

Pretrained audio neural networks for Speech emotion recognition in Portuguese

Marcelo Matheus Gaury¹, Marcelo Finger¹

¹Universidade de São Paulo, Rua do Matão 1010, São Paulo, Brazil

Abstract

The goal of speech emotion recognition (SER) is to identify the emotional aspects of speech. The SER challenge for Brazilian Portuguese speech was proposed with short snippets of Portuguese which are classified as neutral, non-neutral female and non-neutral male according to paralinguistic elements (laughing, crying, etc). This dataset contains about 50 minutes of Brazilian Portuguese speech. As the dataset leans on the small side, we investigate whether a combination of transfer learning and data augmentation techniques can produce positive results. Thus, by combining a data augmentation technique called SpecAugment, with the use of Pretrained Audio Neural Networks (PANNs) for transfer learning we are able to obtain interesting results. The PANNs (CNN6, CNN10 and CNN14) are pretrained on a large dataset called AudioSet containing more than 5000 hours of audio. They were finetuned on the SER dataset and the best performing model (CNN10) on the validation set was submitted to the challenge, achieving an $F1$ score of 0.73 up from 0.54 from the baselines provided by the challenge. Moreover, we also tested the use of Transformer neural architecture, pretrained on about 600 hours of Brazilian Portuguese audio data. Transformers, as well as more complex models of PANNs (CNN14), fail to generalize to the test set in the SER dataset and do not beat the baseline. Considering the limitation of the dataset sizes, currently the best approach for SER is using PANNs (specifically, CNN6 and CNN10).

Keywords

Speech emotion recognition, Pretrained audio neural networks, Transfer learning, Transformers

1. Introduction

Speech emotion recognition (SER) aims at identifying the emotional aspects of speech independently from the actual semantic content. SER can be used to identify the emotions of humans, e.g., when using mobile phones, an ability that may become crucial in improving human-machine interactions in the future [1]. Several efforts to acquire speech data classified with different emotional labels have been undertaken [2, 3, 4]. These datasets are typically small in size, even for languages such as English. In order to tackle these datasets, the use of transfer learning and data augmentation techniques may be instrumental.


Transfer learning is the method of training a network on a particular problem where there is an abundance of data, with the goal of using the acquired knowledge to obtain better performance on a related problem with limited data available. Transfer learning has been effectively used in many fields of deep learning such as computer vision [5] and language modelling [6]. Data augmentation is the method of increasing the amount of data available by slightly modifying

Proceedings of the First Workshop on Automatic Speech Recognition for Spontaneous and Prepared Speech Speech emotion recognition in Portuguese (SER 2022), co-located with PROPOR 2022. March 21st, 2022 (Online).

✉ marcelomatheusgaury@gmail.com (M. M. Gaury); mfinger@ime.usp.br (M. Finger)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

copies of the data. This can be done, for example, by masking parts of the input or by adding Gaussian noise to it.

In this paper, we use transfer learning and data augmentation techniques to study SER in Brazilian Portuguese speech. We participate in the *shared task SER challenge*, a challenge for Brazilian Portuguese speech emotion recognition. This challenge made available a labeled dataset of 625 audio files as training set for SER. Moreover, a dataset of 308 files was made available as the test set. The training and test datasets consisted of short snippets of Brazilian Portuguese speech, usually less than 15 s long, labeled neutral, non-neutral female and non-neutral male (non-neutral for audios containing laughs, cries, etc).

For transfer learning, we employ *Pretrained audio neural network (PANN)* [7], which are convolutional neural networks trained on a large dataset of audios (AudioSet [8]), consisting of 1.9 million audio clips distributed across 527 sound classes. By using the pretrained models made available by the developers, and finetuning on the SER dataset for Brazilian Portuguese speech, we are able to beat the proposed baselines of prosodic features and wav2vec features. We achieve (via CNN10) F1-score of 0.73, up from 0.54 from the baselines. During finetuning, we employ a data augmentation technique called SpecAugment [9].

We also tested the use of Transformer neural networks, pretrained on a large amount of Brazilian Portuguese audio data [10]. However, we find that, with the current amount of available data for SER, Transformers do not generalize their training performance to the validation and test sets. This holds even while using most common techniques to prevent overfitting. The same behaviour was also observed for more complex PANNs, such as CNN14.

2. Related Work

There is a large literature on SER in English [11, 12, 13, 14, 15, 16, 17, 18]. Moreover, there are a lot of small datasets for SER in English, such as RAVDESS [2], SAVEE [3] and IEMOCAP [4]. To the best of our knowledge, the SER dataset for Brazilian Portuguese speech is the only available dataset on the language. In addition, English datasets are usually classified in a different set of labels. RAVDESS [2], for example, has the classes of calm, happy, angry, sad, fearful, surprise and disgust. This contrasts with the classes of neutral, non-neutral female and non-neutral male present in the SER dataset for Brazilian Portuguese speech. As such, comparisons of our work with the state of the art in English language are not really possible. Nevertheless, the authors of [18], the most recent work, obtain an average recall on RAVDESS of 84.3 percent using wav2vec 2.0 [19]. On IEMOCAP, they obtain an average recall of 67.2 percent, also using wav2vec 2.0.

Transfer learning is a very common technique in situations where the dataset available is small in size. It has been effectively employed in computer vision [5, 20], language modelling [6, 21] and audio tasks [7, 22, 18]. In the original PANN paper [7], authors propose several convolutional neural networks pretrained on AudioSet which can be finetuned on other smaller datasets. In [18] the authors use wav2vec 2.0 pretrained on Librispeech and finetuned on either RAVDESS or IEMOCAP for speech emotion recognition. Finally, in [22] the authors provide a comprehensive review on transfer learning methods used for speech and language processing tasks.

3. Methodology

3.1. SER Dataset

To perform SER on Brazilian Portuguese speech, we use the training dataset (CORAA SER version 1.0) provided for the challenge. This dataset was built from the C-ORAL-BRASIL I corpus [23], with 625 audio files, typically less than 15s-long, containing informal spontaneous Brazilian Portuguese speech. These audio files are labeled neutral, non-neutral female, non-neutral male. An audio is labeled non-neutral male if it is a male speaker and it contains paralinguistic elements in the speech (such as laughing, crying, etc). Similarly, an audio is labeled non-neutral female if it is a female speaker and the speech contains such paralinguistic elements.

We split the official training dataset into training (80%), validation (10%) and test sets (10%). The split was done in an arbitrary way to ensure that the three datasets were balanced (i.e. contained relatively the same proportion of neutral, non-neutral female and non-neutral male files). The training dataset consisted of 500 files, the validation dataset consisted of 63 files and the test set of 62 files. The results we report are for the validation and test set performance.

As the official test dataset made available did not have labels, we have labeled it ourselves, out of curiosity and to enable more consistent tests of the performance of the networks. While the labels may not be perfect, they provide a close enough picture, so the performance of the models can be measured as an average over multiple experiments (as we were observing high variance). As such, we also provide results for the official test set with our unofficial labels. We stress that we did not use the test set labels for any form of model or parameter selection.

Lastly, the PANNs we use have been trained on the AudioSet [8] dataset containing more than 5000 hours of audio distributed across 527 classes.

3.2. PANN Architectures

Table 1 describes the three architectures we use. They are named CNN6, CNN10 and CNN14 after the 6-layer, 10-layer and 14-layer CNNs they represent. These are the same CNN network architectures used in [7]. We take their pretrained models on AudioSet [8] to allow us to obtain better generalization performances on the SER dataset.

The audios are preprocessed in the following way: the audios are first resampled to $32kHz$. After that, we apply short-time Fourier transform [24] (with a window size of 1024 frames and hop size of 320 frames) to the standard time-domain waveforms to obtain spectrograms. Then, Mel filter banks are applied to spectrograms, followed by a logarithm operation to obtain log Mel spectrograms. These preprocessing steps are commonly done when using CNNs for audio [25, 26].

As described in Table 1, the CNN architectures used are composed of convolutional layers with kernel 5×5 for CNN6, and 3×3 for CNN10 and CNN14. Each convolutional layer is followed by batch normalization [27] and ReLU non-linearity [28] is used to allow for better training convergence. Each such convolutional block is present 4 times in CNN6 and, in between, an average pooling 2×2 layer is applied (average pooling is observed to be better than max pooling [29]). In CNN10 and CNN14, the convolutional blocks are always used in pairs before an average pooling layer is applied. CNN10 contains 8 such convolutional blocks (4 pairs) and

Table 1

PANN architectures. We describe the layers of CNN6, CNN10 and CNN14.

CNN6	CNN10	CNN14
Log Mel Spectrogram n frames \times 64 mel bins		
$\begin{pmatrix} 5 \times 5 @ 64 \\ (BN, ReLU) \end{pmatrix}$	$\begin{pmatrix} 3 \times 3 @ 64 \\ (BN, ReLU) \end{pmatrix} \times 2$	$\begin{pmatrix} 3 \times 3 @ 64 \\ (BN, ReLU) \end{pmatrix} \times 2$
Avg Pooling 2×2	Avg Pooling 2×2	Avg Pooling 2×2
$\begin{pmatrix} 5 \times 5 @ 128 \\ (BN, ReLU) \end{pmatrix}$	$\begin{pmatrix} 3 \times 3 @ 128 \\ (BN, ReLU) \end{pmatrix} \times 2$	$\begin{pmatrix} 3 \times 3 @ 128 \\ (BN, ReLU) \end{pmatrix} \times 2$
Avg Pooling 2×2	Avg Pooling 2×2	Avg Pooling 2×2
$\begin{pmatrix} 5 \times 5 @ 256 \\ (BN, ReLU) \end{pmatrix}$	$\begin{pmatrix} 3 \times 3 @ 256 \\ (BN, ReLU) \end{pmatrix} \times 2$	$\begin{pmatrix} 3 \times 3 @ 256 \\ (BN, ReLU) \end{pmatrix} \times 2$
Avg Pooling 2×2	Avg Pooling 2×2	Avg Pooling 2×2
$\begin{pmatrix} 5 \times 5 @ 512 \\ (BN, ReLU) \end{pmatrix}$	$\begin{pmatrix} 3 \times 3 @ 512 \\ (BN, ReLU) \end{pmatrix} \times 2$	$\begin{pmatrix} 3 \times 3 @ 512 \\ (BN, ReLU) \end{pmatrix} \times 2$
Global Avg Pooling	Global Avg Pooling	Avg Pooling 2×2
FC 512, ReLU	FC 512, ReLU	$\begin{pmatrix} 3 \times 3 @ 1024 \\ (BN, ReLU) \end{pmatrix} \times 2$
FC 527, Sigmoid	FC 527, Sigmoid	Avg Pooling 2×2
		$\begin{pmatrix} 3 \times 3 @ 2048 \\ (BN, ReLU) \end{pmatrix} \times 2$
		Global Avg Pooling
		FC 2048, ReLU
		FC 527, Sigmoid

CNN14 contains 12 such convolutional blocks (6 pairs). All networks have a penultimate fully connected layer to add extra representation ability, as well as a final 527 units fully connected layer where a sigmoid is applied to obtain the probabilities for each class. In Table 1, the first line describes the input of the networks, that is, n frames of Log Mel Spectrogram with 64 mel bins for each frame. Each subsequent line represents a layer of the networks. The numbers following the @ sign represent the quantity of 5×5 or 3×3 feature maps used.

3.3. Transformer Encoder Architecture

In addition to experimenting with the PANNs, we also attempt to extract good performances from Transformers. The Transformer architecture we use is equivalent to the Transformer Encoder architecture from [30]. That is, we use a three-layer Transformer with multi-head self attention. Each encoder layer is composed of two sub-layers. The first is a multi-head self-attention network and the second is a fully connected feed-forward layer. Each sub-layer has a residual connection followed by layer normalization [31]. The encoder layers and sub-layers produce outputs of dimension d (in experiments d is either 128 or 512). The fully connected feed forward network within each encoder layer has an inner dimension of $4d$. We feed the Transformer Encoders the MFCC-gram of the audios, with each token fed to the Transformer corresponding to a frame of the MFCC-gram [32]. We name these Transformers, the MFCC-gram Transformers [32]. We use sinusoidal positional encoding so the Transformer has access to the order of the sequence fed [30, 33]. The input frames are projected linearly to a hidden layer of dimension d , as direct addition of acoustic features to positional encoding may lead to training failure [33].

Typically, Transformers undergo two training phases: pretraining and finetuning. In the pretraining phase, we make use of a technique called Time alteration [33] to pretrain the Transformer in about 600 hours of Brazilian Portuguese audio data (in other words, we use pretrained models from [10]). Time alteration is a technique that masks random spans of frames of the MFCC-gram similarly to how time masking functions in SpecAugment (described in subsection 3.4). During pretraining, the model is trained to reconstruct the masked frames. For Brazilian Portuguese audio data, we use the corpora of NURC-São Paulo[34], NURC-Recife [35], ALIP [36], SP2010 [37] and Programa Certas Palavras [38]. In the experiments, we also show the performance of Transformers which do not undergo pretraining, that is, which we initialize at random and do finetuning directly. We name those Transformers the Baseline MFCC-gram Transformers. After pretraining, the Transformers are finetuned on the SER dataset.

3.4. Data augmentation: SpecAugment

The SER training dataset used for the challenge leans on the small side and contains about 50 minutes of audio. To mitigate the potential overfitting effects of a small training dataset, we perform a common audio data augmentation technique called SpecAugment [9] on the Mel spectrogram (or MFCC-gram) of the audio files before feeding it to the network's layers. SpecAugment consists in masking random spans of consecutive segments of the spectrogram of the audios. Masking can be done along the time dimension (that is, on spans of consecutive frames), or along the frequency dimension (that is, on spans of consecutive frequency channels).

Following [7], time masking is done by selecting a uniform length ℓ (chosen between 0 and 64) and a uniform frame start t (chosen between 0 and $T - \ell$, where T is the total number of frames of the audio) and proceeding to mask the frames from t to $t + \ell - 1$. We mask two such blocks of consecutive frames. Frequency masking is similar to time masking but done along the frequency dimension. So, a random uniform length ℓ is chosen (between 0 and 8) and a uniform frequency band f is chosen (between 0 and $F - \ell$ where F is the total number of Mel frequency bins). The frequency bands from f to $f + \ell - 1$ are masked to zero. As with time masking, we mask two such blocks of consecutive frequency bands.

4. Results and Discussion

We will check the performance of the three proposed PANNs (CNN6, CNN10 and CNN14) on the SER training and test datasets. In order to take advantage of the large pretraining done on the AudioSet [8] dataset, we will use the pretrained models of CNN6, CNN10 and CNN14 made available by the authors of [7]. These can be found in [Zenodo](#). These pretrained models will be finetuned on the SER training dataset in order to achieve better performance than the baseline.

Moreover, to showcase the massive level of transfer learning that is happening via the pretrained models, we will show the performance of the three networks (CNN6, CNN10 and CNN14) without the use of a pretrained model, that is, initializing their weights at random and not making use of the AudioSet [8] pretraining. We call these three models the Baseline CNN6, the Baseline CNN10 and the Baseline CNN14.

Lastly, we show the performance of three Transformers models. We analyze MFCC-gram Transformers pretrained on about 600 hours of Brazilian Portuguese audio data, as well as,

Baseline MFCC-gram Transformers (without pretraining) containing 512 and 128 units per Encoder layer.

As mentioned before, the SER training dataset is split into a training (80%), validation (10%) and test sets (10%). In Table 2, we report the $F1$ score performance of the nine models in the validation and test datasets as well as in the official dataset (which was labeled by us). The results in the table are averaged across 25 experiments, to better control the generally high $F1$ score variance between different experiments. Each experiment consisted of training the model for 100 epochs for CNNs and 20 epochs for Transformers¹ in the training set and the best validation performance model (checked after each epoch) was saved and later analyzed on the test set and official test set. The batch size used was 16 and the learning rate was 10^{-4} for the CNNs and we use a warmup learning rate schedule according to the formula $d^{-0.5} \times \min(\text{stepnumber}^{-0.5}, \text{stepnumber} \times \text{warmupsteps}^{-1.5})$ for the Transformers as is standard [6]. We use $\text{warmupsteps} = 4000$.

As can be seen on Table 2, the best results in the test set were attained by the CNN6 (0.62 $F1$ score). Moreover, it seems that the test set built by us was inherently harder than the official test set. In the official test set, the best result was obtained by CNN10 (0.74 $F1$ score), in line with it achieving also the best results on the validation set.

We observe that CNN14’s performance was significantly worse both on validation and test. However, representation ability wise it is the most powerful of the PANN models. It is likely that the SER dataset being so small meant CNN14 suffered from overfitting.

We also experienced overfitting issues when attempting MFCC-gram Transformers based models. There, using pretraining techniques did not yield better performance. This is likely because the pretraining data contained primarily voice, without laughs or cries, so the important markers were not present in pretrained data. Moreover, no common technique (such as dropout [39], L1 or L2 regularization [40], data augmentation techniques as SpecAugment [9] and Mixup [41]) to prevent overfitting yielded good results. It seems that the reduced size of the SER dataset is currently hindering performance in more complex networks, so a likely way of dramatically improving results would be to increase the size of the available dataset.

Lastly, note that the three baseline PANN models are far away from beating the baselines provided by the challenge. There is noticeable transfer learning benefit in using the pretrained models on AudioSet [8]. This large difference illustrates again the fact that the SER dataset is so small (50 minutes of audio) and that these networks suffer to generalize on it.

We have sent for evaluation in the challenge, the model which attained best test performance (a CNN6 which officially reported 0.66 $F1$ -score) and the model which attained best validation performance (a CNN10 which officially reported 0.73 $F1$ -score). Moreover, out of curiosity, we show the confusion matrix of the CNN10 model sent for evaluation in Table 3. Observe that the model classifies the vast majority of neutral and non-neutral females files correctly. Most of the errors are done classifying non-neutral male files (often wrongly classified as neutral).

¹As Transformer does not generalize, no advantage exists in training it for longer than 20 epochs.

Table 2

The mean and standard deviation of the $F1$ score is shown in the table below for the nine models (CNN6, CNN10 and CNN14 and their respective baseline version, i.e., their versions without pretraining on AudioSet [8], as well as MFCC-gram Transformers with and without pretraining and a smaller version of MFCC-gram Transformers). The results shown are for the validation set, the test set and the official test set. Labels for the official test set were created by us.

Model	$F1$ score Validation performance	$F1$ score Test performance	$F1$ score Official test performance
Baseline CNN6	0.45 ± 0.06	0.36 ± 0.05	0.33 ± 0.03
Baseline CNN10	0.58 ± 0.06	0.41 ± 0.09	0.42 ± 0.05
Baseline CNN14	0.38 ± 0.06	0.33 ± 0.04	0.32 ± 0.03
CNN6	0.78 ± 0.05	0.62 ± 0.06	0.69 ± 0.04
CNN10	0.80 ± 0.06	0.57 ± 0.06	0.74 ± 0.04
CNN14	0.61 ± 0.11	0.54 ± 0.06	0.52 ± 0.10
MFCC-gram Transformers 512 units	0.50 ± 0.04	0.36 ± 0.06	0.38 ± 0.03
Baseline MFCC-gram Transformers 512 units	0.57 ± 0.04	0.43 ± 0.08	0.43 ± 0.06
Baseline MFCC-gram Transformers 128 units	0.60 ± 0.05	0.45 ± 0.07	0.44 ± 0.04

Table 3

We plot the confusion matrix for the CNN10 model which was submitted to the challenge and attained an $F1$ score of 0.73 (on official labels). Note that the model has the most difficulty classifying non-neutral male files correctly.

Confusion Matrix	predicted neutral	predicted non-neutral male	predicted non-neutral female
Neutral	244	2	5
Non-neutral male	14	8	2
Non-neutral female	6	1	26

5. Conclusion

In this paper, we have effectively used transfer learning to beat the proposed baselines in the shared task SER challenge in Brazilian Portuguese speech. By using, the PANNs CNN6 and CNN10, we have attained $F1$ score of 0.73 up from 0.54 from the baselines. We have also

observed that more complex networks, such as CNN14 and Transformers, while being in theory more capable of attaining better performances, suffer from overfitting. As such, we determine that probably the best way of improving results is by increasing the size of the training set.

Future work could involve increasing the size of the training set so that Transformers and CNN14 generalize their training performances to the test set. In addition, pretraining Transformers with audio data containing specifically laughs, cries and so on may prove useful. Moreover, other data augmentation techniques could be used which might provide additional benefit in terms of preventing overfitting.

Acknowledgments

This work was supported by FAPESP grant number 2020/16543-7 (POSDOC) and project 06443-5 (SPIRA). MF was partly supported by CNPq grant PQ 303609/2018-4, Fapesp 2014/12236-1 (Animals) and 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. This work was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001.

References

- [1] E. Andre, M. Rehm, W. Minker, D. Bühler, Endowing spoken language dialogue systems with emotional intelligence, in: *Tutorial and Research Workshop on Affective Dialogue Systems*, Springer, 2004, pp. 178–187.
- [2] S. R. Livingstone, 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* 13 (2018) e0196391.
- [3] W. Wang, *Machine Audition: Principles, Algorithms and Systems: Principles, Algorithms and Systems*, IGI Global, 2010.
- [4] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, S. S. Narayanan, Iemocap: Interactive emotional dyadic motion capture database, *Language resources and evaluation* 42 (2008) 335–359.
- [5] A. Voulodimos, N. Doulamis, A. Doulamis, E. Protopapadakis, Deep learning for computer vision: A brief review, *Computational intelligence and neuroscience* 2018 (2018).
- [6] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, *arXiv preprint arXiv:1810.04805* (2018).
- [7] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, M. D. Plumbley, Panns: Large-scale pretrained audio neural networks for audio pattern recognition, *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 28 (2020) 2880–2894.
- [8] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, M. Ritter, Audio set: An ontology and human-labeled dataset for audio events, in: *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2017, pp. 776–780.

- [9] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, Q. V. Le, Specaugment: A simple data augmentation method for automatic speech recognition, arXiv preprint arXiv:1904.08779 (2019).
- [10] M. Gauy, M. Finger, Acoustic models for brazilian portuguese speech based on neural transformers, IN PREPARATION (2022).
- [11] M. Lech, M. Stolar, C. Best, R. Bolia, Real-time speech emotion recognition using a pre-trained image classification network: Effects of bandwidth reduction and companding, *Frontiers in Computer Science* 2 (2020) 14.
- [12] S. Yoon, S. Byun, S. Dey, K. Jung, Speech emotion recognition using multi-hop attention mechanism, in: ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, 2019, pp. 2822–2826.
- [13] H. Xu, H. Zhang, K. Han, Y. Wang, Y. Peng, X. Li, Learning alignment for multimodal emotion recognition from speech, arXiv preprint arXiv:1909.05645 (2019).
- [14] S. Yoon, S. Byun, K. Jung, Multimodal speech emotion recognition using audio and text, in: 2018 IEEE Spoken Language Technology Workshop (SLT), IEEE, 2018, pp. 112–118.
- [15] A. Satt, S. Rozenberg, R. Hoory, Efficient emotion recognition from speech using deep learning on spectrograms., in: Interspeech, 2017, pp. 1089–1093.
- [16] D. Issa, M. F. Demirci, A. Yazici, Speech emotion recognition with deep convolutional neural networks, *Biomedical Signal Processing and Control* 59 (2020) 101894.
- [17] Z. Peng, Y. Lu, S. Pan, Y. Liu, Efficient speech emotion recognition using multi-scale cnn and attention, in: ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, 2021, pp. 3020–3024.
- [18] L. Pepino, P. Riera, L. Ferrer, Emotion recognition from speech using wav2vec 2.0 embeddings, arXiv preprint arXiv:2104.03502 (2021).
- [19] A. Baevski, Y. Zhou, A. Mohamed, M. Auli, wav2vec 2.0: A framework for self-supervised learning of speech representations, *Advances in Neural Information Processing Systems* 33 (2020) 12449–12460.
- [20] A. Krizhevsky, I. Sutskever, G. E. Hinton, Imagenet classification with deep convolutional neural networks, *Advances in neural information processing systems* 25 (2012) 1097–1105.
- [21] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, arXiv preprint arXiv:2005.14165 (2020).
- [22] D. Wang, T. F. Zheng, Transfer learning for speech and language processing, in: 2015 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA), IEEE, 2015, pp. 1225–1237.
- [23] T. Raso, H. Mello, The c-oral-brasil i: reference corpus for informal spoken brazilian portuguese, in: *International Conference on Computational Processing of the Portuguese Language*, Springer, 2012, pp. 362–367.
- [24] E. O. Brigham, R. Morrow, The fast fourier transform, *IEEE spectrum* 4 (1967) 63–70.
- [25] K. Choi, G. Fazekas, M. Sandler, Automatic tagging using deep convolutional neural networks, arXiv preprint arXiv:1606.00298 (2016).
- [26] Q. Kong, C. Yu, Y. Xu, T. Iqbal, W. Wang, M. D. Plumbley, Weakly labelled audioset tagging with attention neural networks, *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 27 (2019) 1791–1802.

- [27] S. Ioffe, C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shift, in: International conference on machine learning, PMLR, 2015, pp. 448–456.
- [28] V. Nair, G. E. Hinton, Rectified linear units improve restricted boltzmann machines, in: Icm1, 2010.
- [29] Q. Kong, Y. Cao, T. Iqbal, Y. Xu, W. Wang, M. D. Plumbley, Cross-task learning for audio tagging, sound event detection and spatial localization: Dcase 2019 baseline systems, arXiv preprint arXiv:1904.03476 (2019).
- [30] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, Advances in neural information processing systems 30 (2017) 5998–6008.
- [31] J. L. Ba, J. R. Kiros, G. E. Hinton, Layer normalization, arXiv preprint arXiv:1607.06450 (2016).
- [32] M. M. Gauy, M. Finger, Audio mfcc-gram transformers for respiratory insufficiency detection in covid-19, in: STIL 2021 (), 2021. URL: <http://XXXXXX/219270.pdf>.
- [33] A. T. Liu, S.-w. Yang, P.-H. Chi, P.-c. Hsu, H.-y. Lee, Mockingjay: Unsupervised speech representation learning with deep bidirectional transformer encoders, in: ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, 2020, pp. 6419–6423.
- [34] A. T. d. Castilho, D. Pretti, A linguagem falada culta na cidade de são paulo: materiais para seu estudo., 1986.
- [35] M. Oliviera Jr, et al., Nurc digital um protocolo para a digitalização, anotação, arquivamento e disseminação do material do projeto da norma urbana linguística culta (nurc), CHIMERA: Revista de Corpus de Lenguas Romances y Estudios Lingüísticos 3 (2016) 149–174.
- [36] S. C. L. Gonçalves, Projeto alip (amostra linguística do interior paulista) e banco de dados iboruna: 10 anos de contribuição com a descrição do português brasileiro, Estudos Linguísticos (São Paulo. 1978) 48 (2019) 276–297.
- [37] R. B. Mendes, Projeto sp2010: Amostra da fala paulistana, <http://projetosp2010.fflch.usp.br>. Acesso em 1 (2013) 2013.
- [38] C. S. P. Teixeira, Acervo Certas Palavras- Catálogo 1981-1996., Unicamp Cedae, 1997.
- [39] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: a simple way to prevent neural networks from overfitting, The journal of machine learning research 15 (2014) 1929–1958.
- [40] I. Goodfellow, Y. Bengio, A. Courville, Deep learning, MIT press, 2016.
- [41] H. Zhang, M. Cisse, Y. N. Dauphin, D. Lopez-Paz, mixup: Beyond empirical risk minimization, arXiv preprint arXiv:1710.09412 (2017).