

Enhancing Low Resource Language Text-to-Speech with Native Corpus Training: A High-Quality WaveGlow Vocoder Approach

Abdulahdy Abas Abdullah
Artificial Intelligence and Innovation
Centre, University of Kurdistan Hewlêr
Erbil, Kurdistan Region, Iraq

Sabat Salih Muhamad
Computer Science Department, Faculty
of Science, Soran University
Soran, Kurdistan, Iraq

Hadi Veisi
School of Intelligent Systems
College of Interdisciplinary Science and
Technologies University of Tehran, Iran

Abstract:

The ability to synthesize spoken language from text has greatly facilitated access to digital content with advances in text-to-speech technology. However, effective TTS development for low-resource languages, such as Central Kurdish (CKB), still faces many challenges due mainly to the lack of linguistic information and dedicated resources. In this paper, we improve the Kurdish TTS system based on Tacotron by training the Kurdish WaveGlow vocoder on a 21-hour central Kurdish speech corpus instead of using a pre-trained English vocoder WaveGlow. Vocoder training on the target language corpus is required to accurately and fluently adapt phonetic and prosodic changes in the Kurdish language. The effectiveness of these enhancements is that our model is significantly better than the baseline system with English pre-trained models. Our adaptive WaveGlow model achieves an impressive MOS of 4.91, which sets a new benchmark for Kurdish speech synthesis. On one hand, this study empowers the advanced features of the TTS system for Central Kurdish, and on the other hand, it opens the doors for other dialects in Kurdish and other related languages to further develop.

Keywords: Kurdish Language, Speech Synthesis, WaveGlow vocoder, Deep Learning.

1. Introduction

Text-to-speech has developed and evolved in the recent past, such that written text is converted to spoken language with most features of the intonation and cadence of a natural speaker. [1] Therefore, it is an important application in many fields, such as assistive technology for people with reading difficulties, interactive chatbots, advanced conversational AI, language learning platforms, or entertainment. At the heart of TTS research is speech synthesis, which is the ability to generate audio from text in a way that makes it natural speech of human-level quality. Recent deep learning breakthroughs led to new models covering end-to-end speech synthesis.

Speech synthesis technologies such as Char2Wav by [2], Tacotron by [3], Tacotron2 by [4], DeepVoice by [5], Transformer TTS by [6], FastSpeech by [7], ParaNet by [8], Neural Codec Language Models by [9], Matcha-TTS [10].

AI-powered TTS engines are advancements that have uplifted synthesized speech quality beyond traditional statistical parametric methods. These recent works utilize state-of-the-art neural vocoders, a source-filter vocoder like WORLD [11], which can synthesize audio quality much better and sound very similar to natural human speech.

However, despite these major advances in TTS, it is still a challenge for some languages to build high-quality and human-level systems, especially for low-resource languages, one of which is Kurdish. Kurdish, spoken by over 30 million people, is one such language. It is divided into three main dialects: Central Kurdish, Northern Kurdish “Kurmanji”, and Southern Kurdish. In the last few years, several works have been done for the Kurdish language. In this direction, [12] have worked on overcoming the problems of TTS in Kurdish by developing a Central Kurdish corpus and using Tacotron 2 with the HiFi-GAN vocoder for synthesizing high-quality human-like speech. Using English character embeddings and 10 hours of recorded audio, the system was rated as having a mean opinion score of 4.1, which is like other state-of-the-art synthesizers in other languages. The more advanced end-to-end TTS model for the Kurdish language, in this case, for the Sorani dialect, introduced the use of variational auto-encoders combined with adversarial learning and a stochastic duration predictor. The latter enhances text-to-speech conversion by aligning the latent distributions and including diverse rhythms, hence yielding a high MOS of 3.94, which has outperformed all existing one-stage and two-stage systems with the measures used in naturalness and audio quality for RNN-TTS systems [13].

Another research paper proposes an end-to-end TTS system for Central Kurdish (CKB), hence handling data sparsity. Three different training experiments were performed for Tacotron2 on a dataset of 21 hours of female voice. The one trained from scratch on the full dataset scored the highest in terms of naturalness and intelligibility according to the Mean Opinion Score [5] of 4.78 out of 5 compared to other models [14]. In addition, the development of the Kurdish TTS for the Central Kurdish dialect still needs to be improved. It is characterized as a low-resource language with different phonetic and linguistic characteristics, and Vocoder's general models do not adequately cover the characteristics of Kurdish speech due to the phonetic differences between the source language and the target language. The lack of a vocoder trained specifically for the Kurdish language is one of the obstacles so far to the Kurdish TTS model. In general, the relevance of designing dedicated vocoders that better represent Kurdish speech lies within the need for targeted development for improved quality in TTS of low-resource languages.

As a result, this paper makes several groundbreaking contributions to Kurdish speech synthesis. Firstly, it introduces the first TTS vocoder based on 21 hours of detailed speech data, which marks a significant advancement in Kurdish language technology. Secondly, we have successfully adapted the WaveGlow deep-learning architecture [15] to Kurdish, optimizing it for the unique acoustic properties of the language to ensure clear, natural speech output. Additionally, we have implemented advanced prosody modeling techniques to improve the rhythm, stress, and intonation of the synthesized speech, crucial for achieving lifelike speech quality. These enhancements not only push the boundaries of TTS for Kurdish but also offer scalable methodologies that can be applied to other Kurdish dialect languages, broadening the impact of this work across different linguistic communities.

In the remainder of this paper, we elaborate on these themes in several sections: Section 2 discusses some recent advancements in TTS technology, focusing on Normalizing Flows and WaveGlow technology applied for developing low-resource languages such as Kurdish. Section 3 presents the specific architecture of the adapted WaveGlow model for the Kurdish TTS. Section 4 presents the experimental results, and Section 5 concludes with a summary of our findings and some future lines that this area of research might take.

2. Related Works

With advancements in neural network-based text-to-speech (TTS) systems, generating high-quality, natural-sounding speech has become increasingly feasible. However, because of a lack of data and resources, expanding TTS systems for low-resource languages like Kurdish continues to be difficult. This section examines relevant TTS research, emphasizing methods that make use of vocoder and its variants. We examine several approaches and works in the field, highlighting their applicability and possible modifications for Kurdish text-to-speech. Furthermore, we pinpoint the shortcomings and difficulties in the current literature, which directs our strategy for enhancing voice synthesis vocoders for low-resource languages.

WaveGlow introduced [15] a flow-based generative network for voice synthesis that integrates knowledge from WaveNet and Glow. With this approach, the Mel spectrograms are transformed into high-quality speech using a single network and cost function. The realizations of WaveGlow are the efficient MOSs, the high sample rates of audio samples, and the equalized and consecutive training model. However, the model is very large and computationally expensive, which can be considered as a major drawback of the presented approach.

[16] proposed a new glottal neural vocoder in 2018. It employs glottal-source filters combined with vocal-tract filters by utilizing the hybrid neural network. While this vocoder delivers a good amount of voice synthesis that is of good quality, it comes alongside implementation issues. Other research that can be mentioned in this context is the work of [16] that explored WaveGlow voice synthesis for ultrasound-based articulatory-to-acoustic mapping. WaveGlow, specifically developed for speakers and silent speech interfaces targeting clients with speaking disorders, outperforms regular vocoders and has more realistic synthesized speech, research showed. A major limitation is the use of ultrasonic technology.

A neural vocoder based on WaveNet was designed and presented by [17]; which is a very efficient neural vocoder attained especially for smooth real-time speech synthesizing. While this vocoder offers a very good quality synthesis it is very demanding in terms of the processing power needed. Extending the source-filter model into a flow-based deep generative model, which was named ExcitGlow is a variant of WaveGlow, and it specifically targets the distribution of the excitation signal rather than the speech signal. This approach contributes to a drastic reduction of the model size from 87.73 million to 15. When choosing between GAN-based and TTS systems, our study determines that the latter can achieve an adequate model size of 60 million parameters while maintaining the perceivable quality of synthesized speech. Nevertheless, the closed-loop training framework that has been proposed by ExcitGlow is not easy to adopt. The study completed by [18] aims to demonstrate the application of RNNs to sequence-to-sequence for voice synthesis by utilizing a continuous vocoder. The findings of the study show that RNN-based models including LSTM, BLSTM, and GRU are superior to feed-forward DNNs in aspects of naiveness and prediction capability and have the potential of producing the highest performance in SPSS. An inherent disadvantage of this method is the high computational complexities of RNN-based models which may cause slow loading. Some studies that have been made include [19] the study used Tacotron2 and WaveGlow to assess the effectiveness of TTS systems involving Sanskrit in low-resource contexts. Their strategy was focused on transfer learning and fine-tuning the available models to achieve excellent voice synthesis with a limited amount of available data. The major limitation of this method is its reliance on pre-trained models.

[19] Ping and colleagues introduced WaveFlow for raw audio in 2020 in a flowing form that is concise and called flowing. This model adopts a dilated 2-D convolutional structure which captures long-range structure as well as autoregressive functions to handle local changes. WaveFlow is significantly smaller and computationally faster than WaveGlow which only requires 5.71 billion parameters and can generate high-fidelity audio at a rate faster than real-time without the presence of a deliberately engineered inference kernel. However, the testing of this method has been limited to specific datasets, thus its application might be limited in other applications. Proposed efficient version of the WaveGlow, called Efficient WaveGlow or EWG for short. This improvement is built on top of the previously mentioned model with the primary intention of reducing a number of parameters and increasing speed while keeping the overall quality of the generated speech at the descent level. In the EWG model, the WaveNet style transform network is replaced with the dilated convolution network utilizing the FFTNet protocol. It also involves group convolution, and it conveys local information between two layers. While it is conceivable that there may be trade-offs in quality for the large parameter reduction made, the improvements made are evident with alleviation in computing and inference time all while retaining high-quality synthesis. The contribution of this work is the development of an enhanced technique for TTS synthesis for LRLs and provides a roadmap for further investigation.

A paper that is like the earlier study was done by [12] in April of 2022 and plans on creating a Central Kurdish text-to-voice (TTS) system. The present system used transfer learning techniques borrowed from an English model that had been pre-trained. In addition, 10 hours of voice data in Central Kurdish were also used in total, and the HiFi-GAN vocoder was used. Many vocoders such as Griffin-Lim, WaveGlow, and HiFi-GAN are explored by [20]. to examine their applied use in TTS systems that were designed for extensively spoken low-resource languages with a focus on the Indian region. WaveGlow has been identified as famous for being capable of producing a vocal output of superior quality, however, the authors explained that there are more enhanced vocoders notably; HiFi-GAN who are becoming popular due to efficiency and higher quality.

[13] came up with a new Terminated Text Synthesis (TTS) in Central Kurdish in 2024 based on end-to-end transformers. The suggested technique enhances the realism and intelligibility of the mapping of text-to-speech by learning to map directly while using data augmentation methods to solve the problem of data shortage. Unquestionably, such an approach was quite useful; nonetheless, the method demanded much processing capacity, and its success depended on the quality of the available corpus. In another previous study [14], we employed transfer learning from an English pre-trained model to train our initial model. However, we utilized a larger dataset (21 h) and substituted the HiFi-GAN vocoder with the WaveGlow vocoder, which was trained on an English pre-trained model. This modification resulted in the generation of speech with superior quality and a more authentic sound. A summary of the related studies is provided, highlighting their important elements in Table I.

TABLE I: SUMMARY OF REVIEW ON VOCODER TECHNIQUE FOR SPEECH SYNTHESIS.

| Reference | Methodology | Advantages | Disadvantages |
|-----------|---|--|---|
| [16] | Glottal source + vocal tract filters, hybrid neural network | High-quality, efficient synthesis | Implementation complexity |
| [15] | Glow + WaveNet, single cost function | High-quality speech, stable training | High computational cost, large model size |
| [18] | Sequence-to-sequence RNNs | Improved naturalness, state-of-the-art performance | Computationally intensive |
| [17] | WaveNet-style transform, optimized vocoder | Real-time, high-quality synthesis | High resource requirements |
| [21] | FFTNet-style, group convolution, shared local conditions | Reduced computation, faster inference, smaller model | Possible quality trade-offs |
| [8] | 2-D convolutions, autoregressive functions | Smaller model, faster synthesis, high fidelity | Limited dataset testing |
| [22] | Source-filter model, flow-based generative model | Reduced model size, high-quality synthesis | Complex training framework |
| [19] | Transfer learning, fine-tuning Tacotron2 + WaveGlow | High-quality synthesis with limited data | Dependency on pre-trained models |
| [12] | Transfer learning, Tacotron2 + HiFi-GAN | High-quality synthesis with limited data | Dependency on pre-trained models |

| | | | |
|------|---|--|--|
| [20] | Tacotron 2 + WaveGlow, Griffin-Lim, and HiFi-GAN vocoders | Applied several vocoders, conducting a comparative analysis. Protocols for collecting scalable data. | Data quality obtained from crowdsourcing can exhibit variability. Restricted discourse on difficulties particular to language. Vocoders of superior quality may experience scaling challenges. |
| [13] | end-to-end Transformer model | High-quality synthesis Better Performance in Low-Resource Settings | High Computational Requirements Dependency on Data Quality |
| [14] | Tacotron2 + WaveGlow | High-quality speech | Dependency on pre-trained models |

In conclusion, for low-resource languages like Kurdish, selecting an appropriate vocoder is crucial to balance quality and efficiency. Based on a review of previous works, WaveGlow stands out as a suitable choice. WaveGlow significantly reduces computational costs and model size while maintaining high-quality speech synthesis, making it well-suited for resource-constrained environments.

3. Methodology

In this section, we discuss the speech corpus for the target language, with the proposed approach for enhancing TTS by training new WaveGlow for the Kurdish language.

1. Kurdish Speech Corpus

Train phase: In this study, we utilized the existing "Sabat Speech Corpus" [14] for training the vocoder WaveGlow, enhancing its capability to synthesize Central Kurdish speech. The corpus, which was previously compiled contains 10,979 utterances across diverse categories such as news, sports, linguistics, psychology, poetry, health, scientific topics, general knowledge, interviews, politics, education, literature, narratives, tourism, and miscellaneous subjects.

Table I presents a comprehensive breakdown of the different types of utterances and their respective counts, demonstrating the wide range of topics covered in our corpus.

TABLE II: THE TOTAL OF THE TRAIN UTTERANCES (Muhamad et al., 2024)

| Category | No. of Utterances |
|---------------------------|-------------------|
| linguistics | 1760 |
| questions and exclamation | 1393 |
| story | 1092 |
| poem | 916 |
| tourism | 782 |
| miscellaneous | 700 |
| sport | 683 |
| education and literature | 619 |
| news | 608 |
| science | 543 |
| health | 483 |
| politics | 483 |
| General information | 461 |

| | |
|------------------|--------|
| interview | 456 |
| Total | 10,979 |

This broad coverage is crucial for training a vocoder like WaveGlow, as it ensures that the system can accurately reproduce a wide range of phonetic and intonational nuances inherent to the Kurdish language. By leveraging this pre-existing corpus, we aim to improve the naturalness and expressiveness of synthesized Kurdish speech, addressing the unique challenges of speech synthesis in low-resource language settings. The summary Overview of the Sabat Speech Corpus is shown in Table I.

TABLE III: OVERVIEW OF THE SABAT SPEECH CORPUS (Muhamad et al., 2024)

| Feature | Details |
|-----------------------|---|
| Total Utterances | 10,979 utterances |
| Audio Length | 21 hours total |
| Sampling Rate | 22,050 Hz |
| Bit Depth | 16 bits |
| Channel | Mono |
| File Format | WAV |
| Categories | News, Sport, Linguistics, Poem, Health, etc. |
| Recording Environment | Professional studio |
| Speaker Profile | Female, from Sulaymaniyah, in her thirties |
| Text Normalization | Applied AsoSoft Library Normalization tool ¹ |

Test Phase: For the test set, 110 sentences were meticulously selected from texts covering 17 distinct subject areas (Muhamad et al., 2024), ensuring they differed from the training sentences to effectively gauge the TTS model's performance across diverse contexts. These sentences were sourced from several websites and refined to align with Central Kurdish orthography standards. This careful curation aims to test the TTS system's ability to handle a variety of linguistic challenges and assess its generalization capabilities across different types of content listed in Table I.

TABLE IV: DISTRIBUTION OF TEST SET SENTENCES ACROSS TOPICS (Muhamad et al., 2024)

| Topics | No. of Sentences |
|---------------------|-------------------------|
| News | 10 |
| Formal Letter | 10 |
| Sport | 9 |
| Poem | 8 |
| Questions | 7 |
| Psychology | 6 |
| Health | 6 |
| Science | 6 |
| Miscellaneous | 6 |
| General Information | 6 |
| Story | 6 |
| Tourism | 6 |

¹ <https://github.com/AsoSoft/AsoSoft-Library>

| | |
|--------------------------|------------|
| Linguistics | 5 |
| Interview | 5 |
| Politics | 5 |
| Education and Literature | 5 |
| Exclamation | 4 |
| Total | 110 |

2. Text to Speech

The general design for this TTS model involves a two-stage process. [23]. The first model converts text data into Mel spectrograms, and the second model takes these Mel spectrograms as input, producing binary sound wave representations. These sound waves, when played, generate human speech. This structure is illustrated in the accompanying diagram and shown in Figure. 1.




Figure. 1. Text to Mel to Waveform

3. Tacotron 2

The methodology presented in our previous work [14], as shown in Figure. 2, entails generating a dataset consisting of "text, audio" pairs from a single female speaker. This dataset is then subjected to preprocessing procedures, including text normalization, to prepare the data for our TTS system. For the prediction of Mel spectrograms, we utilized a sequence-to-sequence synthesis network that is based on Tacotron2. [4]. These spectrograms are then transformed into audible sounds using the WaveGlow vocoder [15]. Tacotron2's streamlined end-to-end design, including its acoustic model and vocoder stages, is well-suited for Kurdish because of its straightforward phonetic structure, which requires less input compared to more intricate languages. The proposed model utilized a recurrent neural network to produce Mel spectrograms from given textual input. Subsequently, the WaveGlow vocoder converts these spectrograms into waveforms. The Tacotron2 model and WaveGlow vocoder are trained separately on our comprehensive Kurdish voice corpus.

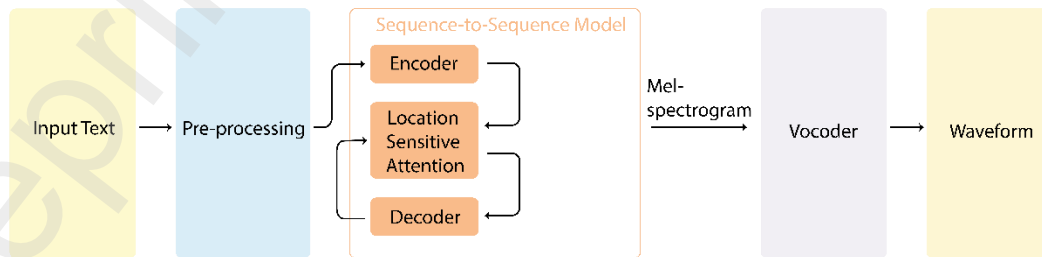


Figure. 2. The block diagram from our previously published methodology

4. Vocoder Architecture

Figure. 3 shows the conversion of the Mel spectrogram into speech waveform with the help of the WaveGlow vocoder. The first input is a Mel spectrogram, which graphically depicts the density of the frequencies in the audio signal concerning the time. As it will be explained later in this paper, the Mel spectrogram retains features of the audio that are critical for signal processing and analysis because they correspond with characteristics of human auditory processing. Subsequently, the Mel spectrogram undergoes the process called “Vocoder WaveGlow” which forms the end part of the proposed model. WaveGlow is a state-of-the-art generative model, aimed at the high-quality audio generation from Mel spectrograms. [22]. A set of invertible transformations, the normalizing flows, is used to transform the data, the Mel spectrogram that has a simpler distribution than the distribution of the audio waveform. In so doing, WaveGlow transforms the Mel spectrogram back into the original signal and all its details into a waveform. The final product from the WaveGlow vocoder is a “Speech Waveform”. This waveform is the reconstructed audio signal that can be played and heard by human intervention as speeches are being made. It has all the differential attributes of natural speech that are synthesized from the features embedded in the Mel spectrogram. The speech waveform is the result that listeners hear, and it is generated to closely mimic natural human speech. It depicts the progression from the Mel spectrogram, which is a text-based or visual representation of sound frequencies, to the speech waveform, which is the audible output. Figure.3 below demonstrates how the WaveGlow model bridges the gap between the abstract representation of audio in the form of a Mel spectrogram and the concrete, listenable speech waveform. The process underscores the capability of the WaveGlow vocoder to synthesize high-quality, natural-sounding speech from a structured, frequency-based input.

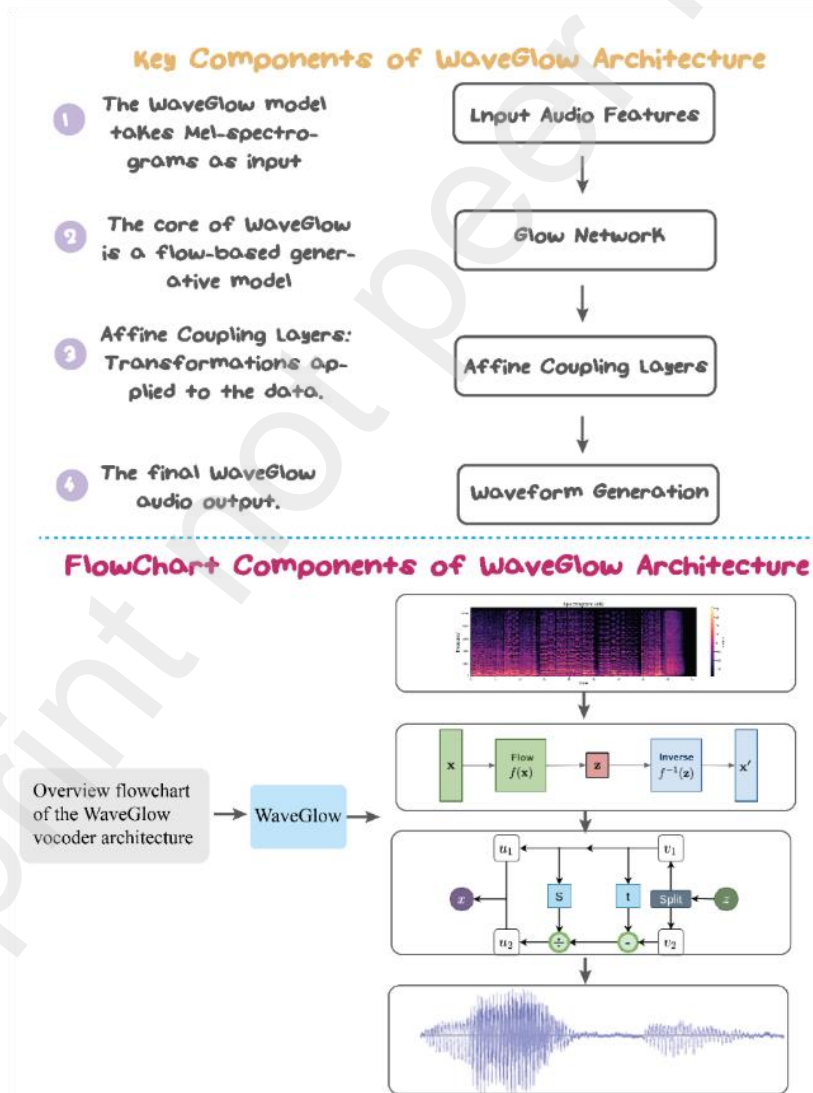


Figure. 3. Overview key components and flowchart vocoder WavGlow Architecture

5. WaveGlow

WaveGlow, inspired by the Glow concept, offers faster synthesis compared to WaveNet. [24]. It employs an invertible transformation between blocks of eight time-domain audio samples and a standard normal distribution, conditioned on the log Mel-spectrogram. [28]. This transformation enables audio generation by sampling from this Gaussian density. The invertible transformation consists of a sequence of individual invertible transformations, known as normalizing flows. In WaveGlow, each flow is composed of a 1x1 convolutional layer followed by an affine coupling layer. The affine coupling layer is a neural transformation that predicts a scale and bias based on the input speech x and Mel-spectrogram. X . Let W_k represent the learned weight matrix for the k -th 1x1 convolutional layer and $s_j(x, X)$ be the predicted scale value at the j -th affine coupling layer. For inference, WaveGlow samples z from a uniform Gaussian distribution and applies the inverse transformations (f^{-1}) conditioned on the Mel spectrogram X to retrieve the speech sample x . Due to the simplicity of parallel sampling from the Gaussian distribution, all audio samples can be generated simultaneously. The model is trained to minimize the log-likelihood of the clean speech samples x :

$$\ln P(x | X) = \ln P(z) - \sum_{j=0}^J \ln s_j(x, X) - \sum_{k=0}^K \ln |W_k| \quad Eq(1)$$

where J is the number of coupling transformations and K is the number of convolution layers. In $P(z)$, the parameter σ^2 refers to the variance of the Gaussian distribution, and $\sigma = 1$ is used during training. WaveGlow uses an alternative to the logistic function in equation (1) for sampling. This system was implemented using 12 coupling layers, each with 8 layers of dilated convolution with 512 residual channels. The same architecture was adapted for the PR system using WaveGlow as its vocoder, referred to as PR-WaveGlow. The architecture of the WaveGlow model is shown in Figure. 4.

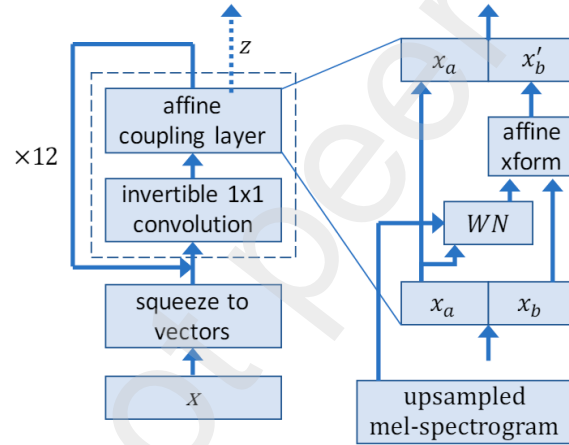


Figure. 4. The architecture of the WaveGlow model [15].

4. Results and Discussion

1. TTS measure

MOS is a commonly used measure in evaluating the quality of synthesized speech. It involves human listeners rating the quality of speech samples on a predefined scale, typically from 1 to 5, where:

- 1: Bad
- 2: Poor
- 3: Fair
- 4: Good
- 5: Excellent

The MOS is calculated by averaging the scores given by all listeners for each speech sample. Equation (2) is utilized to ascertain the assessment of intelligibility and naturalness for MOS.

$$MOS = \frac{1}{N} \sum_{i=0}^N s_i \quad Eq(2)$$

Where:

N is the number of listeners.

s_i is the score given by the i – th listener.

For our evaluation, native Kurdish speakers rated the naturalness and intelligibility of the synthesized speech samples.

2. Implementation Tools and Hyper-Parameters

The training and implementation of the Kurdish WaveGlow model were conducted using two NVIDIA RTX 4090 GPUs, providing a total of 48 GB of GPU memory. The system also utilized a total of 290 GB of RAM. The training parameters used for developing the Kurdish WaveGlow model are summarized in Table V.

TABLE V

TRAINING PARAMETERS

| Parameter | Value |
|---------------------------|---|
| Batch Size | 22 |
| Learning Rate | 1e-4 with an exponential decay |
| Optimizer | Adam optimizer (beta1=0.9, beta2=0.999) |
| Number of Epochs | 100000 |
| Sigma | 1.0 |
| Iterations per Checkpoint | 2000 |
| Seed | 1234 |

The parameters in Table V describe the settings used to train the Kurdish WaveGlow model. The batch size indicates the number of samples processed before the model is updated. The learning rate defines the step size for weight updates. The optimizer used is Adam, known for its efficiency and adaptive learning rate capabilities. Training was conducted for 100,000 epochs to ensure comprehensive learning, with intermediate checkpoints saved every 2000 iterations. A fixed seed ensures the reproducibility of the results.

TABLE VI: DATA CONFIGURATION PARAMETERS

| Parameter | Value |
|-------------------|----------|
| Segment Length | 16000 |
| Sampling Rate | 22050 Hz |
| Filter Length | 1024 |
| Hop Length | 256 |
| Win Length | 1024 |
| Mel Frequency Min | 0.0 |
| Mel Frequency Max | 8000.0 |

Table VI details the data configuration parameters used during the preprocessing stage. These settings ensure the audio data is correctly formatted and segmented for training. The segment length specifies the number of audio samples per segment, while the sampling rate defines the number of samples per second. Filter length, hop length, and win length are used in the Short-Time Fourier Transform (STFT) process to convert audio signals into Mel-spectrograms. Mel frequency parameters set the range of frequencies to be included in the Mel-spectrogram.

TABLE VII
WAVEGLOW CONFIGURATION PARAMETERS

| Parameter | Value |
|------------------------|-------|
| Number of Mel Channels | 80 |
| Number of Flows | 12 |

| | |
|---------------|-----|
| Group Size | 8 |
| Early Every | 4 |
| Early Size | 2 |
| - Layers | 8 |
| - Channels | 256 |
| - Kernel Size | 3 |

Table VII provides the configuration parameters specific to the WaveGlow model. The number of Mel channels indicates the dimensionality of the Mel-spectrogram features. The number of flows specifies the layers of transformations used to map the Mel-spectrogram to audio. Group size determines the number of samples processed together in the model, while early every and early size parameters control the early output of samples to improve efficiency. The WaveNet configuration includes the number of layers, channels, and kernel size used in the WaveGlow network to generate high-fidelity audio.

1. Experimental Results

The Kurdish WaveGlow model was trained from scratch using the Sabat Speech Corpus, which consists of 21 hours of high-quality, annotated speech data. The training process did not involve any pre-trained models or fine-tuning; instead, the model was developed entirely from the ground up to ensure it could fully adapt to the unique characteristics of the Kurdish language. The training process spanned a total of 5 days, with each day comprising 24 hours of continuous training, amounting to 120 hours in total. The model was configured to save checkpoints every 2000 iterations, resulting in a total of 2,702,000 checkpoints throughout the training period. This extensive training and checkpointing allowed for meticulous tracking of the model's performance and adjustments as needed. Throughout the training period, the model demonstrated steady convergence, with a consistent decrease in training loss. This indicates that the model effectively learned the acoustic properties and linguistic nuances of Kurdish speech from the Sabat Corpus.

3. Evaluations

For the MOS evaluation, 110 random sentences from various categories were selected, ensuring these sentences were not part of the training set. The categories included news, sports, linguistics, psychology, poetry, health, questions, exclamations, science, miscellaneous, general information, interviews, politics, education and literature, stories, tourism, and SMS. This diverse selection ensured a comprehensive assessment of the model's performance across different types of content.

Raters: Twelve native Kurdish speakers (7 males and 5 females, aged 21 to 46) participated in the MOS evaluation. They listened to the synthesized sentences using headphones to ensure consistent audio quality and rated each sentence on the 5-point Likert scale. The results of the MOS evaluation for each category for the Kurdish Tacotron2-Scratch (Muhamad et al., 2024) WaveGlow Kurdish-Scratch model are shown in Figure. 5.




Figure. 5. MOS Results for Kurdish Tacotron2-Scratch (Muhamad et al., 2024) WaveGlow Kurdish-Scratch

These MOS results highlight the high performance of the Kurdish Tacotron2-Scratch (Muhamad et al., 2024) WaveGlow Kurdish-Scratch model across various content categories, reflecting its ability to generate natural and intelligible Kurdish speech.

4. Comparative Analysis of Kurdish TTS: WaveGlow (Kurdish) vs. Vcoders (English Pre-trained)

This time the comparison was made between the TTS models trained from scratch on the Kurdish-specific corpus and the models incorporated with the English pre-trained WaveGlow. The comparison is made in view of the average MOS of all the models and of each genre and sub-genre of content. This analysis gives performance information of the synthesized speech and the extent of naturalness of the synthesized speech by the different training approaches. Table VIII below also demonstrates the MOS for different content classifications. These results highlight the performance differences between genuine voice samples and the three models: Kurdish Tacotron2 pre-trained with HiFi-GAN English pre-trained and Kurdish Tacotron2 scratch with WaveGlow English pre-trained and the proposed model Kurdish Tacotron2 scratch with WaveGlow vocoder that train for Kurdish dataset only.

TABLE VIII
MOS RESULTS FOR DIFFERENT KURDISH MODELS BENCHMARK BY CATEGORY

| Category | MOS (Genuine Voice) | Kurdish Tacotron2-Pretrain with HiFi-GAN English Pre-trained [12] | Kurdish Tacotron2-Scratch with WaveGlow English Pre-trained (Muhamad et al., 2024) | Our model (Kurdish Tacotron2-Scratch with WaveGlow Kurdish-Scratch) |
|-------------|---------------------|---|--|---|
| News | 5.0 | 4.2 | 4.5 | 4.92 |
| Sports | 4.9 | 4.1 | 4.4 | 4.75 |
| Linguistics | 4.8 | 4.0 | 4.3 | 4.93 |
| Psychology | 5.0 | 4.3 | 4.6 | 4.97 |
| Poem | 4.9 | 4.2 | 4.5 | 4.83 |
| Health | 4.8 | 4.1 | 4.4 | 4.99 |

| | | | | |
|-----------------|-----|-----|-----|------|
| Questions | 5.0 | 4.2 | 4.5 | 4.96 |
| Exclamation | 5.0 | 4.2 | 4.5 | 4.96 |
| Science | 4.9 | 4.1 | 4.4 | 5.00 |
| Miscellaneous | 4.8 | 4.0 | 4.3 | 4.94 |
| General Info | 4.9 | 4.2 | 4.5 | 4.92 |
| Interviews | 4.9 | 4.2 | 4.5 | 4.88 |
| Politics | 5.0 | 4.3 | 4.6 | 4.78 |
| Education & Lit | 4.8 | 4.0 | 4.3 | 4.93 |
| Story | 4.9 | 4.2 | 4.5 | 4.75 |
| Tourism | 4.9 | 4.2 | 4.5 | 5.00 |
| SMS | 4.8 | 4.0 | 4.3 | 4.91 |

These results demonstrate that the Tacotron2-Scratch (WaveGlow Kurdish-Scratch) model consistently outperformed the other models across all categories. The genuine voice samples, as expected, received the highest MOS scores. However, the Kurdish-Scratch model closely followed, indicating a high level of naturalness and intelligibility in the synthesized speech. To provide a comprehensive comparison, the overall MOS results for each model were averaged across all categories. Figure. 6 analysis highlights the general performance and quality of the synthesized speech for each model.

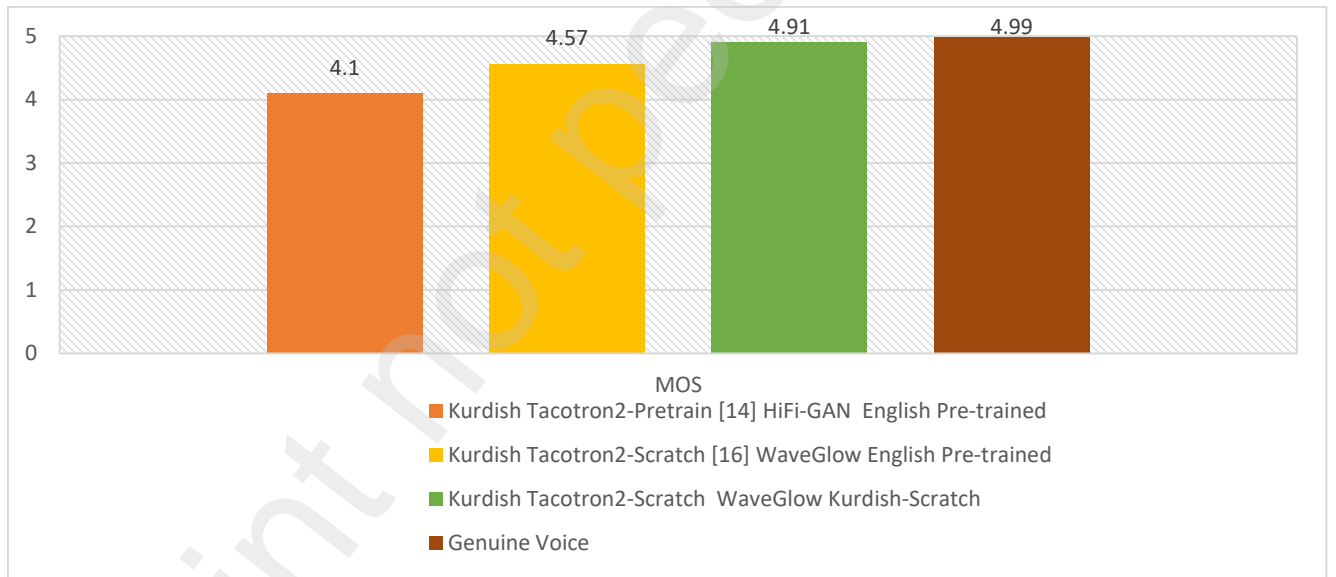


Figure. 6. Comparative Analysis of MOS Results for Kurdish Models Benchmark

The MOS results clearly indicate that the Tacotron2-Scratch (WaveGlow Kurdish-Scratch) model achieved significantly higher scores compared to the models based on English pre-trained WaveGlow. The Kurdish-Scratch model's MOS of 4.78 highlights its ability to generate highly natural and intelligible Kurdish speech, closely approximating genuine voice samples. The lower MOS scores for the Tacotron2-Pre-train (HiFi-GAN) English Pre-trained [12] model reflect the limitations of adapting a pre-trained model to a low-resource language like Kurdish. While it still performed reasonably well, it struggled with pronunciation and naturalness compared to the Kurdish-specific model trained from scratch. Overall, these results emphasize the importance of training TTS models from scratch using language-specific data to achieve the highest quality of synthesized speech. The Kurdish-Scratch model's superior performance demonstrates the effectiveness of this approach for low-resource languages.

5. Summary and Future Works

This work has made significant improvements for TTS technology for Kurdish language especially for Kurdish WaveGlow vocoder to synthesis Kurdish speech. The outcome of our study is the integration of an individual TTS system with the help of 21hrs of clean labelled speech from the Sabat Speech Corpus. This system is therefore devised to accommodate the architectural and development feature of the Kurdish language which is both a popular teaching language and a restricted resource language. These improvements were made through the application of more advanced techniques of prosody modeling to the WaveGlow model: experimental modifications to synthesize Kurdish. All these improvements are thus necessary to achieve an accurate and realistic quality. Besides enhancing the intelligibility and the naturalness of the voice output, the obtained TTS system generates a versatile paradigm that can be extended to other LR languages and, consequently, may revolutionize TTS solutions in overlooked linguistic environments. They also statistically substantiate that our Kurdish-specific proposed model has relatively larger [5] in various contents compared to the models with employ the English pre-trained systems identified in the experiments. This underlines the inherent problem of the lack of language-specific TTS models as it is the only way to ensure the correctness and accurate translation of generated speech. The successful implementation and evaluation of this Kurdish TTS system set the benchmark for future works inside the field of speech synthesis and particularly for languages with limited resources. This precedes possibilities to create speech-based systems that are more accessible to those with disabilities, as well as cater for the fact that people across the globe use different languages. Thus, in the future, the focus will be on improving these approaches, expanding the database including various linguistic data and increasing the efficiency of the system to relate to the complex linguistic and acoustic challenges of the Kurdish language and other low-resource languages. Besides extending the progress of technology in the enhancement of Kurdish digital resources our project also contributes to the discourse on language conservation and digital opportunity.

6. Acknowledgment

The authors would like to express their heartfelt gratitude to the Artificial Intelligence and Innovation Centre (AIIC) at the University of Kurdistan Hewlêr [25] for their invaluable support and resources that made this research possible. The center's state-of-the-art facilities, access to high-performance computing resources, and collaborative environment significantly contributed to the successful development and implementation of this project. Special thanks are extended to the leadership and staff of AIIC for their encouragement and guidance throughout the research process.

7. References

- [1] X. Tan *et al.*, "Naturalspeech: End-to-end text-to-speech synthesis with human-level quality," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [2] J. Sotelo *et al.*, "Char2wav: End-to-end speech synthesis," 2017.
- [3] Y. Wang *et al.*, "Tacotron: Towards end-to-end speech synthesis," *arXiv preprint arXiv:1703.10135*, 2017.
- [4] J. Shen *et al.*, "Natural tts synthesis by conditioning wavenet on mel spectrogram predictions," in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 2018, pp. 4779-4783: IEEE.
- [5] S. Ö. Arik *et al.*, "Deep voice: Real-time neural text-to-speech," in *International conference on machine learning*, 2017, pp. 195-204: PMLR.
- [6] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu, "Neural speech synthesis with transformer network," in *Proceedings of the AAAI conference on artificial intelligence*, 2019, vol. 33, no. 01, pp. 6706-6713.
- [7] Y. Ren *et al.*, "Fastspeech: Fast, robust and controllable text to speech," *Advances in neural information processing systems*, vol. 32, 2019.
- [8] K. Peng, W. Ping, Z. Song, and K. Zhao, "Non-autoregressive neural text-to-speech," in *International conference on machine learning*, 2020, pp. 7586-7598: PMLR.
- [9] X. Wang *et al.*, "Speechx: Neural codec language model as a versatile speech transformer," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024.
- [10] S. Mehta, R. Tu, J. Beskow, É. Székely, and G. E. Henter, "Matcha-TTS: A fast TTS architecture with conditional flow matching," in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 11341-11345: IEEE.
- [11] F. V. Arthur and T. G. Csapó, "Speech synthesis from intracranial stereotactic Electroencephalography using a neural vocoder," *INFOCOMMUNICATIONS JOURNAL: A PUBLICATION OF THE SCIENTIFIC ASSOCIATION FOR INFOCOMMUNICATIONS (HTE)*, vol. 16, no. 1, pp. 47-55, 2024.
- [12] S. Muhamad and H. Veisi, "End-to-End Kurdish Speech Synthesis Based on Transfer Learning," *Passer Journal of Basic and Applied Sciences*, vol. 4, no. 2, pp. 150-160, 2022.

- [13] H. A. Ahmad and T. A. Rashid, "Central Kurdish Text-to-Speech Synthesis with Novel End-to-End Transformer Training," *Algorithms*, vol. 17, no. 7, p. 292, 2024.
- [14] S. S. Muhamad, H. Veisi, A. Mahmudi, A. A. Abdullah, and F. Rahimi, "Kurdish end-to-end speech synthesis using deep neural networks," *Natural Language Processing Journal*, vol. 8, p. 100096, 2024.
- [15] R. Prenger, R. Valle, and B. Catanzaro, "Waveglow: A flow-based generative network for speech synthesis," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 3617-3621: IEEE.
- [16] Y. Cui, X. Wang, L. He, and F. K. Soong, "A New Glottal Neural Vocoder for Speech Synthesis," in *Interspeech*, 2018, pp. 2017-2021.
- [17] K. Oura, K. Nakamura, K. Hashimoto, Y. Nankaku, and K. Tokuda, "Deep neural network based real-time speech vocoder with periodic and aperiodic inputs," *Proc. SSW10*, vol. 32, pp. 13-18, 2019.
- [18] M. S. Al-Radhi, T. G. Csapó, and G. Németh, "RNN-based speech synthesis using a continuous sinusoidal model," in *2019 International Joint Conference on Neural Networks (IJCNN)*, 2019, pp. 1-8: IEEE.
- [19] A. Debnath, S. S. Patil, G. Nadiger, and R. A. Ganesan, "Low-resource end-to-end sanskrit tts using tacotron2, waveglow and transfer learning," in *2020 IEEE 17th India Council International Conference (INDICON)*, 2020, pp. 1-5: IEEE.
- [20] G. K. Kumar, S. Praveen, P. Kumar, M. M. Khapra, and K. Nandakumar, "Towards building text-to-speech systems for the next billion users," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1-5: IEEE.
- [21] W. Song, G. Xu, Z. Zhang, C. Zhang, X. He, and B. Zhou, "Efficient WaveGlow: An Improved WaveGlow Vocoder with Enhanced Speed," in *INTERSPEECH*, 2020, pp. 225-229.
- [22] S. Oh, H. Lim, K. Byun, M.-J. Hwang, E. Song, and H.-G. Kang, "ExcitGlow: Improving a WaveGlow-based neural vocoder with linear prediction analysis," in *2020 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, 2020, pp. 831-836: IEEE.
- [23] Y. A. Li, C. Han, V. Raghavan, G. Mischler, and N. Mesgarani, "Styletts 2: Towards human-level text-to-speech through style diffusion and adversarial training with large speech language models," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [24] A. Van Den Oord *et al.*, "Wavenet: A generative model for raw audio," *arXiv preprint arXiv:1609.03499*, vol. 12, 2016.
- [25] G. Deshmukh, H. S. Mukerjee, and U. D. Prasad, "Risk management in global CRM IT projects," *Business Perspectives and Research*, vol. 8, no. 2, pp. 156-172, 2020.