
RECOGNIZING MORE EMOTIONS WITH LESS DATA USING SELF-SUPERVISED TRANSFER LEARNING

A PREPRINT

Jonathan Boigne
Wonder Technology
Beijing, China
jonathan@wondertech.com.cn

Biman Najika Liyanage
Wonder Technology
Beijing, China
biman@wondertech.com.cn

Ted Östrem
Wonder Technology
Beijing, China
ted@wondertech.com.cn

August 28, 2020

ABSTRACT

We propose a novel transfer learning method for speech emotion recognition allowing us to obtain promising results when only few training data is available. With as low as 125 examples per emotion class, we were able to reach a higher accuracy than a strong baseline trained on 8 times more data. Our method leverages knowledge contained in pre-trained speech representations extracted from models trained on a more general self-supervised task which doesn't require human annotations, such as the wav2vec model. We provide detailed insights on the benefits of our approach by varying the training data size, which can help labeling teams to work more efficiently. We compare performance with other popular methods on the IEMOCAP dataset, a well-benchmarked dataset among the Speech Emotion Recognition (SER) research community.

Furthermore, we demonstrate that results can be greatly improved by combining acoustic and linguistic knowledge from transfer learning. We align acoustic pre-trained representations with semantic representations from the BERT model through an attention-based recurrent neural network. Performance improves significantly when combining both modalities and scales with the amount of data. When trained on the full IEMOCAP dataset, we reach a new state-of-the-art of 73.9% unweighted accuracy (UA).

Keywords Emotion AI · Speech Emotion Recognition · Transfer Learning · Self-Supervised Learning · Low-Resource

1 Introduction

Emotion recognition has been gaining traction for the last decade with the interest of providing conversational agents with high EQ when interacting with users. Such systems have applications in healthcare, non-invasive mental health diagnostics and screening, automotive, as well as education. Recognizing emotions has been a challenging task in the speech domain, mainly due to the lack of large-enough labeled datasets to successfully apply proven deep learning techniques, like it has been done in tasks with more resources such as Automatic Speech Recognition (ASR) [7, 3]. This is even more problematic for non-English languages with fewer resources. It is thus important to find alternative ways to train accurate classifiers in situations where data is scarce. Moreover, emotion recognition is a data problem where having a finer emotion classification system leads to datasets of uneven sample sizes or data with multiple labels.

There have been various lines of work attempting to train accurate emotion classifiers with a low amount of labeled data. Some works have proposed to impose stronger restrictions on convolutional layers to better fit raw speech data and prevent overfitting on small data [14]. Multi-task learning has also been proposed as a way to mitigate overfitting to a small dataset by simultaneously learning different tasks [21, 24]. However, these approaches still require a large enough amount of data to be able to learn efficient filters from scratch or to train all adjacent tasks.

Transfer learning is a growing area of research in deep learning and has the potential to help alleviate this problem of label scarcity. Particularly re-using pre-trained models or representations trained on more general tasks has been successfully applied in other domains than speech [5]. In computer vision, it is now standard practice to train a deep

convolutional model first on the large ImageNet classification dataset and then re-use those weights for fine-tuning on a different vision task where less labeled data is available [9]. In natural language processing, the pre-training of large language models has been widely adopted and led to greatly superior results on many tasks [13, 4].

Despite the clear benefits of pre-training in computer vision and natural language processing, this approach has not yet taken off in the speech domain, mainly because it is still unclear which task is the most suitable for pre-training. Yet recent works have started to show promising results, notably the wav2vec model [17] which helped reach a new state-of-the-art for speech recognition by training rich representations from unlabeled data through a contrastive-type of loss. More specifically, for the task of emotion recognition, [10] obtained strong results by re-using representations from a pre-trained ASR model. However, this approach still requires a large amount of labeled data to first train a good speech transcription model, which makes it impracticable for languages or domains where data is scarce.

In this work, we focus on unsupervised pre-training, which does not require labeled data for learning pre-trained representations, and as such we decide to re-use the wav2vec representations to apply for emotion recognition. We also experiment with combining two different kinds of pre-trained representations for both speech and text. The addition of BERT representations allows us to improve our accuracy by almost 10%.

2 Approach

2.1 Speech Emotion Recognition

The current standard approach to classify emotions from speech is to first extract low-level features, often called low-level descriptors (LLDs) from short frames of speech of duration ranging from 20 to 50 milliseconds, denoted $\mathbf{x}_{1:T} = (x_1, x_2, \dots, x_T)$, where $x_i \in R^n$, n being the number of features. A high-level aggregation transformation is then applied to convert these frame-level features to an utterance-level representation $x_{utterance} \in R^n$. Finally, a softmax layer is applied on top of this new global representation to classify the utterance along the possible emotion classes.

Different methods have been used to transform the variable-length sequence of low-level descriptors to a fixed-length representation at the utterance level. Traditionally, statistical aggregation functions were applied to each of the LLDs over the duration of the utterance such as mean, max, variance, etc. These were concatenated into a long feature-vector to represent a single utterance. With the emergence of neural networks as the preferred approach, there have been experiments with various strategies such as recurrent neural networks or attention-based models [11].

2.1.1 Acoustic features

Traditionally most works use hand-crafted features that have been proven to work well for other speech-related tasks. Commonly used features include Mel-frequency cepstrum coefficients (MFCCs), zero-crossing rate, energy, pitch, voicing probability, etc. Another line of work has been focused on using raw spectrogram magnitudes, log-Mel spectrograms, or even raw waveform [18, 15]. Yet although this works well for tasks where data is sufficient like Automatic Speech Recognition, these methods are still difficult to apply for emotion recognition due to a lack of labeled data.

We suggest that working with pre-trained features trained on raw waveform or raw spectrograms could help alleviate this lack of data, and we suggest using wav2vec features that were trained on large data in a self-supervised fashion.

2.2 Pre-trained Wav2vec Representations

The wav2vec [17] model is made of two simple convolutional neural networks: the encoder network and the context network. The encoder network $f : X \mapsto Z$ takes raw audio samples $x_i \in X$ as input and outputs low frequency feature representations (z_1, z_2, \dots, z_T) which encode about 30ms of 16kHz audio every 10ms. The context network $g : Z \mapsto C$ transforms these low-frequency representations into a higher-level contextual representation $c_i = g(z_i \dots z_{i-v})$ for a receptive field v . The total receptive field after passing through both networks is 210ms for the base version and 810ms for the large version.

The model was trained on about 1,000 hours of unlabeled English speech with a noise contrastive binary classification task, where the objective was to distinguish true future samples from distractors. The motivation behind wav2vec was to learn effective pre-trained representations to be used for Automatic Speech Recognition (ASR). Training a model on the resulting representations allowed the team to outperform Deep Speech 2 [1], the best reported character-based system at the time while using two orders of magnitude less data.




Figure 1: Architectures used for comparing emotion recognition performance using hand-engineered features or pre-trained audio representations. (a) Mean pooling in time. (b) Mean+Max pooling. (c) Attention-weighted pooling. (d) MLP with mean pooling.

Our assumption is that these representations not only contain relevant information for speech recognition but also para-acoustic information that could be useful for detecting emotions in speech.

2.3 Model Architectures

We compare our approach with a baseline of commonly-used hand-engineered features to see how much more emotional information is contained in the pre-trained representations. To better understand the differences between both kinds of features, we experiment with various model architectures described below and in Figure 1. Due to our experiments on a very small amount of data, we selected simpler models rather than more complex recurrent neural networks to avoid overfitting. However, we do compare our solution to more advanced baselines trained on more data in Section 4.

- (a) **Mean pooling:** this approach simply averages each feature along the time dimension. We make the simplifying assumption that the emotional information is constant-enough along time to meaningfully represent the utterance-level overall emotion. We use short-enough audio segments to make sure this is the case.
- (b) **Mean-max pooling:** similar to mean pooling except that we leverage both the mean and the max of features across the time dimension. We then concatenate both vectors before feeding it to the softmax layer. This approach effectively doubles the number of parameters of our model from just mean pooling.
- (c) **Attention pooling:** the model learns by itself a weighted average whose weights are determined by a simple attention mechanism based on logistic regression. This allows the model to learn the most efficient pooling function for the task, including the mean pooling approach above.
- (d) **MLP with pooling:** this last approach allows the model to learn more complex frame-level features before the pooling operation. We build a Multi-Layer Perceptron (MLP) on top of individual frame-level features x_1, x_2, \dots, x_T before mean pooling across the time dimension. The MLP parameters are shared across time. We scale the number of layers and units to match the model capacity for the two types of features in order to compensate for the smaller dimension of LLD feature vectors compared to wav2vec representations.

2.4 Bimodal Emotion Recognition

Several works have explored combining information from linguistic features and acoustic features. Yoon et al. [23], Heusser et al. [8] combined utterance-level information from audio and textual embeddings before classifying through the last softmax layer. Lu et al. [10] also used a similar approach to ours except that they used pre-trained representations from an ASR model, which thus already contains semantic information.

In this work, we decide to use the same model as in [22], where we align both audio and textual pre-trained representations through an attention mechanism on top of a bidirectional recurrent neural network. The only difference is the replacement of hand-engineered features by wav2vec embeddings and of textual GloVe embeddings [12] by BERT embeddings. BERT [4] is a Transformer [19] encoder-only model which has been shown to capture meaningful information for text classification tasks. Specifically, it is able to encode word embeddings which take the context into account.

3 Experimental Setup

To compare emotion recognition performance between traditional features and embeddings from self-supervised pre-trained models, we use the IEMOCAP dataset [2], which contains 12 hours of audiovisual data, including video,

Table 1: Unweighted accuracy when training on only 500 examples using different features and model architectures

Model	LLD Features	Pre-trained Representations	Improvement
Mean pooling	49.7%	54.9%	+5.2%
Mean+Max pooling	49.9%	57.4%	+7.5%
Attention pooling	49.7%	55.3%	+5.6%
MLP with pooling	51.6%	58.5%	+6.9%

speech, motion capture of face, text transcriptions. We use scripted and improvised dialogs from all 5 sessions. We only use the audio (and transcriptions for the bi-modal experiment) for our evaluation.

Similarly to previous works on IEMOCAP, we only use 4 emotion classes (neutral, happy, sad, and angry) where at least 2 annotators agree. We also merge excitement together with happiness, resulting in a total of 5,531 utterances for the full dataset. We evaluate our model using speaker-independent 5-fold cross-validation: for each split, 4 sessions are used for training, and the remaining one is used for testing. We report our model’s accuracy as the average of the accuracy on the validation set over all folds. This is consistent with previous works.

We use the pyAudioAnalysis library [6] to extract a set of 34 commonly-used features as in [22]. For wav2vec representations, we use the large version of the model since they have a wider receptive field which is more likely to contain emotion-salient features. For both audio representations, we crop all utterances to a maximum duration of 5 seconds. For BERT embeddings, we use the base model and extract embeddings from the provided IEMOCAP transcription files with the HuggingFace library [20]. We take embeddings from the second-to-last layer since these contain more general information and are less tied to BERT’s particular training objective.

We train our models with a small batch size of 16 and we use the Adam optimizer with learning rate 10^{-4} . We don’t use any form of early stopping. We add dropout regularization for all models after each layer.

4 Results

4.1 Training on Few Data

We compare generalization performance in a setting where training data is very scarce. We use only 500 examples for training, 125 per class. By using only a simple mean pooling and a softmax layer, we were able to reach 56.7% unweighted accuracy (UA) thanks to the wav2vec representations. This is 5.2% more than the same model trained on standard hand-engineered features. Table 1 report our results across different models and compare them with hand-engineered features. Our best model was able to reach an unweighted accuracy of 58.5%, higher than the same model trained on 8 times more data using hand-engineered features.

Interestingly, our results almost match a more advanced Bi-RNN model [11], trained on the entire dataset, that allows each frame to attend to the entire utterance (see Table 2). In comparison, wav2vec representations only have a receptive field of about 810ms, less than 1/6 of the maximum utterance length of 5 seconds.

4.2 Performance Scaling

In a second phase, we progressively increase the amount of training data on our best model (MLP with pooling). Figure 2 shows the evolution of the unweighted accuracy performance as we increase the amount of data. The UA increases in a log fashion with respect to the number of training examples.

Starting from 2,000 examples, our approach strongly outperforms other acoustic-only baselines that we benchmark against (see Table 2). The models we compare with were not only trained on more than twice the data, but they also present much more complex architectures compared to our simple MLP with pooling.

4.3 Bi-Modal Transfer Learning

Last, we experiment with combining pre-trained embeddings for both audio and text. We align wav2vec representations and sub-words embeddings from BERT through an attention-based recurrent neural network to align both representations in time, similar to [22]. The resulting model is much larger than previous ones, and to avoid over-fitting we only train it on the full dataset. We report the unweighted accuracy in Table 2, where our approach outperforms other models and reach a new state-of-the-art unweighted accuracy of 73.9%.




Figure 2: Evolution of performances for our best model (MLP with mean pooling) in function of the amount of training data

Table 2: Performance of our acoustic-only and bi-modal models when scaled on the full IEMOCAP dataset. We report results from previous state-of-the-art and often-cited works for comparison.

Model	Features	UA
Acoustic only		
Bi-LSTM with attention pooling [11]	Low-level descriptors	58.8%
CNN with Bi-LSTM [16]	Raw spectrograms	59.4%
TDNN with LSTM and attention [15]	Raw waveform	60.7%
MLP with mean pooling	Pre-trained wav2vec representations	64.3%
Acoustic and Textual		
Bi-LSTM with attention alignment [22]	LLDs and GloVe embeddings	70.9%
Bi-LSTM with multi-head self-attention [10]	pre-trained ASR representations	72.6%
Bi-LSTM with attention alignment	Pre-trained wav2vec and BERT representations	73.9%

5 Conclusion

In this paper, we compare performance for emotion recognition when using pre-trained embeddings learned in a self-supervised setting. We demonstrate the superior performance and sample-efficiency of our technique compared to identical models with commonly-used hand-engineered features. Our model is able to reach a higher accuracy with 8 times less data than if it was trained from scratch in a supervised setting.

We report performance as we scale up the training data, and build a final model on two modalities: audio and text. Both modalities use pre-trained features from self-supervised models, and we reach a new state-of-the-art unweighted accuracy of 73.9%.

References

- [1] Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, Jie Chen, Jingdong Chen, Zhijie Chen, Mike Chrzanowski, Adam Coates, Greg Diamos, Ke Ding, Niandong Du, Erich Elsen, Jesse Engel, Weiwei Fang, Linxi Fan, Christopher Fougner, Liang Gao, Caixia Gong, Awni Hannun, Tony Han, Lappi Johannes, Bing Jiang, Cai Ju, Billy Jun, Patrick LeGresley, Libby Lin, Junjie Liu, Yang Liu, Weigao Li, Xiangang Li, Dongpeng Ma, Sharan Narang, Andrew Ng, Sherjil Ozair, Yiping Peng, Ryan Prenger, Sheng Qian, Zongfeng Quan, Jonathan Raiman, Vinay Rao, Sanjeev Satheesh, David Seetapun, Shubho Sengupta, Kavya Srinet, Anuroop Sriram, Haiyuan Tang, Liliang Tang, Chong Wang, Jidong Wang, Kaifu Wang, Yi Wang, Zhijian Wang, Zhiqian Wang, Shuang Wu, Likai Wei, Bo Xiao, Wen Xie, Yan Xie, Dani Yogatama, Bin Yuan, Jun Zhan, and Zhenyao Zhu. Deep Speech 2 : End-to-End Speech Recognition in English and Mandarin. In *International Conference on Machine Learning*, pages 173–182, June 2016. URL <http://proceedings.mlr.press/v48/amodei16.html>.
- [2] Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette N. Chang, Sungbok Lee, and Shrikanth S. Narayanan. IEMOCAP: interactive emotional dyadic motion capture database. *Language Resources and Evaluation*, 42(4):335–359, December 2008. ISSN 1574-020X, 1574-0218. doi: 10.1007/s10579-008-9076-6. URL <http://link.springer.com/10.1007/s10579-008-9076-6>.
- [3] Ronan Collobert, Christian Puhersch, and Gabriel Synnaeve. Wav2Letter: an End-to-End ConvNet-based Speech Recognition System. *arXiv:1609.03193 [cs]*, September 2016. URL <http://arxiv.org/abs/1609.03193>.
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://www.aclweb.org/anthology/N19-1423>.
- [5] Dumitru Erhan, Aaron Courville, Yoshua Bengio, and Pascal Vincent. Why Does Unsupervised Pre-training Help Deep Learning? In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 201–208. JMLR Workshop and Conference Proceedings, March 2010. URL <http://proceedings.mlr.press/v9/erhan10a.html>.
- [6] Theodoros Giannakopoulos. pyAudioAnalysis: An Open-Source Python Library for Audio Signal Analysis. *PLOS ONE*, 10(12):e0144610, December 2015. ISSN 1932-6203. doi: 10.1371/journal.pone.0144610. URL <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0144610>.
- [7] Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, and Andrew Y. Ng. Deep Speech: Scaling up end-to-end speech recognition. *arXiv:1412.5567 [cs]*, December 2014. URL <http://arxiv.org/abs/1412.5567>.
- [8] Verena Heusser, Niklas Freymuth, Stefan Constantin, and Alex Waibel. Bimodal Speech Emotion Recognition Using Pre-Trained Language Models. *arXiv:1912.02610 [cs, eess, stat]*, November 2019. URL <http://arxiv.org/abs/1912.02610>.
- [9] Minyoung Huh, Pulkit Agrawal, and Alexei A. Efros. What makes ImageNet good for transfer learning? *arXiv:1608.08614 [cs]*, December 2016. URL <http://arxiv.org/abs/1608.08614>.
- [10] Zhiyun Lu, Liangliang Cao, Yu Zhang, Chung-Cheng Chiu, and James Fan. Speech Sentiment Analysis via Pre-Trained Features from End-to-End ASR Models. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7149–7153, May 2020. doi: 10.1109/ICASSP40776.2020.9052937.
- [11] Seyedmahdad Mirsamadi, Emad Barsoum, and Cha Zhang. Automatic speech emotion recognition using recurrent neural networks with local attention. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2227–2231, New Orleans, LA, March 2017. IEEE. ISBN 978-1-5090-4117-6. doi: 10.1109/ICASSP.2017.7952552. URL <http://ieeexplore.ieee.org/document/7952552/>.
- [12] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, 2014. Association for Computational Linguistics. doi: 10.3115/v1/D14-1162. URL <http://aclweb.org/anthology/D14-1162>.
- [13] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving Language Understanding by Generative Pre-Training. 2018.
- [14] Mirco Ravanelli and Yoshua Bengio. Speaker Recognition from Raw Waveform with SincNet. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 1021–1028, December 2018. doi: 10.1109/SLT.2018.8639585.

- [15] Mousmita Sarma, Pegah Ghahremani, Daniel Povey, Nagendra Kumar Goel, Kandarpa Kumar Sarma, and Najim Dehak. Emotion Identification from Raw Speech Signals Using DNNs. In *Interspeech 2018*, pages 3097–3101. ISCA, September 2018. doi: 10.21437/Interspeech.2018-1353. URL http://www.isca-speech.org/archive/Interspeech_2018/abstracts/1353.html.
- [16] Aharon Satt, Shai Rozenberg, and Ron Hoory. Efficient Emotion Recognition from Speech Using Deep Learning on Spectrograms. In *Interspeech 2017*, pages 1089–1093. ISCA, August 2017. doi: 10.21437/Interspeech.2017-200. URL http://www.isca-speech.org/archive/Interspeech_2017/abstracts/0200.html.
- [17] Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli. wav2vec: Unsupervised Pre-Training for Speech Recognition. In *Interspeech 2019*, pages 3465–3469. ISCA, September 2019. doi: 10.21437/Interspeech.2019-1873. URL http://www.isca-speech.org/archive/Interspeech_2019/abstracts/1873.html.
- [18] George Trigeorgis, Fabien Ringeval, Raymond Brueckner, Erik Marchi, Mihalis A. Nicolaou, Bjorn Schuller, and Stefanos Zafeiriou. Adieu features? End-to-end speech emotion recognition using a deep convolutional recurrent network. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5200–5204, Shanghai, March 2016. IEEE. ISBN 978-1-4799-9988-0. doi: 10.1109/ICASSP.2016.7472669. URL <http://ieeexplore.ieee.org/document/7472669/>.
- [19] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf>.
- [20] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. HuggingFace’s Transformers: State-of-the-art Natural Language Processing. *arXiv e-prints*, 1910:arXiv:1910.03771, October 2019. URL <http://adsabs.harvard.edu/abs/2019arXiv191003771W>.
- [21] Rui Xia and Yang Liu. A Multi-Task Learning Framework for Emotion Recognition Using 2D Continuous Space. *IEEE Transactions on Affective Computing*, 8(1):3–14, January 2017. ISSN 1949-3045. doi: 10.1109/TAFFC.2015.2512598.
- [22] Haiyang Xu, Hui Zhang, Kun Han, Yun Wang, Yiping Peng, and Xiangang Li. Learning Alignment for Multimodal Emotion Recognition from Speech. In *Interspeech 2019*, pages 3569–3573. ISCA, September 2019. doi: 10.21437/Interspeech.2019-3247. URL http://www.isca-speech.org/archive/Interspeech_2019/abstracts/3247.html.
- [23] Seunghyun Yoon, Seokhyun Byun, and Kyomin Jung. Multimodal Speech Emotion Recognition Using Audio and Text. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 112–118, December 2018. doi: 10.1109/SLT.2018.8639583.
- [24] Zixing Zhang, Bingwen Wu, and Björn Schuller. Attention-augmented End-to-end Multi-task Learning for Emotion Prediction from Speech. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6705–6709, May 2019. doi: 10.1109/ICASSP.2019.8682896.