASC) and balanced Arabic speech corpus (BAC). We trained Glow-TTS and Grad-TTS with CIArTTS corpus and ASC explicitly. The system was evaluated using subjective metrics, Mean Opinion Score, and objective metrics such as MCD, BAP, Lf0 RMSE, PESQ, and speaker similarity. The obtained results indicated a better quality of synthesized speech when using the CIArTTS corpus compared to ASC. In addition to this, we made available the CIArTTS corpus to the public domain for research purposes along with an Arabic TTS demo and Hi-Fi GAN pre-trained neural vocoder model. In the future, we would like to use transfer learning methods to exploit the large-scale CIArTTS corpus for speaker adaptation and voice-cloning in the Arabic language.

8. References

- [1] A. J. Hunt and A. W. Black, "Unit selection in a concatenative speech synthesis system using a large speech database," in *1996 IEEE International Conference on Acoustics, Speech, and Signal Processing Conference Proceedings*, vol. 1. IEEE, 1996, pp. 373–376.
- [2] H. Zen, T. Nose, J. Yamagishi, S. Sako, T. Masuko, A. W. Black, and K. Tokuda, "The hmm-based speech synthesis system (hts) version 2.0," *SSW*, vol. 6, pp. 294–299, 2007.
- [3] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio, Q. V. Le, Y. Agiomyrgiannakis, R. Clark, and R. Saurous, "Tacotron: Towards end-to-end speech synthesis," in *proceedings of INTER-SPEECH*, 2017.
- [4] R. J. Weiss, R. Skerry-Ryan, E. Battenberg, S. Mariooryad, and D. P. Kingma, "Wave-tacotron: Spectrogram-free end-to-end text-to-speech synthesis," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5679–5683.
- [5] N. Halabi, "Modern standard arabic phonetics for speech synthesis," Ph.D. dissertation, UNIVERSITY OF SOUTHAMPTON, 2016.
- [6] A. Amrouche, A. Abed, K. Ferrat, K. N. Boubakeur, Y. Bentrchia, and L. Falek, "Balanced arabic corpus design for speech synthesis," *International Journal of Speech Technology*, vol. 24, no. 3, pp. 747–759, 2021.
- [7] T. Lambert, N. Braunschweiler, and S. Buchholz, "How (not) to select your voice corpus: random selection vs. phonologically balanced," in *SSW*. Citeseer, 2007, pp. 264–269.
- [8] E. Bakhturina, V. Lavrukhin, B. Ginsburg, and Y. Zhang, "Hi-fi multi-speaker english tts dataset," in *Interspeech*, 2021.
- [9] R. Abdelmalek and Z. Mnasri, "High quality arabic text-to-speech synthesis using unit selection," *2016 13th International Multi-Conference on Systems, Signals & Devices (SSD)*, pp. 1–5, 2016.
- [10] A. A. Shalaby, O. A. Dakkak, and N. Ghneim, "An arabic text to speech based on semi-syllable concatenation," *International Review on Computers and Software*, vol. 11, pp. 1178–1186, 2016.
- [11] O. O. Khalifa, M. Z. Obaid, A. W. Naji, and J. I. Daoud, "A rule-based arabic text-to-speech system based on hybrid synthesis technique," 2011.
- [12] N. Halabi, "Modern standard arabic phonetics for speech synthesis," 2016.
- [13] M. Baali, T. Hayashi, H. Mubarak, S. Maiti, S. Watanabe, W. El-Hajj, and A. Ali, "Unsupervised data selection for tts: Using arabic broadcast news as a case study," *ArXiv*, vol. abs/2301.09099, 2023.
- [14] A. Abdelali, N. Durrani, C. Demiroğlu, F. Dalvi, H. Mubarak, and K. Darwish, "Natiq: An end-to-end text-to-speech system for arabic," *ArXiv*, vol. abs/2206.07373, 2022.
- [15] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerry-Ryan, R. A. Saurous, Y. Agiomyrgiannakis, 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)*, 2018, pp. 4779–4783.
- [16] O. Zine and A. Meziane, "Novel approach for quality enhancement of arabic text to speech synthesis," *2017 International Conference on Advanced Technologies for Signal and Image Processing (ATSIP)*, pp. 1–6, 2017.
- [17] A. Amrouche, A. Abed, K. Ferrat, K. N. Boubakeur, Y. Bentrchia, and L. Falek, "Balanced arabic corpus design for speech synthesis," *International Journal of Speech Technology*, vol. 24, pp. 747–759, 2021.
- [18] T. Zerrouki and A. Balla, "Tashkeela: Novel corpus of arabic vocalized texts, data for auto-diacritization systems," *Data in Brief*, vol. 11, pp. 147–151, 2017.
- [19] V. Popov, I. Vovk, V. Gogoryan, T. Sadekova, and M. A. Kudinov, "Grad-TTS: A diffusion probabilistic model for text-to-speech," *proceedings of International Conference on Machine Learning (ICML)*, 2021.
- [20] J. Kim, S. Kim, J. Kong, and S. Yoon, "Glow-TTS: A generative Flow for text-to-speech via monotonic alignment search," in *proceedings of Conference on Neural Information Processing Systems (NIPS)*, 2020.
- [21] J. Kong, J. Kim, and J. Bae, "HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis," in *proceedings of Conference on Neural Information Processing Systems (NIPS)*, 2020.
- [22] M. D. Polkosky and J. R. Lewis, "Expanding the MOS: development and psychometric evaluation of the MOS-R and MOS-X," *International Journal of Speech Technology*, vol. 6, pp. 161–182, 2003.
- [23] A. W. Rix, J. G. Beerends, M. Hollier, and A. P. Hekstra, "Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs," *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No.01CH37221)*, vol. 2, pp. 749–752 vol.2, 2001.
- [24] A. Kulkarni, V. Colotte, and D. Jouvet, "Analysis of expressivity transfer in non-autoregressive end-to-end multispeaker tts systems," in *Interspeech*, 2022.
- [25] B. Desplanques, J. Thienpondt, and K. Demuynck, "Ecapa-tdnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification," in *Interspeech*, 2020.