

A Methodology for Controlling the Emotional Expressiveness in Synthetic Speech - a Deep Learning approach

Noé Tits

Numediart Institute

University of Mons

Mons, Belgium

noe.tits@umons.ac.be

Abstract—In this project, we aim to build a Text-to-Speech system able to produce speech with a controllable emotional expressiveness. We propose a methodology for solving this problem in three main steps. The first is the collection of emotional speech data. We discuss the various formats of existing datasets and their usability in speech generation. The second step is the development of a system to automatically annotate data with emotion/expressiveness features. We compare several techniques using transfer learning to extract such a representation through other tasks and propose a method to visualize and interpret the correlation between vocal and emotional features. The third step is the development of a deep learning-based system taking text and emotion/expressiveness as input and producing speech as output. We study the impact of fine tuning from a neutral TTS towards an emotional TTS in terms of intelligibility and perception of the emotion.

Index Terms—expressive speech synthesis, affective computing, deep learning, latent space, style embeddings, supervised, unsupervised

I. INTRODUCTION & MOTIVATIONS

Text-To-Speech (TTS) synthesis systems, which synthesize speech from text, have been around for decades and have improved very recently with the advent of new machine learning techniques such as Deep Neural Networks (DNN). Companies such as Google and Amazon provide free DNN-based speech synthesis systems such as WaveNet¹ [1] or Amazon Polly². These systems offer an excellent quality of speech, with voices obtained by analyzing tens of hours of neutral speech, therefore not very expressive.

The challenge faced by researchers today has therefore evolved towards the the field of expressive speech synthesis [2].

The objective of this project is to produce, not a neutral voice, but remarkable voices, similar to those of the actors, possessing a specific grain and a great capacity for expressiveness. This will make it possible to create virtual agents communicating in a more natural way, and thus to improve

the quality of the interaction with a machine, by making it closer to a human-human interaction.

In what follows, the background and related work of the field is briefly summarized and remaining challenges are emphasized.

Depending on the synthesis technique used [2], the voice is more or less natural and the synthesis parameters are more or less numerous. These parameters allow to create variations in the voice. The number of parameters is therefore important for the synthesis of expressive speech.

While formant synthesis can control many parameters, the resulting voice is unnatural. Synthesizers using the principle of concatenation of speech segments seem more natural but allow the control of few parameters.

The statistical approaches based on the use of HMM allow to obtain a natural synthesis as well as a control of many parameters [3]. The most recent statistical approach uses DNN [4] and is the basis of new speech synthesis systems such as WaveNet [1] and Tacotron [5]. The improvement provided by this technique [6] comes from the replacement of decision trees by DNNs and the replacement of state prediction (HMM) by frame prediction. This improvement is very interesting for the expression of emotions.

There has been numerous studies to assess DNN based systems ability to synthesize speech with high quality and naturalness such as Wavenet [1], Tacotron [5], Deep Convolutional TTS [7] (DCTTS), WaveRNN [8], Char2Wav [9] and Deep Voice [10]. However, important issue in the literature is the need of tens of hours of speech data and a lot of computational power.

DCTTS needs less computational resources than the others. In [7], they have suggested that they were able to train a TTS model in 15 hours using a regular computer with two GPUs, resulting in nearly acceptable speech synthesis.

A closer look to the literature on control of speech quality and emotional content, however, reveals a number of gaps. Access to a large amount of data remains a problem. The recording of high quality emotional speech datasets is expensive and time consuming. The amount of data available is therefore

¹<https://cloud.google.com/text-to-speech/>

²<https://aws.amazon.com/fr/polly/>

relatively limited compared to the needs of deep learning algorithms. This problem has previously been addressed using promising methods related to knowledge transfer such as transfer learning [11], fine-tuning and multi-task learning.

Concerning the controllable aspect of TTS system, an important issue is the labelization of speech data with style or emotional information. Recent studies were conducted on techniques of unsupervised learning to achieve controllable speech synthesis without the need of labels.

In [12], the authors investigated on an extension to the Tacotron speech synthesis system. The technique they propose is a system that learns a latent embedding space by encoding audio features into a vector that is concatenated to text information and then fed to Tacotron. These latent embeddings model the remaining variability of speech after accounting for variation due to phonetics, speaker identity, and channel effects.

In [13], a Variational Auto-encoder (VAE) is combined with VoiceLoop, a TTS system.

A series of recent studies have used the concept of VAE [14]–[16]. In [14], they combine a VAE with Gaussian Mixture Model (GMM) and name it GMVAE. In [15], they combine Vector Quantization with a VAE (VQ-VAE).

These works has provided evidence that building a latent space can lead to variables useful to control style in speech synthesis. In [14], they demonstrate that their system can synthesize spectrograms with different rhythms, speaking rates and F_0 's from a single text, proving a control on these parameters.

A weakness of these studies is that do not provide insights about the relationships between the resulting latent space and the audio features possible to control. We aim to fill this gap.

Section II exposes the methodology we propose to achieve the objectives of the project in three main points. The contributions made to these three tasks are presented in Section III, IV and V. Section III addresses the problem of data collection. Section IV presents techniques of emotion representation and their applicability for the control of speech synthesis. Section V considers the problem of speech synthesis with emotional content. Finally, Section VI concludes and describes the remaining challenges and plans for the future of this project.

II. PLAN

To synthesize expressive speech, the main problem is the variability of the vocal expression of emotions. Deep Learning has been effective at handling complex data but requires a large amount of these annotated data. The difficulty is then to annotate large databases with expressive metadata that is very subjective and not well defined.

The goal is to create an automatic system for annotating large expressive vocal databases, and synthesizing expressive speech.

Our work plan contains 3 main tasks:

- Collect emotional speech data

- Build a system able to extract a representation of emotional expressiveness in speech
- Build a system able to synthesize expressive speech based on the data collected and the representation extracted

This mapping will be used to create a tool that will control the expressiveness of synthesized speech by varying the dimensions continuously. This system will then be integrated into a speech synthesis instrument from which one will be able to play an utterance as he wishes.

In a second part of the thesis, we will focus on the automatic control of the expressive dimensions. The goal will then be to synthesize expressive speech based on the text to be produced and the emotion of the interlocutor. The tool must respond to the interlocutor imitating its expressiveness.

III. DATASETS

As of today, there are a lot of speech datasets available online. A good source of speech data are audiobooks (LJ-Speech³, LibriSpeech [17], LibriTTS [18]) in which we have automatically transcriptions corresponding to speech.

Other datasets come from research projects. Some are designed for speech generation purpose (voice conversion [19], speech synthesis [20]–[22]) while others are designed for a speech analysis goal (ASR, emotion in speech etc.). The last often consist of conversational setups and contain overlaps in speech as well as noise compared to datasets for speech generation. This is the reason why they cannot be used for emotional speech synthesis.

These datasets contain a fair amount of data for the use of deep learning algorithms. However, either they are not suited for synthesis purpose, or they are poor in expressiveness.

Other datasets offer emotionally rich content with a high quality, but in a limited amount [23]–[26].

To fill this gap, we collected a speech dataset called EmoV-DB Database [27] rich in expressiveness and suited for speech generation purpose in a fair amount to use deep learning algorithms.

It contains data for male and female actors in English and a male actor in French. The English actors were recorded in two 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 database for English speakers and from the SIWIS database for the French speaker. The database covers 5 classes (amusement, anger, disgust, sleepiness and neutral) with the goal to build, in the future, synthesis and voice transformation systems with the potential to control the emotional dimension in a continuous way.

It is available online and there is a script available for using gentle's forced alignment ⁴.

³<https://keithito.com/LJ-Speech-Dataset/>

⁴<https://github.com/numediart/EmoV-DB>

TABLE I
DESCRIPTION OF EMOV-DB

Type of data	Audio, text and emotion category
How data was acquired	Audio recorded in anechoic chambers of the North-eastern University campus.
Data format	Segmented in sentences, associated with transcriptions (CMU-Artic/SIWIS), classified in emotional categories
Experimental features	Recordings of sentences uttered by 2 male and 2 female speakers in 5 different emotions, making a total of 7000 sentences
Data accessibility	https://github.com/numediart/EmoV-DB

IV. EMOTION REPRESENTATION

Despite the great improvements that signal processing field has known, the field of affective computing remains challenging. As mentioned in Section II, we want to build a system able to extract a representation of emotional expressiveness in speech. In this Section we present two studies. The first one [28] investigates if a DNN trained on an Automatic Speech Recognition (ASR) task learns features useful for predicting emotion dimensions (valence, arousal).

The second one [29] investigates a good way to represent expressiveness in speech without emotional label. We evaluate these representations with a correlation analysis with the eGeMAPS feature set [30], a feature set designed for affective computing tasks.

A. ASR-based features for Emotion recognition

In this experiment, we study whether a deep learning based ASR can be used as a feature extractor for an emotion recognition task. An ASR system is a mapping between audio features and text features. To have a good accuracy, an ASR has to take into account the variability in speech to be able to recover text information from all kinds of speaker identity, style etc. Therefore, an intermediate representation extracted by the system may contain useful information to predict emotion in speech.

We used an open source ASR available online ⁵. It consists of a stack of 15 dilated convolutional layers. Each layer is constituted by 128 gated convolutional units (GCU) with skip connection. This architecture is inspired by Wavenet [1]. The system was trained with MFCCs extracted from VCTK dataset [19].

The analyzed data comes from IEMOCAP database [31]. It contains 5 sessions of dialogues between a male and a female, making a total of 10 speakers. These conversations were acted either with a script or improvised from a given situation described to the actors. The dataset is segmented in utterances and each utterance was annotated in emotional dimensions (valence, arousal, dominance) and emotion categories.

The analysis reveals that predicting valence and arousal with a linear model trained based on the output of GCUs outperform a similar model based on the eGeMAPS [30] features extracted

with opensmile ⁶, the state of the art of handcrafted features for emotion recognition.

B. Interpretation of latent spaces

In this Section, the aim is to investigate ways to represent expressiveness in speech without the need of any labels that are difficult and expensive to collect. Moreover, it is of interest to have an interpretable representation to be able to build controllable speech synthesis systems with an understandable and predictable behaviour.

We present a methodology to obtain latent spaces, visualize them and analyze their relationship with audio features by correlation. We then explain a way to visualize the influence of variation of embeddings on audio features.

We compare three techniques of encoding mel-spectrogram. The principle is to encode the mel-spectrogram by training deep learning algorithms on three different tasks. The encoding is used for:

- classification in categories of style,
- speaker classification,
- conditioning a Text-to-Speech synthesis along with the text.

The mel-spectrogram is encoded to vector of length of 8 that contains expressiveness information. This vector is computed for each utterance of the dataset to obtain a collection of embeddings.

Dimensionality reduction is then performed with UMAP to have a collection of 2D points that one can visualize (see Figure 1). We observed that points in a given area of space correspond to a given style.

To have an interpretation of the space, we analyzed the trend of variation of some audio features. To that aim, for each audio feature, i.e. the F0 mean is computed for each utterance of the dataset. Each F0 mean corresponds to one of the 2d-points of the collection. A plane $F0_{mean} = f(x, y) = ax + by + c$ is then approximated.

We hope this plan $f(x, y)$ to be a good approximation of $F0_{mean}$. If it is, it means that going in a certain direction of the space will increase F0 mean linearly. We check that by computing the correlation between the approximations $f(x, y)$ with the real values of F0 mean.

The gradient of the plane is $(\frac{\partial f}{\partial x} = a, \frac{\partial f}{\partial y} = b)$ and corresponds to the direction of its greatest slope. It thus gives the direction of the variation of the audio feature approximated by the plane.

The embedding space trained to condition a TTS system has several advantages: it does not require any label about the style and it is already in use with a TTS system.

V. SYNTHESIS

As stated in Section I, a recurrent problem while using deep learning algorithms is the need of a large amount of data, especially in fields in which data collection is costly and difficult. This is the case of emotional TTS field. In what

⁵<https://github.com/buriburisuri/speech-to-text-wavenet>

⁶<https://www.audeering.com/opensmile/>




Fig. 1. Latent space with directions of gradients of audio features

follows, we aim to study the impact of fine-tuning applied to emotional TTS.

A. Synthesis with Emotion adaptation

In this study [32], we aim to find out whether it is possible to obtain an emotional TTS system by fine-tuning a neutral TTS system with a small emotional speech dataset. We study the impact of this fine-tuning on the intelligibility of generated speech and the subjective perception of the generated speech.

The experiments were performed with an open source implementation⁷ of DCTTS [7].

A first model is trained with LJ-speech dataset, a dataset that consists of sentences uttered by a female speaker, with a total of 24 hours of speech.

To obtain emotional TTS models by fine-tuning, we proceed in two steps. The data used is a subset of EmoV-DB corresponding to one speaker. The two steps are:

- Fine-tuning of the model with the neutral part of the subset
- Fine-tuning copies of the resulting model separately with subsets of each emotional categories

To evaluate the first step, we use an objective measure of intelligibility proposed in [33] on 100 sentences synthesized with three systems: the first model trained on LJ-speech dataset, the model fine-tuned with the neutral part of the subset and a model trained only with the neutral part of the subset. The average accuracies are respectively 0.630 ± 0.042 , 0.517 ± 0.048 , and 0.004 ± 0.004 . These results suggest that the fine-tuning on a small subset tends to degrade a little the intelligibility. It also suggests however that using only the small dataset leads to unintelligible speech like sounds, therefore this fine-tuning is useful to generate more intelligible speech.

To evaluate the second step, we performed a subjective MOS test to measure the perception of the emotion perceived for each emotional model. The results show that the emotion are perceivable although less intelligible.

⁷https://github.com/Kyubyong/dc_tts

B. Controlling the intensity of emotional categories

As a demonstration, we developed a multi-emotional TTS system with the possibility to control the intensity of emotional categories. We implemented a modified version of DCTTS that takes an encoding of the emotion category at the input of the decoder. During training, a simple one-hot encoding is used. But at synthesis stage, we can modify the intensity of an emotion category by inputting encodings that are not one-hot.

The demonstration is available online⁸.

In this demonstration, only the number corresponding to neutral and the category with which we interpolate are non-zero. The sum of the numbers of a code is one.

VI. CONCLUSION & FUTURE WORKS

We propose a methodology for building a system for speech synthesis that is able to control the expressiveness of speech from user control in real time, thanks to deep learning algorithms.

Our contribution to answer the need for data is the collection of EmoV-DB that will allow to experiment on emotional TTS.

To address the challenging task of extracting emotion/expressiveness features from speech, we present systems able to extract such a representation with supervised and unsupervised techniques. We also provide a way to visualize the trends of audio features in an embedding space.

As for synthesis itself, we studied the impact of transfer learning from neutral TTS to emotional TTS and demonstrated the benefits of this technique.

We now plan to build a system that will take advantage of neutral and emotional speech datasets along with the unsupervised expressiveness representation provided by our TTS algorithm, based on a variational autoencoder (VAE-TTS).

We plan to develop a multi-speaker system with VAE-TTS, to be able to share knowledge from several datasets. The challenge will be to study if the system is able to generalize expressiveness across different speaker identities, which will require it to distinguish the effects of phonetics, speaker characteristics and expressiveness on the generated speech.

We will then integrate the system in a control interface, to provide users with a friendly-user tool to generate speech with various control levels (text, speaker identity and expressiveness).

Our ultimate goal is to use this tool in a conversational context in which the system has to detect the emotional expressiveness on its interlocutor and answer him/her by imitating his/her style applied to some predefined answer.

ACKNOWLEDGMENT

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.

⁸https://github.com/numediart/speak_with_style_demo

REFERENCES

- [1] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," in *SSW*, 2016.
- [2] F. Burkhardt and N. Campbell, "Emotional speech synthesis," in *The Oxford Handbook of Affective Computing*. Oxford University Press New York, 2014, p. 286.
- [3] H. Zen, K. Tokuda, and A. W. Black, "Statistical parametric speech synthesis," *Speech Communication*, vol. 51, no. 11, pp. 1039–1064, 2009.
- [4] H. Zen, A. Senior, and M. Schuster, "Statistical parametric speech synthesis using deep neural networks," in *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*. IEEE, 2013, pp. 7962–7966.
- [5] Y. Wang, R. J. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio, Q. V. Le, Y. Agiomyrgiannakis, R. Clark, and R. A. Saurous, "Tacotron: Towards end-to-end speech synthesis," in *INTERSPEECH*, 2017.
- [6] O. Watts, G. E. Henter, T. Merritt, Z. Wu, and S. King, "From hmms to dnns: where do the improvements come from?" in *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*. IEEE, 2016, pp. 5505–5509.
- [7] H. Tachibana, K. Uenoyama, and S. Aihara, "Efficiently trainable text-to-speech system based on deep convolutional networks with guided attention," *arXiv preprint arXiv:1710.08969*, 2017.
- [8] N. Kalchbrenner, E. Elsen, K. Simonyan, S. Noury, N. Casagrande, E. Lockhart, F. Stimberg, A. v. d. Oord, S. Dieleman, and K. Kavukcuoglu, "Efficient neural audio synthesis," *arXiv preprint arXiv:1802.08435*, 2018.
- [9] J. Sotelo, S. Mehri, K. Kumar, J. F. Santos, K. Kastner, A. Courville, and Y. Bengio, "Char2wav: End-to-end speech synthesis," *ICLR2017 workshop submission*, 2017.
- [10] S. O. Arik, M. Chrzanowski, A. Coates, G. Diamos, A. Gibiansky, Y. Kang, X. Li, J. Miller, A. Ng, J. Raiman *et al.*, "Deep voice: Real-time neural text-to-speech," *arXiv preprint arXiv:1702.07825*, 2017.
- [11] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on knowledge and data engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.
- [12] R. Skerry-Ryan, E. Battenberg, Y. Xiao, Y. Wang, D. Stanton, J. Shor, R. J. Weiss, R. Clark, and R. A. Saurous, "Towards end-to-end prosody transfer for expressive speech synthesis with tacotron," *arXiv preprint arXiv:1803.09047*, 2018.
- [13] K. Akuzawa, Y. Iwasawa, and Y. Matsuo, "Expressive speech synthesis via modeling expressions with variational autoencoder," *arXiv preprint arXiv:1804.02135*, 2018.
- [14] W.-N. Hsu, Y. Zhang, R. J. Weiss, H. Zen, Y. Wu, Y. Wang, Y. Cao, Y. Jia, Z. Chen, J. Shen *et al.*, "Hierarchical generative modeling for controllable speech synthesis," *arXiv preprint arXiv:1810.07217*, 2018.
- [15] G. E. Henter, X. Wang, and J. Yamagishi, "Deep encoder-decoder models for unsupervised learning of controllable speech synthesis," *arXiv preprint arXiv:1807.11470*, 2018.
- [16] Y. Wang, D. Stanton, Y. Zhang, R. Skerry-Ryan, E. Battenberg, J. Shor, Y. Xiao, F. Ren, Y. Jia, and R. A. Saurous, "Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis," *arXiv preprint arXiv:1803.09017*, 2018.
- [17] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [18] H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, "Libritts: A corpus derived from librispeech for text-to-speech," *arXiv preprint arXiv:1904.02882*, 2019.
- [19] C. Veaux, J. Yamagishi, K. MacDonald *et al.*, "Cstr vtck corpus: English multi-speaker corpus for cstr voice cloning toolkit," 2017.
- [20] J. Kominek and A. W. Black, "The cmu arctic speech databases," in *Fifth ISCA Workshop on Speech Synthesis*, 2004.
- [21] P.-E. Honnet, A. Lazaridis, P. N. Garner, and J. Yamagishi, "The siwfs french speech synthesis database ? design and recording of a high quality french database for speech synthesis," *Online Database*, 2017.
- [22] K. El Haddad, I. Torre, E. Gilmartin, H. Çakmak, S. Dupont, T. Dutoit, and N. Campbell, "Introducing amus: The amused speech database," in *Statistical Language and Speech Processing*, N. Camelin, Y. Estève, and C. Martín-Vide, Eds. Cham: Springer International Publishing, 2017, pp. 229–240.
- [23] S. R. Livingstone and F. A. Russo, "The ryerson audio-visual database of emotional speech and song (ravdess): A dynamic, multimodal set of facial and vocal expressions in north american english," *PLOS ONE*, vol. 13, no. 5, pp. 1–35, 05 2018.
- [24] H. Cao, D. G. Cooper, M. K. Keutmann, R. C. Gur, A. Nenkova, and R. Verma, "Crema-d: Crowd-sourced emotional multimodal actors dataset," *IEEE transactions on affective computing*, vol. 5, no. 4, pp. 377–390, 2014.
- [25] T. Bänziger, M. Mortillaro, and K. R. Scherer, "Introducing the geneva multimodal expression corpus for experimental research on emotion perception," *Emotion*, vol. 12, no. 5, p. 1161, 2012.
- [26] F. Burkhardt, A. Paeschke, M. Rolfes, W. F. Sendlmeier, and B. Weiss, "A database of german emotional speech," in *Ninth European Conference on Speech Communication and Technology*, 2005.
- [27] A. Adigwe, N. Tits, K. E. Haddad, S. Ostadabbas, and T. Dutoit, "The emotional voices database: Towards controlling the emotion dimension in voice generation systems," *arXiv preprint arXiv:1806.09514*, 2018.
- [28] N. Tits, K. El Haddad, and T. Dutoit, "Asr-based features for emotion recognition: A transfer learning approach," in *Proceedings of Grand Challenge and Workshop on Human Multimodal Language (Challenge-HML)*. Association for Computational Linguistics, 2018, pp. 48–52. [Online]. Available: <http://aclweb.org/anthology/W18-3307>
- [29] N. Tits, F. Wang, K. E. Haddad, V. Pagel, and T. Dutoit, "Visualization and interpretation of latent spaces for controlling expressive speech synthesis through audio analysis," *arXiv preprint arXiv:1903.11570*, 2019.
- [30] F. Eyben, K. R. Scherer, B. W. Schuller, J. Sundberg, E. André, C. Busso, L. Y. Devillers, J. Epps, P. Laukka, S. S. Narayanan *et al.*, "The geneva minimalistic acoustic parameter set (gemaps) for voice research and affective computing," *IEEE Transactions on Affective Computing*, vol. 7, no. 2, pp. 190–202, 2016.
- [31] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "Iemocap: Interactive emotional dyadic motion capture database," *Language resources and evaluation*, vol. 42, no. 4, p. 335, 2008.
- [32] N. Tits, K. E. Haddad, and T. Dutoit, "Exploring transfer learning for low resource emotional tts," *arXiv preprint arXiv:1901.04276*, 2019.
- [33] J. R. Orozco-Arroyave, J. Vdsquez-Correa, F. Höning, J. D. Arias-Londoño, J. Vargas-Bonilla, S. Skodda, J. 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.