



Queensland University of Technology
Brisbane Australia

This may be the author's version of a work that was submitted/accepted for publication in the following source:

Madduma Kankanamalage Don, Darshana Priyasad, Warnakulasuriya, Tharindu Fernando, Denman, Simon, Sridharan, Sridha, & Fookes, Clinton
(2020)

Attention driven fusion for multi-modal emotion recognition.

In *Proceedings of ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.

Institute of Electrical and Electronics Engineers Inc., United States of America, pp. 3227-3231.

This file was downloaded from: <https://eprints.qut.edu.au/202004/>

© 2020 IEEE

2020 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Notice: *Please note that this document may not be the Version of Record (i.e. published version) of the work. Author manuscript versions (as Submitted for peer review or as Accepted for publication after peer review) can be identified by an absence of publisher branding and/or typeset appearance. If there is any doubt, please refer to the published source.*

<https://doi.org/10.1109/ICASSP40776.2020.9054441>

ATTENTION DRIVEN FUSION FOR MULTI-MODAL EMOTION RECOGNITION

Darshana Priyasad, Tharindu Fernando, Simon Denman, Sridha Sridharan, Clinton Fookes

Speech and Audio Research Lab - SAIVT
Queensland University of Technology, Brisbane, Australia

ABSTRACT

Deep learning has emerged as a powerful alternative to hand-crafted methods for emotion recognition on combined acoustic and text modalities. Baseline systems model emotion information in text and acoustic modes independently using Deep Convolutional Neural Networks (DCNN) and Recurrent Neural Networks (RNN), followed by applying attention, fusion, and classification. In this paper, we present a deep learning-based approach to exploit and fuse text and acoustic data for emotion classification. We utilize a SincNet layer, based on parameterized sinc functions with band-pass filters, to extract acoustic features from raw audio followed by a DCNN. This approach learns filter banks tuned for emotion recognition and provides more effective features compared to directly applying convolutions over the raw speech signal. For text processing, we use two branches (a DCNN and a Bi-direction RNN followed by a DCNN) in parallel where cross attention is introduced to infer the N-gram level correlations on hidden representations received from the Bi-RNN. Following existing state-of-the-art, we evaluate the performance of the proposed system on the IEMOCAP dataset. Experimental results indicate that the proposed system outperforms existing methods, achieving **5.2%** improvement in weighted accuracy.

Index Terms— Speech emotion recognition, deep learning, multi-modal fusion, cross attention, SincNet

1. INTRODUCTION

With the advance of technology, Human Computer Interaction (HCI) has become a major research area. Within this field, automatic emotion recognition is being pursued as a means to improve the level of user experience by tailoring responses to the emotional context, especially in human-machine interactions [1]. However, this remains challenging due to the ambiguity of expressed emotions. An utterance may contain subject dependent auditory clues regarding expressed emotions which are not captured through speech transcripts alone. With deep learning, architectures can extract higher level features and more robust features for accurate speech emotion recognition [2, 3]. In this paper, we propose a model that combines acoustic and textual information for speech emotion recognition.

Recently, multi-modal information has been used in emotion recognition in preference to uni-modal methods [4], since humans express emotion via multiple modes [5, 6, 7, 8]. Most state-of-the-art methods for utterance level emotion recognition have used low-level (energy) and high-level acoustic features (such as Mel Frequency Cepstral Coefficients (MFCC) [5, 9]). However, when the emotion expressed through speech becomes ambiguous, the lexical content may provide complementary information that can address the ambiguity.

Tripathi *et al.*[10] has used a Long-Short Term Memory (LSTM) along with a DCNN to perform joint acoustic and textual emotion recognition. They have used features such as MFCC, Zero Crossing Rate and spectral entropy for acoustic data, while using Glove [11] embeddings to extract a feature vector from speech transcripts. However, the performance gain is minimal due to the lack of robustness in the acoustic features, and the sparse text feature vectors. Yenigalla *et al.*[12] proposed a spectral and phenoms-sequence based DCNN model which is capable of retaining the emotional content of the speech that is lost when converted to text. Yoon *et al.*[5] presented a framework using Bidirectional LSTMs to obtain hidden representations of acoustic and textual information. The resultant features are fused with multi-hop attention, where one modality is used to direct attention for the other mode. Higher performance has been achieved due to the attention and fusion which select relevant segments from the textual model, and complementary features from each mode. Yoon *et al.*[6] have also presented an encoder based method where the fusion of two recurrent encoders is used to combine features from audio and text. However, both methods use manually calculated audio features which limit their accuracy and the robustness of the acoustic model [13, 14].

Gu *et al.*[7] presented a multimodal framework where a hybrid deep multimodal structure that considers spatial and temporal information is employed. Features obtained from each model were fused using a DNN to classify the emotion. Li *et al.*[8] proposed a personalized attribute aware attention mechanism where an attention profile is learned based on acoustic and lexical behavior data. Mirsamadi *et al.*[15] used deep learning along with local attention to automatically extract relevant features where segment level acoustic features are aggregated for utterance level emotion represen-

tation. However, the accuracy could be further improved by fusing the audio features with another modality with complementary information.

In this paper, we present a multi-modal emotion recognition model with combines acoustic and textual information by using DCNNs, and both cross attention and self-attention. Experiments are performed on the IEMOCAP [16] dataset to enable fair comparison with state-of-the-art methods, and a performance gain of 5.2% in weighted accuracy is achieved.

2. METHODOLOGY

2.1. Acoustic Feature Extraction

In our proposed model, we utilize a SincNet filter layer [17] to learn custom filter banks tuned for emotion recognition from speech audio. This layer is shown to have fast convergence and higher interpretability with a smaller number of parameters compared to conventional convolution layers. Formally, this layer can be defined as,

$$y[n] = x[n] * g[n, \theta], \quad (1)$$

where $x[n]$, $y[n]$, g , θ refers to the input signal, filtered output, filter-bank function, and the learnable parameters respectively. In SincNet filters, convolution operations are applied over a raw waveform with predefined functions. Each defined filter-bank is composed of rectangular band-pass filters which can be represented by two low-pass filters with learnable cutoff frequencies. The time-domain representation of the function g can be derived as follows [17],

$$g[n, f_1, f_2] = 2f_2 \operatorname{sinc}(2\pi f_2 n) - 2f_1 \operatorname{sinc}(2\pi f_1 n), \quad (2)$$

where f_1, f_2 refers to low and high cutoff frequencies and $\operatorname{sinc} = \sin x/x$.

The resultant convolution layer outputs are passed through a DCNN which contains several ‘‘Convolution1D’’, ‘‘Batch Normalization’’ and ‘‘fully connected layers’’. During the initial training phase, a random 250ms chunk from the audio signal is selected as the input. During validation and testing, we obtain the final ‘‘softmax’’ response for each chunk and add them together to get the final classification scores, similar to [17]. However, we retrieve a 2048-D feature vector from the final dense layer before the classification layer for each chunk of an utterance, and average these before fusing them with textual features from the corresponding transcript in a later step (see Section 2.3).

2.2. Textual Feature Extraction

In our proposed model, after the sequence vector is passed through a common embedding layer, we utilize two parallel branches for textual feature extraction as illustrated in Figure 1. Bi-RNNs followed by DCNNs have been extensively used in textual emotion analysis [18, 19]. As an alternative, a CNN

based architecture which is capable of considering ‘‘n’’ words at a time (n-grams) can be used [20]. Therefore we use two parallel branches, employing, one using Bi-RNNs with DCNNs and the other DCNNs alone to increase the effectiveness of the learned features (see Figure 1 (B)). The resultant feature vector from the Bi-RNN is passed through three convolutional layers with a filter sizes of 1, 3 and 5; and convolutional layers with the same size filter are used in the parallel branch. We introduce cross-attention where we use convolution layers with the same filter size from the right branch as the attention for the left branch, as illustrated, and jointly train with the other components of the network. The cross-attention is calculated using

$$\alpha_i = \frac{\exp((\mathbf{b}_{i,j})^\top \mathbf{a}_{i,j}^\top)}{\sum_i \exp((\mathbf{b}_{i,j})^\top \mathbf{a}_{i,j}^\top)}, \quad (3)$$

$$H = \sum_i \alpha_i \mathbf{a}_{i,j}^\top, \quad (4)$$

where $\alpha_i, b_{i,j}, a_{i,j}, H$ are the attention score, context vector from the right branch with a filter size of j , output of the convolution layer with filter size j in left branch, and the output.

The convolution layers from both branches are concatenated together and passed through a DNN consisting of fully connected layers for textual emotion classification. We retrieve a 4800-D feature vector from the final dense layer before the classification layer for multi-modal feature fusion.

2.3. Acoustic & Textual Feature Fusion

Mid-level fusion is used to fuse textual and acoustic features obtained from individual networks. A 2048-D feature vector from the acoustic network and a 4800-D feature vector from the textual network are concatenated as illustrated in Figure 1 (C). A neural network with attention is used to identify informative segments in the feature vectors. We have explored using fusion without attention (F-I), attention after fusion (F-II) where self-attention is applied on concatenated features, and attention before fusion (F-III) where attention is applied on individual feature vectors. For F-III, we calculate attention weights and combine the vectors using [21],

$$h_t = \tanh(\mathbf{W}_h \mathbf{c}_t + b_h), \quad (5)$$

$$\beta_t = \frac{\exp(|\mathbf{h}_t|^\top \hat{\mathbf{h}})}{\sum_t \exp(|\mathbf{h}_t|^\top \hat{\mathbf{h}})}, \quad (6)$$

$$q = \sum_t \beta_t \times \mathbf{c}_t. \quad (7)$$

$c_t, h_t, \beta_t, \hat{h}$ and q refer to the merged feature vector, neural network output, attention score, context vector (which is randomly initialized and jointly trained with other components of the network), and the output respectively. Finally the the utterance emotion is classified using a ‘‘softmax’’ activation over the final dense layer of the fusion network.

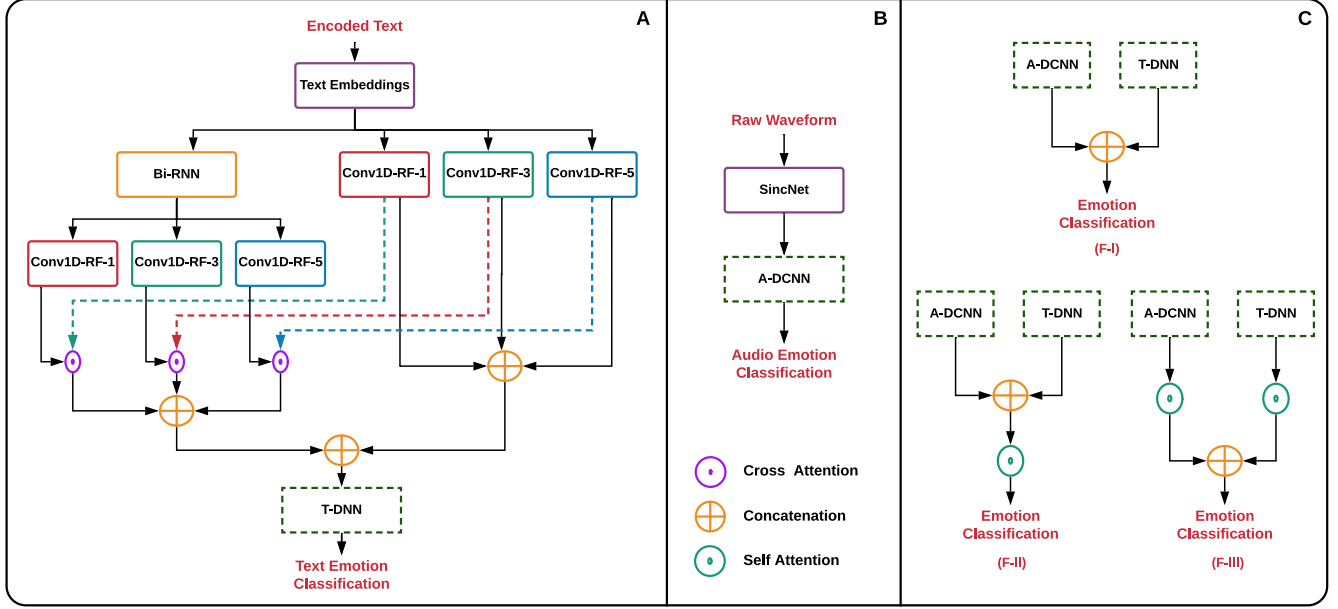


Fig. 1. Proposed architecture - The system contains three main parts: the text network (A); the audio network (B); and the fusion network (C). The raw audio is passed to the acoustic network after amplitude normalization. It contains SincNet filter layers containing parameterized sinc functions with band-pass filters to extract acoustic features, followed by the A-DCNN component containing convolutional and dense layers. The corresponding text is converted to a word vector using Glove embeddings, and then passed through the text network. It contains two parallel branches with bi-LSTM and convolutional layers (12 filters) with different filter sizes to capture n-grams (n-words) in one iteration where $n = \{1,3,5\}$. As shown in the figure, convolutional layers in the right branch are used as cross-attention for the left branch. The two branches are fused, followed by the T-DNN component for textual emotion recognition. A-DCNN and T-DNN are then fused using self-attention to get the final emotion classification result.

3. EXPERIMENTS

3.1. Dataset and Experimental Setup

Experiments are conducted on the Interactive Emotional Dyadic Motion Capture (IEMOCAP) dataset which includes five sessions of utterances for 10 unique speakers. We follow the evaluation protocol of [5, 6], and select utterances annotated with four basic emotions “anger”, “happiness”, “neutral” and “sadness”. Samples with “excitement” are merged with “happiness” as per [5, 6]. The resultant dataset contains 5531 utterances {“anger”:1103, “happiness”:1636, “neutral”:1708, “sadness”:1084}.

Initial training is carried out on both acoustic and textual networks separately before the fusion. The sampling rate of each utterance waveform is set to 16,000Hz while a random segment of 250ms is used in training the acoustic network. During the evaluation, the cumulative sum of all the predictions with a window size and shift of 250ms and 10ms are considered. In the textual network, all the transcripts of the utterances are set to a maximum length of 100 and padded with 0s. Glove-300d embeddings are used to convert the word sequence to a vector of (100, 300). We utilize 10-fold cross-validation with an 8:1:1 split for training, validation, and test

sets respectively. The learning rate and the batch size in each network are fixed at 0.001 and 64 respectively, and the Adam optimiser is used.

3.2. Performance Evaluation and Analysis

Following [5, 6], the performance of our system is measured in terms of weighted accuracy (WA) and unweighted accuracy (UA). Table 1 and Figure 2 present performance of our approach for emotion recognition compared with the state of the art methods.

MDRE [6] has used two RNNs to encode both acoustic and textual data followed by a DNN for classification, while Evcc-MCNN-LSTM [22] has used an RNN and a DCNN to encode both modalities followed by fusion and an SVM for classification. MHA-2 [5] has used two Bidirectional Recurrent Encoders (BRE) for both modalities followed by a multi-hop attention mechanism. MDRE has outperformed MCNN-LSTM by 10.6% and MHA-2 has outperformed MDRE by 6.5%, demonstrating how attention can increase performance.

Our proposed model has achieved a substantial improvement in overall accuracy, with a 5.2% increase compared to MHA-2. We have utilized self-attention before (F-II) and after fusion (F-III) as illustrated in Figure 1. Cross-modal atten-




Fig. 2. Confusion matrices of the proposed architecture for separate fusion methods calculated using a 10-fold cross validation. Left, middle and right figures represent Fusion I, Fusion II and Fusion III respectively.

Model	Modality	WA	UA
Evvec-MCNN-LSTM [22]	$A + T$	64.9%	65.9%
MDRE [6]	$A + T$	71.8%	—
MHA-2 [5]	$A + T$	76.5%	77.6%
Ours - F-I	$A + T$	79.27%	77.85%
Ours - F-II	$A + T$	80.01%	78.98%
Ours - F-III	$A + T$	80.51%	79.22%

Table 1. Recognition accuracy for IEMOCAP using a 10-fold cross-validation, compared with state of the art methods.

Model	Modality	UA
MHA-2 [5]	A	65.2%
MHA-2 [5]	T	70.3%
Ours	A	69.89%
Ours	T	70.81%

Table 2. Recognition accuracy individual modes of the IEMOCAP dataset with 10-fold cross-validation, comparing the proposed approach with [5].

tion has not been utilized after fusion since the dimensionality of the feature vectors from the two modalities are different. A slight increase in the classification accuracy has been obtained by applying self-attention compared to conventional feature fusion (F-I). Furthermore, the highest accuracy has been obtained by **F-III**, outperforming F-II by 0.6%. Given that F-I slightly outperforms MHA-2, we have compared the classification accuracy of the individual modes of MHA-2 with our individual modes in Table 2.

Our acoustic and textual models outperformed the corresponding individual modes of MHA-2 [5], where a substantial improvement of **7.2%** is achieved with the acoustic model.

The SincNet layer in the acoustic model is capable of learning and deriving customized filter banks tuned for emotion recognition. It has been successfully applied for speaker recognition [17] as an alternative to i-vectors. The confusion matrices for F-I, F-II, and F-III are illustrated in Figure 2. A **12.9%** improvement in classification accuracy can be observed when comparing the individual modalities with the fusion network in our model. Given that accuracy is approximately similar for both modalities, each modality has complemented the other to increase recognition accuracy.

4. CONCLUSION

In this paper, we present an attention-based multi-modal emotion recognition model combining acoustic and textual data. The raw audio waveform is utilized in our method, rather than extracting hand-crafted features as done by baseline methods. Combining a DCNN with a SincNet layer, which learns suitable filter parameters over the waveform for emotion recognition, outperforms the hand-crafted feature-based audio emotion detection of the baselines. Cross attention is applied to text-based feature extraction to guide the features derived by RNNs using N-gram level features extracted by a parallel branch. We have used self-attention on both feature vectors obtained from two networks before the fusion, to attend to the informative segments from each feature vector. We have achieved a weighted accuracy of **80.51%** on the IEMOCAP database, which outperforms the existing state-of-the-art model by **5.2%**.

5. ACKNOWLEDGEMENTS

This research was supported by an Australia Research Council (ARC) Discovery grant DP140100793.

6. REFERENCES

- [1] A. Mohanta and U. Sharma, "Detection of human emotion from speech—tools and techniques," in *Speech and Language Processing for Human-Machine Communications*. Springer, 2018, pp. 179–186.
- [2] P. Tzirakis, J. Zhang, and B. W. Schuller, "End-to-end speech emotion recognition using deep neural networks," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5089–5093.
- [3] L. Tarantino, P. N. Garner, and A. Lazaridis, "Self-attention for speech emotion recognition," *Proc. Interspeech 2019*, pp. 2578–2582, 2019.
- [4] M. El Ayadi, M. S. Kamel, and F. Karray, "Survey on speech emotion recognition: Features, classification schemes, and databases," *Pattern Recognition*, vol. 44, no. 3, pp. 572–587, 2011.
- [5] S. Yoon, S. Byun, S. Dey, and 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.
- [6] S. Yoon, S. Byun, and K. Jung, "Multimodal speech emotion recognition using audio and text," in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 112–118.
- [7] Y. Gu, S. Chen, and I. Marsic, "Deep mul timodal learning for emotion recognition in spoken language," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5079–5083.
- [8] J.-L. Li and C.-C. Lee, "Attentive to individual: A multimodal emotion recognition network with personalized attention profile," *Proc. Interspeech 2019*, pp. 211–215, 2019.
- [9] D. Nguyen, K. Nguyen, S. Sridharan, D. Dean, and C. Fookes, "Deep spatio-temporal feature fusion with compact bilinear pooling for multimodal emotion recognition," *Computer Vision and Image Understanding*, vol. 174, pp. 33–42, 2018.
- [10] S. Tripathi and H. Beigi, "Multi-modal emotion recognition on iemocap dataset using deep learning," *arXiv preprint arXiv:1804.05788*, 2018.
- [11] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1532–1543. [Online]. Available: <http://www.aclweb.org/anthology/D14-1162>
- [12] P. Yenigalla, A. Kumar, S. Tripathi, C. Singh, S. Kar, and J. Vepa, "Speech emotion recognition using spectrogram & phoneme embedding," in *Interspeech*, 2018, pp. 3688–3692.
- [13] C. W. Lee, K. Y. Song, J. Jeong, and W. Y. Choi, "Convolutional attention networks for multimodal emotion recognition from speech and text data," *ACL 2018*, p. 28, 2018.
- [14] M. Chen, X. He, J. Yang, and H. Zhang, "3-d convolutional recurrent neural networks with attention model for speech emotion recognition," *IEEE Signal Processing Letters*, vol. 25, no. 10, pp. 1440–1444, 2018.
- [15] S. Mirsamadi, E. Barsoum, and C. Zhang, "Automatic speech emotion recognition using recurrent neural networks with local attention," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 2227–2231.
- [16] 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.
- [17] M. Ravanelli and Y. Bengio, "Speaker recognition from raw waveform with sincnet," in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 1021–1028.
- [18] C. Etienne, G. Fidanza, A. Petrovskii, L. Devillers, and B. Schmauch, "Cnn+ lstm architecture for speech emotion recognition with data augmentation," *arXiv preprint arXiv:1802.05630*, 2018.
- [19] G. Ramet, P. N. Garner, M. Baeriswyl, and A. Lazaridis, "Context-aware attention mechanism for speech emotion recognition," in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 126–131.
- [20] Y. Kim, "Convolutional neural networks for sentence classification," *arXiv preprint arXiv:1408.5882*, 2014.
- [21] D. Priyasad, T. Fernando, S. Denman, S. Sridharan, and C. Fookes, "Learning salient features for multimodal emotion recognition with recurrent neural networks and attention based fusion," *15th International Conference on Auditory Visual Speech Processing (AVSP)*, 2019.
- [22] J. Cho, R. Pappagari, P. Kulkarni, J. Villalba, Y. Carmiel, and N. Dehak, "Deep neural networks for emotion recognition combining audio and transcripts," *Proc. Interspeech 2018*, pp. 247–251, 2018.