



TTS-Portuguese Corpus: a corpus for speech synthesis in Brazilian Portuguese

Edresson Casanova¹ · Arnaldo Candido Junior² ·
Christopher Shulby³ · Frederico Santos de Oliveira⁴ ·
João Paulo Teixeira⁵ · Moacir Antonelli Ponti¹ ·
Sandra Aluísio¹

Accepted: 23 November 2021

© The Author(s), under exclusive licence to Springer Nature B.V. 2021

Abstract Speech provides a natural way for human–computer interaction. In particular, speech synthesis systems are popular in different applications, such as personal assistants, GPS applications, screen readers and accessibility tools. However, not all languages are on the same level when in terms of resources and systems for speech synthesis. This work consists of creating publicly available resources for Brazilian Portuguese in the form of a novel dataset along with deep learning models

✉ Edresson Casanova
edresson@usp.br

Arnaldo Candido Junior
arnaldoc@utfpr.edu.br

Christopher Shulby
christopher@definedcrowd.com

Frederico Santos de Oliveira
fredericosantos@inf.ufg.br

João Paulo Teixeira
joaopt@ipb.pt

Moacir Antonelli Ponti
moacir@icmc.usp.br

Sandra Aluísio
sandra@icmc.usp.br

¹ Instituto de Ciências Matemáticas e de Computação, University of São Paulo, São Carlos, Brazil

² Federal University of Technology – Paraná (UTFPR), Medianeira, Brazil

³ DefinedCrowd Corp., Seattle, USA

⁴ Federal University of Mato Grosso, Cuiabá, Brazil

⁵ Research Center in Digitalization and Intelligent Robotics (CEDRI) - Instituto Politecnico de Braganca, Bragança, Portugal

for end-to-end speech synthesis. Such dataset has 10.5 h from a single speaker, from which a Tacotron 2 model with the RTISI-LA vocoder presented the best performance, achieving a 4.03 MOS value. The obtained results are comparable to related works covering English language and the state-of-the-art in European Portuguese.

Keywords Corpora · Speech synthesis · TTS · Portuguese

1 Introduction

Speech synthesis systems have received a lot of attention in recent years due to the great advance provided by the use of deep learning, which allowed the popularization of virtual assistants, such as Amazon Alexa (Purinton et al., 2017), Google Home (Dempsey, 2017) and Apple Siri (Gruber, 2009).

According to Tachibana et al. (2017), traditional speech synthesis systems are not easy to develop, because these are typically composed of many specific modules, such as, a text analyzer, a grapheme-to-phoneme converter, a duration estimator and an acoustic model. In summary, given an input text, the text analyzer module converts dates, currency symbols, abbreviations, acronyms and numbers into their standard formats to be pronounced or read by the system, i.e. carries out the text normalization and tackles problems like homographs, then with the normalized text the phonetic analyzer converts the grapheme into phonemes. In turn, the duration estimator estimates the duration of each phoneme. Finally, the acoustic model receives the phoneme representation sequence, the prosodic information about phoneme segments' length, the F0 contour and computes the speech signal (Teixeira et al., 2003; Ze et al., 2013). Several acoustic models have been proposed, such as the classical formant model (Klatt, 1980), Linear Prediction Coefficients (LPC) model, the Pitch Synchronous Overlap and Add (PSOLA) models (Charpentier & Stella, 1986) widely used in TTS engines like Microsoft Speech API. In addition, Hidden Markov Model (HMM) based synthesis is still a topic of research (Aroon & Dhonde, 2015; Braude et al., 2013; Tokuda et al., 2000), as well as a variety of Unit Selection Models (Siddhi et al., 2017; Wang & Georgila, 2011).

Deep learning (Goodfellow et al., 2016) allows to integrate all processing steps into a single model and connects them directly from the input text to the synthesized audio output, which is referred to as end-to-end learning. While neural models are sometimes criticized as difficult to interpret, several end-to-end trained speech synthesis systems (Kim et al., 2020; Sotelo et al., 2017; Ping et al., 2017; Shen et al., 2018; Tachibana et al., 2017; Valle et al., 2020; Wang et al., 2017) were shown to be able to estimate spectrograms from text inputs with promising performances. Due to the sequential characteristic of text and audio data, recurrent units were the standard building blocks for speech synthesis, such as in Tacotron 1 and 2 (Shen et al., 2018; Wang et al., 2017). In addition, convolutional layers showed good performance while reducing computational costs as implemented in DeepVoice3 and Deep Convolutional Text To Speech (DCTTS) methods (Ping et al., 2017; Tachibana et al., 2017).

Models based on deep learning require a greater amount of data for training, therefore, languages with low available resources are impaired. For this reason, most current TTS models are designed for the English language (Kim et al., 2020; Ping et al., 2017; Shen et al., 2018; Valle et al., 2020), which is a language with many open resources. In this work we propose to solve this problem for the Brazilian Portuguese language. Although there are some public datasets of speech synthesis for European Portuguese (Teixeira et al., 2001), due to the small amount of speech, approximately 100 min, makes the training of models based on deep learning unfeasible. In addition, simultaneously with this work, two datasets for automatic speech recognition for Portuguese, with good quality, were released. The CETUC dataset (Alencar & Alcaim, 2008), which was made publicly available by Quintanilha et al. (2020), has approximately 145 h of 100 speakers. In this dataset, each speaker uttered a thousand phonetically balanced sentences extracted from journalistic texts; on average each speaker spoke 1.45 h. The Multilingual LibriSpeech (MLS) dataset (Pratap et al., 2020) is derived from LibriVox audiobooks and consists of speech in eight languages including Portuguese. For Portuguese, the authors provided approximately 130 h of 54 speakers, an average of 2.40 h of speech per speaker. Although the quality of both datasets is good, both were made available with a sampling rate of 16 kHz and have no punctuation in their texts, making it difficult to apply them for speech synthesis. In addition, the amount of speech per speaker in both datasets is low, thus making it difficult to obtain a single-speaker dataset with a large vocabulary for single-speaker speech synthesis. For example, the LJ Speech dataset (Ito, 2017), which is derived from audiobooks and is one of the most popular open datasets for single-speaker speech synthesis in English, has approximately 24 h of speech.

In this article, we compare models of TTS available in the literature for a language with low available resources for speech synthesis. The experiments were carried out in Brazilian Portuguese and based on a single-speaker TTS. For this, we created a new public dataset, including 10.5 h of speech. Our contributions are twofold: (i) a new publicly available dataset with more than 10 h of speech recorded by a native speaker of Brazilian Portuguese; (ii) an experimental analysis comparing two publicly available TTS models in Brazilian Portuguese language. In addition, our results and discussions shed light on the matter of training end-to-end methods for a non-English language, in particular Portuguese, and the first public dataset and trained model for this language are made available.

This work is organized as follows. Section 2 presents related work on speech synthesis. Section 3 describes our novel audio dataset. Section 4 details the models and experiments performed. Section 5 compares and discusses the results. Finally, Section 6 presents conclusions of this work and future work.

2 Speech synthesis approaches

With the advent of deep learning, speech synthesis systems have evolved greatly, and are still being intensively studied. Models based on Recurrent Neural Networks, such as Tacotron (Wang et al., 2017), Tacotron 2 (Shen et al., 2018), Deep Voice 1

(Arik et al., 2017) and Deep Voice 2 (Arik et al., 2017), have gained prominence, but as these models use recurrent layers they have high computational costs. This has led to the development of fully convolutional models, such as DCTTS (Tachibana et al., 2017) and Deep Voice 3 (Ping et al., 2017), which sought to reduce computational cost while maintaining good synthesis quality.

Ping et al. (2017) proposed a fully convolutional model for speech synthesis and compared three different vocoders: Griffin–Lim (Griffin & Lim, 1984), WORLD Vocoder (Morise et al., 2016) and WaveNet (Van Den Oord et al., 2016). Their results indicated that WaveNet neural vocoder produced a more natural waveform synthesis. However, WORLD was recommended due to its better runtime, even though WaveNet had better quality. The authors further compared the proposed model (Deep Voice 3) with the Tacotron (Wang et al., 2017) and Deep Voice 2 (Arik et al., 2017) models.

Tachibana et al. (2017) proposed the DCTTS model, a fully convolutional model, consisting of two neural networks. The first, called Text2Mel (text to Mel spectrogram), which aims to generate a Mel spectrogram from an input text and the second, Spectrogram Super-resolution Network (SSRN), which converts a Mel spectrogram to the STFT (Short-time Fourier Transform) spectrogram (Benesty et al., 2011). DCTTS consists of only convolutional layers and uses dilated convolution (Kalchbrenner et al., 2016; Yu & Koltun, 2015) to take long, contextual information into account. DCTTS uses the vocoder RTISI-LA (Real-Time Iterative Spectrogram Inversion with Look-Ahead) (Zhu et al., 2007), which is an adaptation of the Griffin–Lim vocoder (Griffin & Lim, 1984), which aims to increase the speed of the synthesis by slightly sacrificing the quality of the audio generated.

Tacotron 1 (Wang et al., 2017) proposes the use of a single trained end-to-end Deep neural network. The model includes an encoder and a decoder. It uses an attention mechanism (Bahdanau et al., 2014) and also includes a post-processing module. This model uses convolutional filters, skip connections (Srivastava et al., 2015), and Gated Recurrent Units (GRUs) (Chung et al., 2014) neurons. Tacotron also uses Griffin–Lim (1984) algorithm to convert the STFT spectrogram to the wave form.

Tacotron 2 (Shen et al., 2018) combines Tacotron 1 with a modified WaveNet vocoder (Tamamori et al., 2017). Tacotron 2 is composed of a recurrent network of prediction resources from sequence to sequence that maps the incorporation of characters in Mel spectrograms, followed by a modified WaveNet model acting as a vocoder to synthesize waveforms in the time domain from those spectrograms. They also demonstrated that the use of Mel spectrograms as the conditioning input for WaveNet, instead of linguistic characteristics, allows a significant reduction in the size of the WaveNet architecture.

3 TTS-Portuguese Corpus

Portuguese is a language with few publicly available resources for speech synthesis. In Portuguese, as far as we know, there is no public dataset with a large amount of speech and quality available for speech synthesis. Although there are some public

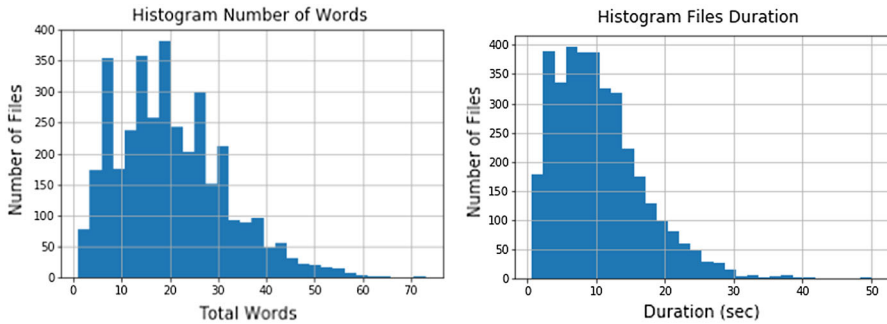


Fig. 1 The figure on the left shows the number of words and on the right the duration of each file

speech datasets for European Portuguese, for example (Teixeira et al., 2001), the work has a small amount of speech, approximately 100 min, which normally is not useful for training deep-learning models. On the other hand, Quintas and Trancoso (2020) explored the training of a model based on deep learning with an in-house dataset, called SSF1, which has approximately 14 h of speech in European Portuguese. Therefore, given the inexistence of an open dataset with a large amount of speech and quality for speech synthesis for Portuguese, we propose the TTS-Portuguese Corpus.

To create the TTS-Portuguese Corpus, public domain texts were used. Initially, seeking to reach a large vocabulary we extracted articles from the Highlights sections of Wikipedia for all knowledge areas. After this extraction, we separated the articles into sentences (considering textual punctuation) and randomly selected sentences from this corpus during the recording. In addition, we used 20 sets of phonetically balanced sentences, each set containing ten sentences proposed by Seara (1994). Finally, in order to increase the number of questions and introduce more expressive speech, we extracted sentences from Chatterbot-corpus¹, a corpus originally created for the construction of chatbots. Therefore, we decided both to have a large vocabulary and also to bring words from different areas. In addition, to have an expressive speech representation with the use of questions and answers from a chatbot dataset.

The recording was made by a male native Brazilian Portuguese speaker, not professional, in quiet environment but without acoustic isolation due to difficulties having access to studios. All the audios were recorded at a sampling frequency of 48 kHz and a 32-bit resolution.

In the dataset, each audio file has its respective textual transcription (phonetic transcription is not provided). The final dataset consists of a total of 71,358 words spoken by the speaker, 13,311 unique words, resulting in 3632 audio files and totaling 10 h and 28 min of speech. Audio files range in length from 0.67 to 50.08 s. The Fig. 1 shows two histograms regarding the number of words and the duration of each file.

¹ <https://github.com/gunthercox/chatterbot-corpus/>.

Table 1 Comparison between LJ Speech, SSF1 and TTS-Portuguese datasets in terms of language, duration, sampling rate and proportion of sentences: interrogative (Int), exclamatory (Exc) and declarative (Dec)

Dataset	Language	Duration (h)	Sampling rate (kHz)	Int. (%)	Exc. (%)	Dec. (%)
LJ Speech	EN	24	22.05	0.58	0.35	99.07
SSF1	PT-PT	14	16	11	0.7	88.3
This work	PT-BR	10	48	3.42	0.38	96.96

To compare the TTS-Portuguese Corpus with datasets used in the literature for speech synthesis, we chose the LJ Speech dataset (Ito, 2017), which is one of the most widely used, publicly available datasets, for training single-speaker models in the English language. Additionally, we present the statistics for the SSF1 dataset, which is a corpus of European Portuguese explored in the work of Quintas and Trancoso (2020). Table 1 shows the language, duration, sampling rate and percentage of interrogative, exclamatory and declarative sentences in the LJ Speech, the SSF1 and the TTS-Portuguese datasets.

The TTS-Portuguese Corpus dataset has a smaller number of hours when compared to the others, it has 14 h less than the LJ Speech and 4 h less than the SSF1.

The sampling rate of 22.05 kHz is widely used in the training of TTS models based on deep learning. However, some works like Shen et al. (2018) use a sampling rate of 24 kHz. In addition, Kumar et al. (2020) showed that it is possible to obtain a 44.1 kHz TTS model by training the NU-GAN model on a dataset sampled at 44.1 kHz. Following this idea, we made available the TTS-Portuguese Corpus at 48 kHz, a sampling rate much higher than the datasets compared here. Finally, in the distribution of interrogative and exclamatory sentences, the TTS-Portuguese Corpus has less coverage for these sentences compared to the SSF1 in-house dataset. However, the TTS-Portuguese Corpus has greater coverage of these sentences than the publicly available dataset LJ Speech.

The TTS-Portuguese Corpus² is open source, and publicly available under the terms of the license Creative Commons Attribution 4.0 (CC BY 4.0)³.

4 Experiments

To evaluate the quality of the TTS-Portuguese Corpus in practice, we explored the speech synthesis using models of prominence in the literature. We chose the models: DCTTS (Tachibana et al., 2017) and Tacotron 2 (Shen et al., 2018).

Here, we compare the models DCTTS and Tacotron 2. To maintain results reproducible, we used open source implementations and tried to replicate related works as faithfully as possible. In the cases where hyper-parameters were not

² Official repository: <https://github.com/Edresson/TTS-Portuguese-Corpus>.

³ <https://creativecommons.org/licenses/by/4.0/>.

specified, we empirically optimized those for our dataset. We have used the following implementations: DCTTS provided by Park (2018) and Tacotron 2 provided by Gölge (2019).

For all experiments, to speed up training, we initialized the model using the weights of the pre-trained model on English, using the LJ Speech dataset and we also use RTISI-LA (Zhu et al., 2007) as a vocoder, which is a variation of the Griffin–Lim (1984) vocoder.

Although the acquisition avoided external noise as best as possible, the audio files were not recorded in a studio setting. Therefore, some noise may be present in part of the files. To reduce the interference with our analysis, we applied RNNoise (Valin, 2017) in all audio files. RNNoise is based on Recurrent Neural Networks; more specifically Gated Recurrent Unit (Cho et al., 2014), and demonstrated good performance for noise suppression.

We report two experiments:

- Experiment 1: replicates the implementation of the DCTTS model, training the model for Brazilian Portuguese with the TTS-Portuguese Corpus. For this experiment, as reported in the DCTTS article, the model receives the text directly as input, so no phonetic transcription is used. As previously mentioned, the original DCTTS paper does not describe any normalization, so for the model to converge we tested different normalization options and decided to use, in all layers, 5% dropout and layer normalization (Ba et al., 2016). We did not use a fixed learning rate as described in the original article. Instead, we used a starting learning rate of 0.001 decaying using Noam’s learning rate decay scheme (Vaswani et al., 2017).
- Experiment 2: this experiment explores Tacotron 2 (Shen et al., 2018) model, for that we use the Mozilla TTS implementation (Gölge, 2019). This model receives phonetic transcription as input instead of text directly. To perform phonetic transcription we use the Phonemizer⁴ library that supports 121 languages and accents.

In experiment 1, two parts of the model are trained separately. The first part of the model, called Text2Mel, is responsible for generating a Mel spectrogram from the input text and this part of the model was induced using the composition of the functions: binary cross-entropy, L1 (Goodfellow et al., 2016) and guided attention loss (Tachibana et al., 2017). The second part, called SSRN, is responsible for the transformation of a mel spectrogram into the complete STFT spectrogram and applies super-resolution in the process and the loss function is composed of the functions L1 and binary cross-entropy.

In experiment 2, no guided attention is used, therefore, the loss function did not include the cost of attention. Since the network is trained end-to-end, the loss depends on the output of two network modules. The first module converts text into Mel spectrogram. The second module is a SSRN-like module called CBHG (1-D Convolution Bank Highway Network Bidirectional Gated Recurrent Unit).

⁴ <https://github.com/bootphon/phonemizer>.

Table 2 Hardware specifications of the computers used in training

Specifications	Computer 1	Computer 2
Processor	i7-8700	i7-7700
RAM memory	16 GB	32 GB
Video Card	Nvidia GeForce Gtx Titan V	Nvidia GeForce Gtx 1080 TI
Operational system	Ubuntu 18.04	Windows 10

Table 3 Model training

Experiment	Training steps	Time
Experiment 1 (Text2Mel/SSRN)	2115 k/2019 k	4 days 19 h/5 days 22 h
Experiment 2	261 k	9 days 7 h

Table 2 shows the hardware specifications of the equipment used for model training. Experiment 1 was trained on computer 2, while experiment 2 were performed using computer 1.

Table 3 presents the training data from the experiments. The metrics presented in the table are: number of training steps, and the time required for training. It is important to note that experiment 1 is trained in two phases, both reported in the table: Text2Mel and SSRN.

5 Results and discussion

To compare and analyze our results we used the Mean Opinion Score (MOS) calculated following the work of Ribeiro et al. (2011). To calculate the MOS, the evaluators were asked to assess the naturalness of generated sentences on a five-point scale (1: Bad, 2: Poor, 3: Fair, 4: Good, 5: Excellent). We chose 20 phonetically balanced sentences (Seara, 1994) not seen in the training, so that our analysis has a good phonetic coverage. However, in this set of sentences there are no interrogative and exclamatory sentences. These sentences were synthesized for each of our experiments. In addition, 20 samples with the pronunciation of these sentences by the original speaker were added as ground truth. Each sample was evaluated by 20 native evaluators. Our Models, synthesized audios, corpus and an interactive demo are public available.⁵

Table 4 presents the MOS values, with their respective 95% confidence intervals, for our experiments and for the best experiment of the Quintas and Trancoso (2020), which can also be seen in Fig. 2.

⁵ <https://edresson.github.io/TTS-Portuguese-Corpus/>.

Table 4 MOS results

Experiment	MOS (rank)
Ground truth—SSF1	4.42 ± 0.60 (–)
Quintas and Trancoso (2020)	3.82 ± 0.69 (–)
Ground truth—Our	4.71 ± 0.16 (–)
Experiment 1	3.03 ± 0.34 (2)
Experiment 2	4.02 ± 0.27 (1)

The results of the main analysis indicate that experiment 2 (Mozilla TTS) presented the best MOS value (4.02). According to Ribeiro et al. (2011), the obtained value indicates a good quality audio, with a barely perceptible, but not annoying, distortion. On the other hand, experiment 1 presented a MOS of 3.03 indicating a perceptible and slightly annoying distortion in the audios.

With respect to previous results in the English language, Shen et al. (2018) (Tacotron 2) can be compared to our experiment 2. The authors trained their model on an in-house US English dataset (24.6 h), reaching a MOS of 4.52 ± 0.06 . Considering the confidence intervals, our model reaches 4.29 in the best case and 3.75 in the worst case. Therefore, our model has a slightly lower MOS, and this can be justified as Shen et al. (2018) uses the WaveNet vocoder which achieves a higher quality in relation to the RTISI-LA/Griffin–Lim vocoder as shown in (Ping et al., 2017).

Finally, regarding DCTTS we obtained 3.03 ± 0.34 MOS. The original paper (Tachibana et al., 2017) had 2.71 ± 0.66 on the LJ Speech dataset. Therefore, the Tachibana et al. (2017) model can at best achieve a MOS of 3.37 and in the worst case 2.05. On the other hand, our model can at best achieve a MOS of 3.37 and in the worst case 2.69. Thus, the DCTTS model trained in the LJ Speech dataset and the TTS-Portuguese Corpus dataset showed similar MOS results.

It is also possible to compare our results with related works in European Portuguese. The current state of the art (SOTA) in European Portuguese (Quintas & Trancoso, 2020) achieved a MOS score of 3.82 ± 0.69 when training Tacotron 2 on the in-house SSF1 dataset. Considering the confidence intervals in the best case, the model of Quintas and Trancoso (2020) can achieve a MOS of 4.51 and in the worst case of 3.13. On the other hand, as previously discussed, our best model can reach 4.29 and 3.75 in the best and worst cases, respectively. These values are compatible since in the work of Quintas and Trancoso (2020) the authors used the neural vocoder WaveNet that generates speech with a higher quality. In addition, our confidence intervals are shorter, which may indicate that our evaluators agreed more during the evaluation. Furthermore, Quintas and Trancoso (2020) used only 8 evaluators in their MOS analysis, while in this work we used 20 evaluators; the number of evaluators can also have an impact on confidence intervals.

Comparing the Ground truth for the SSF1 dataset and the TTS-Portuguese Corpus, we can see that the TTS-Portuguese Corpus can vary from 4.87 to 4.55 in




Fig. 2 MOS analysis chart

the best and worst cases, respectively. On the other hand, the MOS for the SSF1 dataset reported by Quintas and Trancoso (2020) ranges from 5.02 (a value above 5 can be justified by rounding) to 3.82. Considering this MOS analysis, the two datasets are comparable in terms of quality and naturalness.

6 Conclusions and future work

We found that it is possible to train a good quality speech synthesizer for Brazilian Portuguese using our dataset, reaching 4.02 MOS value. Our best results were based on Tacotron 2 model. We had MOS scores comparable to the SOTA paper that explores the use of deep learning in the European Portuguese language (Quintas & Trancoso, 2020), using a in-house dataset. In addition, our results are also comparable to works in the literature that used the English language.

To the best of our knowledge, this is the first publicly available single-speaker synthesis dataset for the language. Similarly, the trained models are a contribution to the Brazilian Portuguese language, since this language has limited open access models based on deep learning.

Future work can investigate the training of flow-based (Durkan et al., 2019; Hoogetboom et al., 2019; Kingma et al., 2016) TTS models, such as Flowtron (Valle et al., 2020), GlowTTS (Kim et al., 2020) and Flow-TTS (Miao et al., 2020), in the TTS-Portuguese Corpus dataset. Also, a more in-depth analysis of the datasets can be carried out to enrich it. For example, obtain the statistics of homographic-heterophone word pairs. Although time-consuming, it would provide a better learning context for the homograph disambiguation.

Acknowledgements This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior—Brasil (CAPES)—Finance Code 001, as well as CNPq (National Council of Technological and Scientific Development) Grants 304266/2020-5. This research was carried out at the Center for Artificial Intelligence (C4AI-USP), with support by the São Paulo Research Foundation (FAPESP Grant #2019/07665-4) and by the IBM Corporation. Also, we would like to thank the partial financial support for this paper provided by CEIA⁶ (Artificial Intelligence Excellence Center, in English) via a funded project by the Cyberlabs Group⁷. We also gratefully acknowledge the support of NVIDIA corporation with the donation of the GPU used in part of the experiments presented in this research.

Authors' contributions EC, ACJ and SA proposed the methodology for collecting the dataset. EC, CS, ACJ, MAP and JPT proposed the experiments. EC and FSdO implemented and trained the models. All authors contributed to the final manuscript.

Funding This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior—Brasil (CAPES)—Finance Code 001, as well as CNPq (National Council of Technological and Scientific Development) Grant 304266/2020-5.

Data availability Dataset is available on the github of the project: <https://github.com/Edresson/TTS-Portuguese-Corpus>.

Code availability The trained models and demos are available on the github of the project: <https://github.com/Edresson/TTS-Portuguese-Corpus>.

Declarations

Conflict of interest The authors have no conflicts of interest to declare.

References

- Alencar, V., & Alcaim, A. (2008). LSF and IPC-derived features for large vocabulary distributed continuous speech recognition in Brazilian Portuguese. In *2008 42nd Asilomar conference on signals, systems and computers* (pp. 1237–1241). IEEE.
- Arik, S. O., Chrzanowski, M., Coates, A., Damos, G., Gibiansky, A., Kang, Y., Li, X., Miller, J., Raiman, J., & Sengupta, S., & Ng, A. (2017). Deep voice: Real-time neural text-to-speech. *arXiv preprint*. <http://arxiv.org/abs/170207825>
- Ank, S. O., Damos, G., Gibiansky, A., Miller, J., Peng, K., Ping, W., Raiman, J., & Zhou, Y. (2017). Deep voice 2: Multi-speaker neural text-to-speech. *arXiv preprint*. <http://arxiv.org/abs/170508947>
- Aron, A., & Dhonde, S. (2015). Statistical parametric speech synthesis: A review. In *2015 IEEE 9th international conference on intelligent systems and control (ISCO)* (pp. 1–5). IEEE.
- Ba, J. L., Kiros, J. R., & Hinton, G. E. (2016). Layer normalization. *arXiv preprint*. <http://arxiv.org/abs/160706450>
- Bahdanau, D., Cho, K., & Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. *arXiv preprint*. <http://arxiv.org/abs/14090473>
- Benesty, J., Chen, J., & Habets, E. A. (2011). *Speech enhancement in the STFT domain*. Springer Science & Business Media.
- Braude, D. A., Shimodaira, H., & Youssef, A. B. (2013). Template-warping based speech driven head motion synthesis. In *Interspeech* (pp. 2763–2767).

⁶ <http://centrodeia.org/>.

⁷ <https://cyberlabs.ai/>.

- Charpentier, F., & Stella, M. (1986). Diphone synthesis using an overlap-add technique for speech waveforms concatenation. In *ICASSP'86. IEEE international conference on acoustics, speech, and signal processing* (Vol. 11, pp. 2015–2018). IEEE.
- Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder–decoder for statistical machine translation. *arXiv preprint*. <http://arxiv.org/abs/14061078>
- Chung, J., Gulcehre, C., Cho, K., & Bengio, Y. (2014). Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint*. <http://arxiv.org/abs/14123555>
- Dempsey, P. (2017). The teardown: Google home personal assistant. *Engineering & Technology*, 12(3), 80–81.
- Durkan, C., Bekasov, A., Murray, I., & Papamakarios, G. (2019). Neural spline flows. In *Advances in neural information processing systems* (pp. 7511–7522)
- Goodfellow, I., Bengio, Y., Courville, A., & Bengio, Y. (2016). *Deep learning* (Vol. 1). MIT Press.
- Griffin, D., & Lim, J. (1984). Signal estimation from modified short-time Fourier transform. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 32(2), 236–243.
- Gruber, T. R. (2009) Siri, a virtual personal assistant-bringing intelligence to the interface. In *Semantic technologies conference*.
- Gölge, E. (2019). Deep learning for text to speech. <https://github.com/mozilla/TTS>
- Hoogeboom, E., Van Den Berg, R., & Welling, M. (2019). Emerging convolutions for generative normalizing flows. *arXiv preprint*. <http://arxiv.org/abs/19011137>
- Ito, K. (2017). The Ij speech dataset. Retrieved April 29, 2020, from <https://keithito.com/LJ-Speech-Dataset/>
- Kalchbrenner, N., Espeholt, L., Simonyan, K., van den Oord, A., Graves, A., & Kavukcuoglu, K. (2016). Neural machine translation in linear time. *arXiv preprint*. <http://arxiv.org/abs/161010099>
- Kim, J., Kim, S., Kong, J., & Yoon, S. (2020). Glow-TTS: A generative flow for text-to-speech via monotonic alignment search. *arXiv preprint*. <http://arxiv.org/abs/200511129>
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., & Welling, M. (2016). Improved variational inference with inverse autoregressive flow. In *Advances in neural information processing systems* (pp. 4743–4751).
- Klatt, D. H. (1980) Software for a cascade/parallel formant synthesizer. *The Journal of the Acoustical Society of America*, 67(3), 971–995.
- Kumar, R., Kumar, K., Anand, V., Bengio, Y., & Courville, A. (2020) NU-GAN: High resolution neural upsampling with GAN. *arXiv preprint*. <http://arxiv.org/abs/201011362>
- Mehri, S., Kumar, K., Gulrajani, I., Kumar, R., Jain, S., Sotelo, J., Courville, A., Bengio, Y., & Sample R. N. N. (2017). Char2wav: End-to-end speech synthesis. In *International conference on learning representations, workshop*.
- Miao, C., Liang, S., Chen, M., Ma, J., Wang, S., & Xiao, J. (2020). Flow-TTS: A non-autoregressive network for text to speech based on flow. In *ICASSP 2020–2020 IEEE international conference on acoustics, speech and signal processing(ICASSP)* (pp. 7209–7213). IEEE.
- Morise, M., Yokomori, F., & Ozawa, K. (2016). World: A vocoder-based high-quality speech synthesis system for real-time applications. *IEICE Transactions on Information and Systems*, 99(7), 1877–1884.
- Park, K. (2018). A tensorflow implementation of DC-TTS. https://github.com/Kyubyong/dc_tts
- Ping, W., Peng, K., Gibiansky, A., Arik, S. O., Kannan, A., Narang, S., Raiman, J., & Miller, J. (2017). Deep voice 3: 2000-speaker neural text-to-speech. *arXiv preprint*. <http://arxiv.org/abs/171007654>
- Pratap, V., Xu, Q., Sriram, A., Synnaeve, G., & Collobert, R. (2020). MIs: A large-scale multilingual dataset for speech research. In *Proceedings of Interspeech 2020* (pp. 2757–2761).
- Purinton, A., Taft, J. G., Sannon, S., Bazarova, N. N., & Taylor, S. H (2017). “Alexa is my new BFF”: Social roles, user satisfaction, and personification of the Amazon echo. In *Proceedings of the 2017 CHI conference extended abstracts on human factors in computing systems* (pp. 2853–2859).
- Quintanilha, I. M., Netto, S. L., & Biscainho, L. W. P. (2020). An open-source end-to-end ASR system for Brazilian Portuguese using DNNs built from newly assembled corpora. *Journal of Communication and Information Systems*, 35(1), 230–242.
- Quintas, S., & Trancoso, I. (2020). Evaluation of deep learning approaches to text-to-speech systems for European Portuguese. In *International conference on computational processing of the Portuguese language* (pp. 34–42). Springer.

- Ribeiro, F., Florêncio, D., Zhang, C., & Seltzer, M. (2011). Crowdmoss: An approach for crowdsourcing mean opinion score studies. In *2011 IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 2416–2419). IEEE.
- Seara, I. (1994). Estudo estatístico dos fonemas do português brasileiro falado na capital de santa catarina para elaboração de frases foneticamente balanceadas. PhD thesis, Dissertação de Mestrado, Universidade Federal de Santa Catarina ...
- Shen, J., Pang, R., Weiss, R. J., Schuster, M., Jaitly, N., Yang, Z., Chen, Z., Zhang, Y., Wang, Y., Skerry-Ryan, R., & Saurous R. A. (2018). Natural TTS synthesis by conditioning WaveNet on Mel spectrogram predictions. In *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 4779–4783). IEEE.
- Siddhi, D., Verghese, J. M., & Bhavik, D. (2017). Survey on various methods of text to speech synthesis. *International Journal of Computer Applications*, *165*(6), 26–30.
- Sotelo, J., Mehri, S., Kumar, K., Santos, J. F., Kastner, K., Courville, A., & Bengio, Y. (2017). Char2wav: End-to-end speech synthesis. In *International conference on learning representations, workshop*.
- Srivastava, R. K., Greff, K., & Schmidhuber, J. (2015). Training very deep networks. In *Advances in neural information processing systems* (pp. 2377–2385).
- Tachibana, H., Uenoyama, K., & Aihara, S. (2017). Efficiently trainable text-to-speech system based on deep convolutional networks with guided attention. *arXiv preprint*. <http://arxiv.org/abs/171008969>
- Tamamori, A., Hayashi, T., Kobayashi, K., Takeda, K., & Toda, T. (2017). Speaker-dependent WaveNet vocoder. In *Proceedings of Interspeech* (pp. 1118–1122).
- Teixeira, J. P., Freitas, D., Braga, D., Barros, M. J., & Latsch, V. (2001). Phonetic events from the labeling the European Portuguese database for speech synthesis, FEUP/IPBDB. In *Seventh European conference on speech communication and technology*.
- Teixeira, J. P., Freitas, D., & Fujisaki, H. (2003). Prediction of Fujisaki model's phrase commands. In *Eighth European conference on speech communication and technology*.
- Tokuda, K., Yoshimura, T., Masuko, T., Kobayashi, T., & Kitamura, T. (2000). Speech parameter generation algorithms for HMM-based speech synthesis. In *2000 IEEE international conference on acoustics, speech, and signal processing. Proceedings (Cat. No. 00CH37100)* (Vol. 3, pp. 1315–1318). IEEE.
- Valin, J. M. (2017). A hybrid DSP/deep learning approach to real-time full-band speech enhancement. *arXiv preprint*. <http://arxiv.org/abs/170908243>
- Valle, R., Shih, K., Prenger, R., & Catanzaro, B. (2020). Flowtron: An autoregressive flow-based generative network for text-to-speech synthesis. *arXiv preprint*. <http://arxiv.org/abs/200505957>
- Van Den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., & Kavukcuoglu, K. (2016). Wavenet: A generative model for raw audio. *arXiv preprint*. <http://arxiv.org/abs/160903499>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems* (pp. 5998–6008).
- Wang, W. Y., & Georgila, K. (2011). Automatic detection of unnatural word-level segments in unit-selection speech synthesis. In *2011 IEEE workshop on automatic speech recognition & understanding* (pp. 289–294). IEEE.
- Wang, Y., Skerry-Ryan, R., Stanton, D., Wu, Y., Weiss, R. J., Jaitly, N., Yang, Z., Xiao, Y., Chen, Z., & Bengio, S., & Le, Q. V. (2017). Tacotron: A fully end-to-end text-to-speech synthesis model. *arXiv preprint*. <http://arxiv.org/abs/170310135>
- Yu, F., & Koltun, V. (2015). Multi-scale context aggregation by dilated convolutions. *arXiv preprint*. <http://arxiv.org/abs/151107122>
- Ze, H., Senior, A., & Schuster, M. (2013). Statistical parametric speech synthesis using deep neural networks. In *2013 IEEE international conference on acoustics, speech and signal processing* (pp. 7962–7966). IEEE.
- Zhu, X., Beauregard, G. T., & Wyse, L. L. (2007). Real-time signal estimation from modified short-time fourier transform magnitude spectra. *IEEE Transactions on Audio, Speech, and Language Processing*, *15*(5), 1645–1653.