



Unsupervised Representation Learning with Future Observation Prediction for Speech Emotion Recognition

Zheng Lian^{1,3}, Jianhua Tao^{1,2,3}, Bin Liu¹, Jian Huang^{1,3}

¹National Laboratory of Pattern Recognition, CASIA, Beijing, China

²CAS Center for Excellence in Brain Science and Intelligence Technology, Beijing, China

³School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China

{zheng.lian, jhtao, liubin, jian.huang}@nlpr.ia.ac.cn

Abstract

Prior works on speech emotion recognition utilize various unsupervised learning approaches to deal with low-resource samples. However, these methods pay less attention to modeling the long-term dynamic dependency, which is important for speech emotion recognition. To deal with this problem, this paper combines the unsupervised representation learning strategy – Future Observation Prediction (FOP), with transfer learning approaches (such as Fine-tuning and Hypercolumns). To verify the effectiveness of the proposed method, we conduct experiments on the IEMOCAP database. Experimental results demonstrate that our method is superior to currently advanced unsupervised learning strategies.

Index Terms: speech emotion recognition, unsupervised learning, Future Observation Prediction, transfer learning

1. Introduction

Obtaining large amounts of realistic data is currently challenging and expensive in speech emotion recognition [1]. Compared with automatic speech recognition which has thousands of hours of transcribed speech, databases annotated in emotional categories are still scarce. For example, the eINTERFACE [2] and EMODB [3] databases contain less than 1000 samples. Although the IEMOCAP [4] and FAU-AIBO [5] databases contain 10039 and 13348 samples respectively, the distribution of emotional categories is extremely unbalanced.

Prior works on speech emotion recognition utilize various methods to deal with low-resource training samples, including unsupervised representation learning [6, 7] and transfer learning [8, 9]. Unsupervised representation learning takes full advantage of the information from unlabeled data. It does not utilize the label information but aims to learn robust representations that can capture intrinsic structures of the data. While transfer learning makes use of additional labeled data from a different but related task. Its main idea is to share the “knowledge” from the source task to the target task.

Most of existing unsupervised learning approaches provide salient representations, leading to notable improvement for speech emotion recognition [6, 7, 10, 11]. A mainstream approach of unsupervised learning is to train autoencoders (AE) [12, 13], including denoising autoencoders (DAE) [14], variational autoencoders (VAE) [15] and autoencoders with ladder structures [16]. These methods [12, 13] take the whole input into account, aiming to learn intermediate feature representations that can reconstruct the input. However, they pay less attention to modeling the long-term dynamic dependency, which is important for speech emotion recognition [17, 18].

To deal with this problem, this paper focuses on another

popular unsupervised learning strategy, whose main idea is to predict future, missing or contextual information [19, 20, 21, 22]. For convenience, it is marked as Future Observation Prediction (FOP). As FOP needs to infer the global information in the prediction process, capturing long-range dynamic information is also important [19, 20]. The main intuition behind FOP is to learn the representations that encode the underlying shared information between different parts of the signal. In time series, FOP uses next step prediction to exploit the local smoothness of the signal, and discards low-level information and noise that are local. When predicting further in the future, the amount of shared information become much lower, and the FOP model is able to infer more global structures [19]. Recent works have successfully used the idea behind FOP to learn robust and generic representations in audio, textual and visual domains [19, 20, 21, 22]. Mikolov et al. [20] learned word representations by predicting neighboring words. Zhang et al. [21] learned image representations by predicting color for grey-scale images. Oord et al. [19] proposed a general framework to learn effective representations for multiple domains.

To capture long-term dynamic dependencies, we propose to use the self-attention mechanism [23] for FOP. As the self-attention mechanism can provide an opportunity for injecting global context of the whole sequence into each input frame, it can attend to longer sequences than many typical RNN-based models [23, 24].

FOP can learn discriminative representations from the inputs. To distill the “knowledge” from FOP to emotion recognition, two transfer learning approaches are utilized, including Fine-tuning [9, 25] and Hypercolumns [26]. Fine-tuning [9, 25] is an approach that fine-tunes either the last or several of the last layers in the pre-trained model and leaves the remaining layers unchanged. While Hypercolumns [26] concatenate embeddings at different layers in a pre-trained model. The target task can benefit from Hypercolumns as they connect low-level features with high-level semantics.

This paper combines the unsupervised representation learning strategy – FOP, with transfer learning approaches (such as Fine-tuning and Hypercolumns) to improve the performance of speech emotion recognition. The main contributions of this paper lie in three aspects: 1) to deal with long-range temporal dependencies, we propose to utilize the self-attention mechanism for FOP; 2) to share the “knowledge” from FOP to speech emotion recognition, two transfer learning strategies are evaluated; 3) our proposed method is superior to other currently advanced unsupervised learning strategies. To the best of our knowledge, it is the first time that FOP is utilized as the unsupervised learning strategy to improve the performance of speech emotion recognition.

2. Proposed Method

Our training procedure consists of two stages. In the first stage, we train a model for FOP to learn discriminative representations from the inputs, which is marked as the **FOP model**. In the second stage, to share the “knowledge” from FOP to speech emotion recognition, we take advantage of transfer learning approaches including Fine-tuning (in Fig. 1(a)) and Hypercolumns (in Fig. 1(b)).

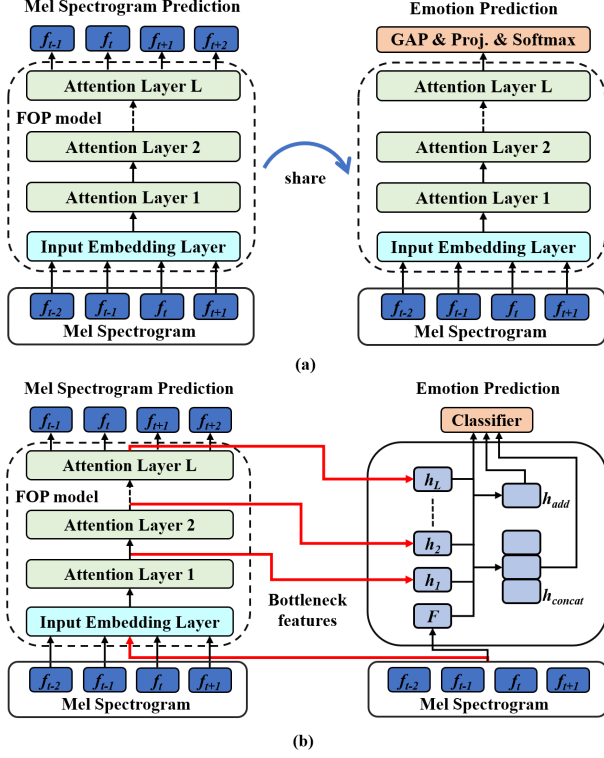


Figure 1: Overall structure of the proposed framework. (a) FOP combined with Fine-tuning: Through unsupervised pre-training, the FOP model can learn discriminative representations from the inputs. Then the output layer of the pre-trained FOP model is replaced by the combination of a global average pooling (GAP) layer and an emotion-specific projection layer. Finally, the whole model is fine-tuned for speech emotion recognition. (b) FOP combined with Hypercolumns: Bottleneck features are extracted from the pre-trained FOP model, followed with additional classifiers for speech emotion recognition.

2.1. Unsupervised pre-training

FOP can learn discriminative representations from the inputs. This paper utilizes a generative model operating on audio features for FOP. As depicted in Fig. 1, given an 80-dimensional mel-scale spectrogram $F = [f_1, f_2, \dots, f_T]$, where T is the number of frames and f_t is the features at timestep t . The joint probability of a feature sequence F is factorized as a product of conditional probabilities, which is computed as follows:

$$P(F) = \prod_{t=1}^T P(f_t | f_1, f_2, \dots, f_{t-1}; \theta) \quad (1)$$

where the conditional probability P is modeled using neural networks with parameters θ . Each f_t is therefore conditioned

on the features at all previous timesteps f_1, f_2, \dots, f_{t-1} .

Capturing long-term dynamic dependencies is important for FOP. Several architectures are evaluated, including convolutional neural networks (CNNs) [27], long-short term memory (LSTM) [28] and the masked multi-head self-attention mechanism [23]. We find that the last one is more suitable for FOP. As the self-attention mechanism can inject global context of the whole sequence into each input frame, it can attend to longer sequences than other two architectures. Therefore, the conditional probability distribution is modeled by a stack of masked multi-head self-attention mechanisms [23].

As depicted in Fig. 1, the FOP model has an input embedding layer and L identical attention layers. The input embedding layer injects position information into the inputs. The attention layer captures long-term dynamic dependencies. We describe these components below.

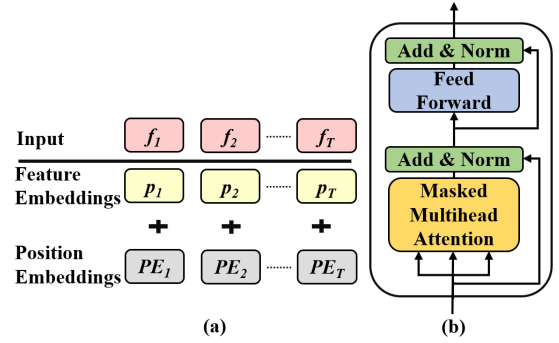


Figure 2: Two components in the FOP model. (a) Input embedding layer: The input embeddings are the sum of corresponding feature embeddings and position embeddings. (b) Attention layer: The attention layer has two sub-modules: a masked multi-head self-attention layer [23] and a feed forward layer.

2.1.1. Input embedding layer

We first describe the input embedding layer, illustrated in Fig. 2(a). The FOP model relies on the self-attention mechanism, which contains no recurrence and no convolution networks. If we shuffle the inputs $F = [f_1, f_2, \dots, f_T]$, we will get the same output [23]. To take the order of the sequence into consideration, we inject the position information into F .

This paper utilizes position embeddings the same with that in [23]. These embeddings are computed as follows:

$$PE_{(pos, 2i)} = \sin(pos/10000^{2i/d^{model}}) \quad (2)$$

$$PE_{(pos, 2i+1)} = \cos(pos/10000^{2i/d^{model}}) \quad (3)$$

where pos is the time step index, i is the dimension and d^{model} is the vector dimension of each frame.

The input embeddings z_t are the sum of corresponding feature embeddings p_t and position embeddings PE_t .

$$p_t = W_e f_t \quad (4)$$

$$z_t = p_t + PE_t \quad (5)$$

where W_e is the feature embedding matrix and f_t is the inputs at timestep t .

2.1.2. Attention layer

The input embeddings for all timesteps $h_0 = [z_1, z_2, \dots, z_T]$ are passed into a stack of L identical attention layers. The outputs of each attention layer are marked as $h_l, l \in [1, L]$. As depicted in Fig. 2(b), each attention layer has two sub-modules. The first is a masked multi-head self-attention mechanism [23], and the second is a feed forward layer. Then we employ a residual connection [29] around two sub-modules, followed with layer normalization [30]. Finally, the outputs of the last layer h_L is utilized to predict the next frame’s features.

The main ingredient of the attention layer is the masked self-attention mechanism [23]. With the help of this mechanism, we make sure that the prediction $P(f_t|f_1, f_2, \dots, f_{t-1})$ at timestep t cannot depend on any of the future timesteps features f_t, f_{t+1}, \dots, f_T .

2.2. Supervised transfer learning

To take advantage of discriminative representations that are learned from FOP, two transfer learning approaches are evaluated, including Fine-tuning and Hypercolumns.

2.2.1. Fine-tuning

As depicted in Fig. 1(a), the output layer of the pre-trained FOP model is replaced by the combination of a global average pooling (GAP) layer and an emotion-specific projection layer. Then the whole model is fine-tuned for speech emotion recognition. Therefore, only parameters in the projection layer is trained from scratch in Fine-tuning.

We assume that a labeled sample consists of a mel-scale spectrogram F , along with an emotional label y . We first pass F into the FOP model to obtain the outputs of the last attention layer h_L . Then h_L is fed into the GAP layer and the projection layer with parameters W_y to predict the emotional state y . This gives us the following objective to maximize:

$$P(y|F) = \text{softmax}(\text{GAP}(h_L)W_y) \quad (6)$$

2.2.2. Hypercolumns

As depicted in Fig. 1(b), the mel-scale spectrogram F and the outputs of each attention layer $h_l, l \in [1, L]$ are utilized as different levels of abstraction for the waveform. As F has the same feature dimensionality as $h_l, l \in [1, L]$, we further add and concatenate these representations together, which are marked as h_{add} and h_{concat} , respectively.

$$h_{add} = F + \sum_{l=1}^L h_l \quad (7)$$

$$h_{concat} = \text{Concat}(F, h_1, \dots, h_L) \quad (8)$$

3. Databases

IEMOCAP: The IEMOCAP [4] database contains about 12.46 hours of audiovisual data. There are five sessions with two actors each (one female and one male) and each session has different actors. Each session has been segmented into utterances, which are labeled into ten discrete labels (e.g., happy, sad, angry). To compared with other currently advanced approaches [31, 32], we form a four-class emotion classification dataset containing *angry*, *happy*, *sad* and *neutral*, where happy and excited categories are merged into a single happy category.

VCTK: The CSTR VCTK Corpus (Center for Speech Technology Voice Cloning Toolkit) [33] consists of 44 hours of data from 109 native speakers of English with various accents. Each speaker reads out about 400 sentences, which are selected from a newspaper plus the Rainbow Passage and an elicitation paragraph.

4. Experiments and Results

4.1. Experimental setup

In the experiments, we train a 2-layer FOP model with masked self-attention heads (256 dimensional states and 4 attention heads). For the feed forward layer in the FOP model, 513 dimensional inner states are chosen. We use the Adam [34] optimization scheme with a learning rate of 0.001 and a batch size of 32. To prevent over-fitting, we use dropout [35] with $p = 0.2$ and early stopping [36]. To alleviate the impact of the weight initialization, each configuration is tested 20 times. To consider the unbalanced number of samples between classes, weighted accuracy (WA) is chosen as our evaluation criterion.

Source task (FOP): We utilize all samples (without labels) in the IEMOCAP and VCTK databases.

Target task (speech emotion recognition): We conduct experiments on the IEMOCAP dataset. There are five sessions and each session has different actors. To ensure that models are trained and tested on speaker independent sets, utterances from the first 2 speakers are used as the testing set and utterances from other speakers are used as the training set.

4.2. Evaluation of FOP combined with Fine-tuning

This section discusses the effectiveness of FOP combined with Fine-tuning. As the FOP model can be pre-trained on different datasets, three pre-training settings are discussed including “VCTK”, “IEMOCAP” and “None”. “VCTK” and “IEMOCAP” represent that the FOP model is pre-trained on the VCTK dataset and the IEMOCAP dataset, respectively. While the pre-training process is not adopted in “None”.

Table 1: Performance of FOP combined with Fine-tuning under different pre-training settings.

Pre-training Settings	WA(%)
None: FOP without pre-training	65.45
VCTK: FOP pre-trained on VCTK	67.37
IEMOCAP: FOP pre-trained on IEMOCAP	68.14

Experimental results in Table 1 show that the pre-training process can improve the performance of speech emotion recognition. It reveals that after the pre-training process, the FOP model that learns to predict future information can also learn discriminative representations from waveforms. These high-level representations can be utilized to improve the recognition performance through Fine-tuning. What’s more, experimental results reveal that “IEMOCAP” gains higher classification accuracy than “VCTK”. As speech emotion recognition is evaluated on the IEMOCAP dataset, the FOP model pre-trained on the IEMOCAP dataset can match with speech emotion recognition better than the FOP model pre-trained on the VCTK dataset. It indicates that, in transfer learning, the similarity of the source task and the target task can improve the performance of the target task.

Therefore, the FOP model is pre-trained on the unlabeled IEMOCAP dataset in the following experiments.

4.3. Evaluation of FOP combined with Hypercolumns

Hypercolumns extract the outputs of each attention layer in the FOP model, followed with additional classifiers for speech emotion recognition. In this section, we discuss the performance of FOP combined with Hypercolumns.

Multiple features (in Sec 2.2) are investigated. Since the mel-scale spectrogram F is extracted without the FOP model, it is considered to be our baseline. In the 2-layer FOP model, the outputs of each attention layer are marked as h_1 and h_2 , respectively. As F has the same feature dimensionality as h_1 and h_2 , we add and concatenate these representations together, which are marked as h_{add} and h_{concat} , respectively. Besides various features, we also evaluate different classifiers for speech emotion recognition, including support vector machine (SVM), random forest (RF) and attention-based LSTM [17].

Table 2: *Weighted accuracy (%) of FOP combined with Hypercolumns under different combinations of features and classifiers. (Note: Baseline F is extracted without the FOP model.)*

Features	SVM	RF	A-LSTM	Best
F (Baseline)	61.04	62.57	58.93	62.57
h_1	63.92	56.66	64.68	64.68
h_2	63.53	57.39	62.38	63.53
h_{add}	63.72	58.35	64.88	64.88
h_{concat}	64.11	58.16	65.54	65.54

Experimental results in Table 2 demonstrate that F gains the worst performance, 62.57%, among all features. It indicates that FOP can learn robust and generic representations (such as h_1 , h_2) from the inputs F . These representations (such as h_1 , h_2) are more effective than the inputs F for speech emotion recognition. Therefore, we can conclude that transferring the “knowledge” from FOP into speech emotion recognition can improve recognition performance.

Furthermore, experimental results in Table 2 show that h_{add} and h_{concat} gain better performance than h_1 , h_2 and F in speech emotion recognition. F , h_1 and h_2 learn different levels of representations from the waveforms. As these representations are relevant to emotion recognition, h_{add} and h_{concat} that combine these emotion-related representations together can gain better performance. What’s more, we find that h_{concat} is superior to h_{add} . The reason lies in that h_{add} may result in the loss of original information in h_1 , h_2 and F . While h_{concat} not only can preserve original information, but also can combine these information properly.

4.4. Comparison to other advanced approaches

To show the effectiveness of the proposed method, we compare our method with currently advanced approaches through the five-folder cross validation.

Table 3: *Weighted accuracy (%) of currently advanced approaches and the proposed approach on the IEMOCAP dataset.*

	Single-folder	Cross-validation
Neumann et al. [31]	—	63.85
Lian et al. [32]	—	62.20
DAE in [37]	—	56.40
Ladder network in [37]	—	59.10
FOP+Hypercolumns	65.54	63.56
FOP+Fine-tuning	68.14	65.03

Compared with our proposed method, these approaches [31, 32] also utilized mel-scale spectrograms as inputs, and showed promising results for speech emotion recognition. Neumann et al. [31] proposed an attentive CNN with multi-view learning objective function for speech emotion recognition. The CNN learned the representations of the audio signal, while the attention layer computed the weighted sum of all the information extracted from different parts. Lian et al. [32] introduced the contrastive loss function for speech emotion recognition. This loss function encouraged intra-class compactness and inter-class separability between different emotional categories. Experimental results in Table 3 demonstrate the effectiveness of the proposed method. Our method outperforms currently advanced approaches [31, 32] in weight accuracy. It proves that FOP, which learns discriminative representations from the original inputs, can be utilized to improve the performance of speech emotion recognition.

Meanwhile, we compare our method with other semi-supervised and unsupervised learning strategies. Huang et al. [37] utilized semi-supervised learning with ladder networks, which outperformed unsupervised DAE above 2% for speech emotion recognition. Experimental results in Table 3 demonstrate that our method shows above 4% performance improvement over the ladder network in [37]. DAE and ladder structures pay less attention to modeling long-term dynamic dependencies. However, temporal information is important for speech emotion recognition [17, 18]. Therefore, our self-attention based FOP model, which can capture long-term dynamic dependencies, is superior to other currently advanced unsupervised learning strategies.

5. Conclusions

This paper combines the unsupervised representation learning strategy – FOP, with transfer learning approaches (such as Fine-tuning and Hypercolumns) for speech emotion recognition. To evaluate the effectiveness of the proposed method, we conduct experiments on the IEMOCAP database. Experimental results reveal that FOP can learn robust and discriminative representations from the mel-scale spectrograms. These representations are more effective than the original mel-scale spectrograms for speech emotion recognition. As our method can capture long-term dynamic dependencies, it also outperforms the mainstream semi-supervised and unsupervised learning strategies (such as DAE and ladder structures) above 4% for speech emotion recognition.

Future investigations include a detailed analysis of input features for FOP. Besides mel-scale spectrograms, other audio features (e.g., Mel Frequency Cepstrum Coefficients (MFCCs), Linear Predictive Codings (LPCs)) should be evaluated. Additionally we aim to investigate and further improve the classification accuracy using other advanced fine-tuning approaches (e.g., discriminative fine-tuning in [38]). Finally, we will apply the proposed method to other audio classification tasks, such as key word spotting and speaker identification.

6. Acknowledgements

This work is supported by the National Key Research & Development Plan of China (No.2017YFC0820602), the National Natural Science Foundation of China (NSFC) (No.61425017, No.61831022, No.61773379, No.61771472), and the Strategic Priority Research Program of Chinese Academy of Sciences (No.XDC02050100).

7. References

- [1] B. Schuller, A. Batliner, S. Steidl, and D. Seppi, "Recognising realistic emotions and affect in speech: State of the art and lessons learnt from the first challenge," *Speech Communication*, vol. 53, no. 9-10, pp. 1062–1087, 2011.
- [2] O. Martin, I. Kotsia, B. Macq, and I. Pitas, "The enterface'05 audio-visual emotion database," in *22nd International Conference on Data Engineering Workshops (ICDEW'06)*. IEEE, 2006, pp. 8–8.
- [3] 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.
- [4] 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.
- [5] A. Batliner, C. Hacker, S. Steidl, E. Nöth, S. D'Arcy, M. J. Russell, and M. Wong, "'you stupid tin box'-children interacting with the aibo robot: A cross-linguistic emotional speech corpus." in *Lrec*, 2004.
- [6] J. Deng, R. Xia, Z. Zhang, Y. Liu, and B. Schuller, "Introducing shared-hidden-layer autoencoders for transfer learning and their application in acoustic emotion recognition," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 4818–4822.
- [7] S. E. Eskimez, Z. Duan, and W. Heintzelman, "Unsupervised learning approach to feature analysis for automatic speech emotion recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5099–5103.
- [8] 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.
- [9] X. Ouyang, S. Kawaai, E. G. H. Goh, S. Shen, W. Ding, H. Ming, and D.-Y. Huang, "Audio-visual emotion recognition using deep transfer learning and multiple temporal models," in *Proceedings of the 19th ACM International Conference on Multimodal Interaction*. ACM, 2017, pp. 577–582.
- [10] R. Xia and Y. Liu, "Using denoising autoencoder for emotion recognition." in *Interspeech*, 2013, pp. 2886–2889.
- [11] S. Ghosh, E. Laksana, L.-P. Morency, and S. Scherer, "Representation learning for speech emotion recognition." in *Interspeech*, 2016, pp. 3603–3607.
- [12] C. Poultney, S. Chopra, Y. L. Cun *et al.*, "Efficient learning of sparse representations with an energy-based model," in *Advances in neural information processing systems*, 2007, pp. 1137–1144.
- [13] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle, "Greedy layer-wise training of deep networks," in *Advances in neural information processing systems*, 2007, pp. 153–160.
- [14] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol, "Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion," *Journal of machine learning research*, vol. 11, no. Dec, pp. 3371–3408, 2010.
- [15] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [16] A. Rasmus, M. Berglund, M. Honkala, H. Valpola, and T. Raiko, "Semi-supervised learning with ladder networks," in *Advances in neural information processing systems*, 2015, pp. 3546–3554.
- [17] C.-W. Huang and S. S. Narayanan, "Attention assisted discovery of sub-utterance structure in speech emotion recognition." in *INTERSPEECH*, 2016, pp. 1387–1391.
- [18] J. Kim, G. Englebienne, K. P. Truong, and V. Evers, "Towards speech emotion recognition 'in the wild' using aggregated corpora and deep multi-task learning," *arXiv preprint arXiv:1708.03920*, 2017.
- [19] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [20] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.
- [21] R. Zhang, P. Isola, and A. A. Efros, "Colorful image colorization," in *European Conference on Computer Vision*. Springer, 2016, pp. 649–666.
- [22] C. Doersch, A. Gupta, and A. A. Efros, "Unsupervised visual representation learning by context prediction," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 1422–1430.
- [23] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [24] P. J. Liu, M. Saleh, E. Pot, B. Goodrich, R. Sepassi, L. Kaiser, and N. Shazeer, "Generating wikipedia by summarizing long sequences," *arXiv preprint arXiv:1801.10198*, 2018.
- [25] S. Thrun and L. Pratt, *Learning to learn*. Springer Science & Business Media, 2012.
- [26] B. Hariharan, P. Arbeláez, R. Girshick, and J. Malik, "Hypercolumns for object segmentation and fine-grained localization," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 447–456.
- [27] Y. LeCun, Y. Bengio *et al.*, "Convolutional networks for images, speech, and time series," *The handbook of brain theory and neural networks*, vol. 3361, no. 10, p. 1995, 1995.
- [28] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [29] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [30] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.
- [31] M. Neumann and N. T. Vu, "Attentive convolutional neural network based speech emotion recognition: A study on the impact of input features, signal length, and acted speech," *INTERSPEECH*, 2017.
- [32] Z. Lian, Y. Li, J. Tao, and H. Jian, "Speech emotion recognition via contrastive loss under siamese networks," in *International Workshop on Affective Social Multimedia Computing*, 2018.
- [33] J. Yamagishi, "English multi-speaker corpus for cstr voice cloning toolkit," URL <http://homepages.inf.ed.ac.uk/jyamagis/page3/page58/page58.html>, 2012.
- [34] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [35] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [36] L. Prechelt, "Automatic early stopping using cross validation: quantifying the criteria," *Neural Networks*, vol. 11, no. 4, pp. 761–767, 1998.
- [37] J. Huang, Y. Li, J. Tao, Z. Lian, M. Niu, and J. Yi, "Speech emotion recognition using semi-supervised learning with ladder networks," in *2018 First Asian Conference on Affective Computing and Intelligent Interaction (ACII Asia)*. IEEE, 2018, pp. 1–5.
- [38] J. Howard and S. Ruder, "Universal language model fine-tuning for text classification," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, vol. 1, 2018, pp. 328–339.