

Exploring Transfer Learning for Low Resource Emotional TTS

Noé Tits, Kevin El Haddad and Thierry Dutoit

University of Mons,
Numediart Institute,
Mons, Belgium 7000
{noe.tits, kevin.elhaddad, thierry.dutoit}@umons.ac.be

Abstract. During the last few years, spoken language technologies have known a big improvement thanks to Deep Learning. However Deep Learning-based algorithms require amounts of data that are often difficult and costly to gather. Particularly, modeling the variability in speech of different speakers, different styles or different emotions with few data remains challenging. In this paper, we investigate how to leverage fine-tuning on a pre-trained Deep Learning-based TTS model to synthesize speech with a small dataset of another speaker. Then we investigate the possibility to adapt this model to have emotional TTS by fine-tuning the neutral TTS model with a small emotional dataset.

Keywords: Speech Synthesis, Emotion, Deep Learning, Transfer Learning, Fine-Tuning

1 INTRODUCTION

The current state of the art of Text-to-Speech (TTS) synthesis is based on deep learning algorithms. These systems are now capable of producing natural human-like speech.

There are more and more deep learning-based TTS systems developed. The Merlin toolkit [1] has played an important part in this development, first with simple architectures based on stacks of Fully Connected layers and then with more complex ones such as Recurrent Neural Networks (RNNs).

Most recent TTS systems, such as Wavenet [2], Tacotron [3], WaveRNN [4], Char2Wav [5] and Deep Voice [6], achieve excellent results in terms of naturalness. However they require tens of hours of speech data and a lot of computational power. A first system that aims to synthesize speech with few computational power is Deep Convolutional TTS [7] (DCTTS). This system is only based on Convolutional Neural Networks (CNNs) and avoids using RNNs known to be difficult to train due to the vanishing gradient issue during gradient descent [8]. In their experiments, the authors of DCTTS were able to train their model in 15 hours using a standard PC with two GPUs, resulting in nearly acceptable speech synthesis.

Moreover it is difficult to have a fine control on speech quality and emotional content with such systems, while this has become an important challenge in speech synthesis. Here again, data availability is an issue. Indeed, high quality speech datasets with emotional content needed for speech synthesis are quite difficult to collect. The amount of data available is therefore relatively limited compared to what deep learning algorithms require to converge.

Promising methods to tackle the problem of quantity of data are those related to knowledge transfer such as transfer learning [9], fine-tuning and multi-task learning. These techniques have proved useful in various applications of deep learning.

In the field of Motion Capture and Analysis, [10] mapped a motion sequence to an RGB image to be able to use a CNN pre-trained for image classification in their motion classification task. They showed that fine-tuning the CNN on their motion data improved classification results.

In a previous work of ours [11], we used an neural Automatic Speech Recognition (ASR) as a feature extractor for emotion recognition. We showed that the mapping between speech and text learned by the ASR system contains information useful for emotion recognition.

In the TTS field as well, transfer learning is being investigated. In [12], they successfully transferred knowledge from a model trained to discriminate between speakers to a multi-speaker TTS model.

These examples motivates our interest to investigate the use of knowledge transferability between models. Our goal is to use this to tackle the inherent problem of the amount of data needed in deep learning in the case of emotional TTS which is a topic of growing interest. We can cite [13] that experimented an unsupervised learning technique to change prosody of synthesized sentences with style tokens or [14] that modified Tacotron’s architecture to synthesize speech given emotional labels. In this paper we explain how to leverage fine-tuning on deep learning-based TTS systems to synthesize emotional speech with a small emotional speech dataset.

2 SYSTEM

Our goal is to study the feasibility of fine-tuning a TTS system pre-trained on a big dataset on few new data and analyze how much the model is able to fit them. This Section describes the whole system. Figure 1 represents its overall idea. First, in Section 2.1, we present the TTS system used as a basis for fine-tuning. We then briefly present the dataset we are using in Section 2.2. In Section 2.3, we explain the pre-processing of our dataset. Finally, in Section 2.4, we detail the fine-tuning procedure applied to obtain emotional TTS models.

2.1 TTS System

The number of deep learning-based TTS system of the state of the art are growing. To carry out our experiments, we chose, DCTTS [7], a system that seems to

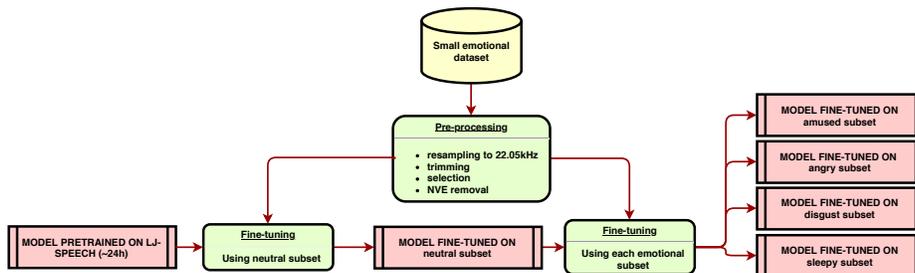


Fig. 1. Block diagram of the system

combine advantages of several systems. DCTTS models a sequence-to-sequence problem with an encoder-decoder structure along with an Attention Mechanism like Tacotron [3]. However, unlike Tacotron, the modules of the architecture are all CNN-based and there is no RNN component. In [7], they compared an open source implementation of Tacotron to DCTTS and have higher Mean Opinion Score (MOS).

In this work, we use the Tensorflow implementation provided in [15].

There are two modules trained separately: Text2Mel and SSRN (for Spectrogram Super-resolution Network). Text2Mel takes care of the mapping between character embeddings and the output of Mel Filter Banks (MFBs) applied on the audio signal, that is, a mel-spectrogram. Then the second module SSRN maps the mel-spectrogram to full resolution spectrogram. Finally, Griffin-Lim is used as a vocoder.

Text2Mel module models the sequence-to-sequence task. It is composed of a Text Encoder, an Audio Encoder, an Attention Mechanism, and an Audio Decoder.

2.2 Dataset Used

The dataset used in this work is EmoV-DB: The Database of Emotional Voices [16] that is available online ¹.

EmoV-DB contains sentences uttered by male and female actors in English and a male actor in French. The English actors were recorded in 2 different anechoic chambers of the Northeastern University campus while the French actor was recorded in an anechoic chamber of the University of Mons.

Each actor was asked to utter a subset of the sentences from the CMU-Arctic [17] database for English speakers and from the SIWIS [18] database for the French speaker. The actors uttered these sentences with 5 emotion classes making it possible to build synthesis and voice transformation systems. For every speaker, the different emotions were recorded in different sessions.

In this work we used one of the English actress to perform emotion adaptation of the TTS system.

¹ <https://github.com/numediart/EmoV-DB>

The experiments performed on this dataset also assess its usability with deep learning algorithms for voice generation systems.

2.3 Pre-processing

Important aspects of the pre-processing to use this model are

- the sample frequency
- the trimming of silences at the beginning and end of audio files
- removal of non-verbal expressions (laughters, yawns, etc.)

As the model was trained with LJ-speech database with a sample frequency of 22050Hz, we should use the same with our database.

The trimming of silences is important because the model use guided attention [7]. It helps the attention mechanism by assuming that the ordering of characters is almost linearly related to the time in the audio file. This is true only if the speech begins from the start of the file without a silence.

We experimented that without this trimming, the synthesized sentences often omitted the first words of the text to pronounce. The implementation of [15] already use trimming with librosa library. However we noticed that default parameters were not suited for our database and changed *top_db* to 20dB.

The same problem occurs due to the nonverbal vocalizations [19] (NVV) present in the audio files such as laughs, yawns or sighs. Indeed in these cases, the hypothesis to use guided attention is not verified either. To overcome this problem, we first manually selected utterances without such NVV for the amused dataset (156 utterances) and the sleepy dataset (361 utterances). To increase the amount of amused sentences, we manually trimmed the laughs of 82 of the remaining amused sentences to have a total of 238 utterances.

2.4 Fine-tuning

In this Section, we explain how we leveraged knowledge transfer on TTS by fine-tuning a part of a pre-trained model on our small dataset.

The pre-training of the model was done using the LJ-Speech dataset. This dataset is available online² and contains 23.9h of speech uttered by a single female speaker.

The fine-tuning was done with the dataset described in Section 2.2.

There are several possibilities of how we could fine-tune the model. First, we can choose which parts of the pre-trained model we want to fine-tune with the new dataset and which part we want to keep fixed.

The second part of the model, SSRN, does the mapping between MFBS and full spectrogram. Therefore, it should not depend on the speaker identity on speaking style as it is just trained to do the mapping between two audio features. However, as the model has been pre-trained on one speaker, there is a

² <https://keithito.com/LJ-Speech-Dataset>

possibility of over-fitting on the characteristics of that specific speaker. The first question we want to answer is whether the SSRN can generalize the mapping to other speaking styles.

As for Text2Mel module, it is composed of a Text Encoder, Audio Encoder, Attention, and Audio Decoder. As the text does not depend on characteristics of the speaker or his speaking style, we tried to train only the Audio part. However we found that there are some problems of rhythm in synthesized speech. We believe this is because Attention module is not adapted to the new speaking style. As a consequence, we chose to fine-tune on the entire Text2Mel module.

3 EXPERIMENT

In this Section, we detail the experiments performed on the system.

In the first experiment, we evaluate the usefulness of the fine-tuning technique compared to a random initialization of the parameters of the model. The evaluation is based on a measure of intelligibility of the synthesized speech in terms of word accuracy proposed in [20].

In the second experiment, we evaluate the quality of the emotional speech synthesized through a MOS test for each emotion according to the confidence in the perception of the emotion specified.

The amount of speech data used for the experiments are showed in Table 1. Durations are rounded to the minute. The values between parentheses correspond to the amount of data before selection and NVV removal.

Table 1. Amount of data available for each emotion in terms of total duration and number of utterances

	Total duration [min]	Number of utterances
Amused	15 (20)	238 (296)
Angry	19	304
Disgusted	29	303
Neutral	23	357
Sleepy	36 (51)	361 (496)

3.1 Objective measures

In this experiment, we synthesized 100 sentences of the Harvard sentences [21] with several models. Then an objective measure of the intelligibility of every sentence was computed in terms of word accuracy [20]. The measure consists of using an ASR to recognize speech and compute a word accuracy by comparing the result to the text label. The mean word accuracy with 95% confidence interval for all models are summarized in Table 2. The first line show the word accuracy

Table 2. Intelligibility in terms of Word Accuracy.

	Word Accuracy
LJ-speech	0.630 ± 0.042
Neutral (random initialization)	0.004 ± 0.004
Neutral (fine-tuning)	0.517 ± 0.048

of the pre-trained model (using LJ-speech). The second line corresponds to the model trained only on the neutral subset. Finally, the third line corresponds to the pre-trained model fine-tuned on the neutral subset.

These measures allow us to compare the fine-tuning of the parameters of a pre-trained model with the random initialization of the parameters of the model.

The experiments clearly shows that model trained with the neutral subset of 20 min is unable to generate intelligible speech if the parameters are randomly initialized. However, if the initialization of the parameters

3.2 Perception tests

After fine-tuning from the neutral model to emotional models, we synthesized 5 sentences not seen during training with each of these models. These sentences were used in a Mean Opinion Score (MOS) test.

During this test, the participants were asked to complete a form.

This survey contained 5 sections. Each section was dedicated to one emotion. In every section, the participants were asked to rate utterances between 0 and 5 for the confidence in the perception of the emotion specified (0=we can not hear the emotion specified, 5=we perfectly hear the emotion specified).

This test was performed on both original files from the dataset and synthesized files. Table 3 gives MOS with 95% confidence interval for the original files and Table 4 gives them for the synthesized files.

Table 3. MOS test results of original files

	Confidence
Amused	4.60 ± 0.20
Angry	4.22 ± 0.25
Disgusted	3.28 ± 0.27
Neutral	4.37 ± 0.23
Sleepy	3.80 ± 0.27

Results of Table 3 should be considered as higher bounds as they represent the opinion about original files of the dataset. Results from Table 4 should be

Table 4. MOS test results of synthesized files

	Confidence
Amused	2.00 ± 0.27
Angry	2.10 ± 0.28
Disgusted	2.27 ± 0.30
Neutral	3.59 ± 0.24
Sleepy	3.29 ± 0.26

compared to these higher bounds. One can see that these higher bounds do not meet the maximum value of 5.

The confidence of the perception of an emotion is altered by the recording quality, the playing quality and the emotion expressed. Moreover, the test was done on non-native English speakers. In [22], they conducted an experiment showing a disproportionate disadvantage for the non-native English speaker when listening to synthesized speech compared to their native English speaker.

In Table 4, for the confidence of the perception of an emotion, the Neutral category has the higher value. A possible explanation of this is that the pre-trained model used has been trained with a neutral corpus and is therefore closer to the Neutral subset used for fine-tuning. For the other emotional categories, the values are a little degraded but still significantly higher than 0 which shows that the emotion is perceptible in the synthesized speech.

4 CONCLUSIONS AND FUTURE WORK

In this paper, we present a technique allowing to synthesize emotional speech using a small emotional speech dataset. This technique is based on the fine-tuning of a deep learning-based TTS model with the neutral subset of the small dataset and then fine-tuning the resulting model with each emotional subset to obtain one model per emotional category.

In the first experiment, we show that training the model with random initialization of the parameters gives completely unintelligible speech synthesis. However using a pre-trained model as initialization and fine-tuning it allows to have intelligible speech.

In the second experiment, we show through perception tests that the speech synthesized is correctly perceived as emotional.

To improve these results, we plan to try a multi-speaker and multi-emotional system to be able to share the knowledge of emotional content of several speakers and several emotions of EmoV-DB. The approach in [12] seems really interesting for this.

Acknowledgments

Noé Tits is funded through a PhD grant from the Fonds pour la Formation à la Recherche dans l'Industrie et l'Agriculture (FRIA), Belgium.

References

1. Zhizheng Wu, Oliver Watts, and Simon King, “Merlin: An open source neural network speech synthesis system,” *Proc. SSW, Sunnyvale, USA*, 2016.
2. Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W. Senior, and Koray Kavukcuoglu, “Wavenet: A generative model for raw audio,” in *SSW*, 2016.
3. Yuxuan Wang, R. J. Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J. Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, Quoc V. Le, Yannis Agiomyrgiannakis, Rob Clark, and Rif A. Saurous, “Tacotron: Towards end-to-end speech synthesis,” in *INTERSPEECH*, 2017.
4. Nal Kalchbrenner, Erich Elsen, Karen Simonyan, Seb Noury, Norman Casagrande, Edward Lockhart, Florian Stimberg, Aaron van den Oord, Sander Dieleman, and Koray Kavukcuoglu, “Efficient neural audio synthesis,” *arXiv preprint arXiv:1802.08435*, 2018.
5. Jose Sotelo, Soroush Mehri, Kundan Kumar, Joao Felipe Santos, Kyle Kastner, Aaron Courville, and Yoshua Bengio, “Char2wav: End-to-end speech synthesis,” *ICLR2017 workshop submission*, 2017.
6. Sercan O Arik, Mike Chrzanowski, Adam Coates, Gregory Diamos, Andrew Gibiansky, Yongguo Kang, Xian Li, John Miller, Andrew Ng, Jonathan Raiman, et al., “Deep voice: Real-time neural text-to-speech,” *arXiv preprint arXiv:1702.07825*, 2017.
7. Hideyuki Tachibana, Katsuya Uenoyama, and Shunsuke Aihara, “Efficiently trainable text-to-speech system based on deep convolutional networks with guided attention,” *arXiv preprint arXiv:1710.08969*, 2017.
8. Sepp Hochreiter, “The vanishing gradient problem during learning recurrent neural nets and problem solutions,” *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 6, no. 02, pp. 107–116, 1998.
9. Sinno Jialin Pan and Qiang Yang, “A survey on transfer learning,” *IEEE Transactions on knowledge and data engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.
10. Sohaib Laraba, Mohammed Brahimi, Joëlle Tilmann, and Thierry Dutoit, “3d skeleton-based action recognition by representing motion capture sequences as 2d-rgb images,” *Computer Animation and Virtual Worlds*, vol. 28, no. 3-4, 2017.
11. Noé Tits, Kevin El Haddad, and Thierry Dutoit, “Asr-based features for emotion recognition: A transfer learning approach,” *arXiv preprint arXiv:1805.09197*, 2018.
12. Ye Jia, Yu Zhang, Ron J Weiss, Quan Wang, Jonathan Shen, Fei Ren, Zhifeng Chen, Patrick Nguyen, Ruoming Pang, Ignacio Lopez Moreno, et al., “Transfer learning from speaker verification to multispeaker text-to-speech synthesis,” *arXiv preprint arXiv:1806.04558*, 2018.
13. Jonathan Shen, Ruoming Pang, Ron J. Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, R. J. Skerry-Ryan, Rif A. Saurous, Yannis Agiomyrgiannakis, and Yonghui Wu, “Natural tts synthesis by conditioning wavenet on mel spectrogram predictions,” *CoRR*, vol. abs/1712.05884, 2017.
14. Younggun Lee, Azam Rabiee, and Soo-Young Lee, “Emotional end-to-end neural speech synthesizer,” *arXiv preprint arXiv:1711.05447*, 2017.
15. Park Kyubyong, “A tensorflow implementation of dc-tts: yet another text-to-speech model,” <https://github.com/Kyubyong/dc.tts>, 2018.
16. Adaeze Adigwe, Noé Tits, Kevin El Haddad, Sarah Ostadabbas, and Thierry Dutoit, “The emotional voices database: Towards controlling the emotion dimension in voice generation systems,” *arXiv preprint arXiv:1806.09514*, 2018.

17. John Kominek and Alan W Black, “The cmu arctic speech databases,” in *Fifth ISCA Workshop on Speech Synthesis*, 2004.
18. Pierre-Edouard Honnet, Alexandros Lazaridis, Philip N. Garner, and Junichi Yamagishi, “The siwis french speech synthesis database ? design and recording of a high quality french database for speech synthesis,” *Online Database*, 2017.
19. Kevin El Haddad, No Tits, and Thierry Dutoit, “Annotating nonverbal conversation expressions in interaction datasets,” in *Proceedings of Laughter Workshop 2018*, 09 2018.
20. Juan Rafael Orozco-Arroyave, JC Vdsquez-Correa, Florian Hönig, Julián D Arias-Londoño, JF Vargas-Bonilla, Sabine Skodda, Jan Ruzs, and E Noth, “Towards an automatic monitoring of the neurological state of parkinson’s patients from speech,” in *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on. IEEE*, 2016, pp. 6490–6494.
21. EH Rothausser, “Ieee recommended practice for speech quality measurements,” *IEEE Trans. on Audio and Electroacoustics*, vol. 17, pp. 225–246, 1969.
22. Diane Mayasari Alamsaputra, Kathryn J Kohnert, Benjamin Munson, and Joe Reichle, “Synthesized speech intelligibility among native speakers and non-native speakers of english,” *Augmentative and Alternative Communication*, vol. 22, no. 4, pp. 258–268, 2006.