

Group Gated Fusion on Attention-based Bidirectional Alignment for Multimodal Emotion Recognition

Pengfei Liu¹, Kun Li¹ and Helen Meng²

¹SpeechX Limited, Shenzhen, China

²The Chinese University of Hong Kong, Hong Kong SAR, China

pfliu@speechx.cn, kli@speechx.cn, hmmeng@se.cuhk.edu.hk

Abstract

Emotion recognition is a challenging and actively-studied research area that plays a critical role in emotion-aware human-computer interaction systems. In a multimodal setting, temporal alignment between different modalities has not been well investigated yet. This paper presents a new model named as Gated Bidirectional Alignment Network (GBAN), which consists of an attention-based bidirectional alignment network over LSTM hidden states to explicitly capture the alignment relationship between speech and text, and a novel group gated fusion (GGF) layer to integrate the representations of different modalities. We empirically show that the attention-aligned representations outperform the last-hidden-states of LSTM significantly, and the proposed GBAN model outperforms existing state-of-the-art multimodal approaches on the IEMOCAP dataset.

Index Terms: multimodal emotion recognition, attention models, information fusion

1. Introduction

Emotion recognition is a core component in any emotion-aware human-computer interaction system, such as an intelligent virtual assistant, and an affective spoken dialog system. An emotion recognizer typically analyzes speech, text or images. For example, a speech emotion recognizer aims to identify the emotion carried in speech, often in terms of a set of emotion categories such as happy, angry, sad and neutral [1–5]. However, this is a nontrivial task because emotions are manifested in various factors such as conversational discourse, linguistic content and prosodic features [6–8]. It is difficult to predict the emotion of an utterance based only on the acoustic features from speech, or only on the discrete word sequence in spoken text. Furthermore, there is a lack of large-scale emotive datasets of speech or text, as these are costly and difficult to collect and label. One should therefore explore the use of multiple modalities for emotion recognition.

Multimodal emotion recognition has received a lot of attention in recent years [8–17]. Modalities such as speech, text and/or images have been exploited for better recognition performance. Existing approaches to multimodal learning are either *early-fusion* which fuses low-level feature interactions between different modalities before making a prediction decision; or *late-fusion* which models each modality independently and combines the decisions from each model [13]. Recent early-fusion approaches have focused on different mathematical formulations to fuse acoustic and lexical features, such as multimodal pooling fusion [13], modality hierarchical fusion [16], conversational neural network [18], word-level concatenation of acoustic and lexical features [8], etc. To capture inter-modality dynamics in multimodal sentiment analysis, Zadeh et al. [12]

introduces the Tensor Fusion Network (TFN) model using tensor fusion to explicitly aggregate unimodal, bimodal and trimodal interactions between different features. Arevalo et al. [19] presents a Gated Multimodal Unit (GMU) model to find an intermediate representation based on a combination of features from different modalities, where it learns to decide how modalities influence the activation of the unit using multiplicative gates. These methods are generally applicable in information fusion from different sources. However, they do not take into account the particular temporal correlation property of speech and text in an utterance, where a sequence of speech frames are aligned with a sequence of words temporally.

Given a sequence of speech frames, it is beneficial to know the corresponding words to learn a more discriminative representation for emotion recognition, and vice versa. For example, consider a speaker expressing *happiness* in the utterance “*That’s great!*”, and *anger* in “*That’s unfair!*”. An emotion recognizer should learn to pay more attention to the words *great* and *unfair* and their *corresponding* speech frames, as there is complementary information across modalities for emotion recognition. This calls for a mechanism to capture the alignment between words and their speech frames. Yoon et al. [14] presents a deep dual recurrent neural network to encode information from speech and text sequences, and captures the alignment by a dot product attention to focus on specific words of a text sequence conditioned on speech for emotion classification. Similarly, Xu et al. [17] learns the alignment between speech and text using an additive attention, and adopts an LSTM layer fed with the concatenation of the aligned speech representation and the hidden-state text representation for emotion classification. However, both approaches [14, 17] consider only unidirectional alignment, and we will show empirically that bidirectional alignment brings better recognition performance. Furthermore, the concatenation-based fusion methods in [14, 17] may not have sufficient expressive power to exploit the complementary information across modalities, which in turn reflects the need for more effective ways to fuse multiple representations.

In this paper, we propose a model named Gated Bidirectional Alignment Network (GBAN), which consists of a bidirectional attention-based alignment network to capture the alignment information between speech and text, and a novel group gated fusion (GGF) layer to automatically learn the contribution of each modality in determining the final emotion category. We summarize the major contributions of this paper as follows:

- (1) Experimental results show that the proposed bidirectional alignment network leads to more discriminative representations of speech and text for emotion recognition;
- (2) The proposed GGF method allows integration of multiple representations effectively, and obtains interpretable contribution weights of each modality in emotion recognition;

- (3) The proposed GBAN model outperforms existing state-of-the-art multimodal approaches [14, 17] using both speech and text on the IEMOCAP dataset [20].

2. Gated Bidirectional Alignment Network

We propose an approach named Gated Bidirectional Alignment Network (GBAN) for multimodal emotion recognition. As illustrated in Figure 1, the GBAN model consists of three major parts: (1) two separate CNN-LSTM encoders to extract features from speech and text respectively; (2) an attention-based bidirectional alignment network to capture temporal correlations between speech and text; and (3) a group gated fusion layer to learn the contribution of each representation automatically.




Figure 1: The block diagram of the proposed GBAN model.

2.1. CNN-LSTM Encoders

We adopt two separate CNN-LSTM encoders for speech and text respectively, where CNN layers extract local features and a bidirectional LSTM layer captures the global dependencies over time in either the speech acoustics or the text of an utterance.

Speech Representation. The speech signal in an utterance is represented as a sequence of vectors $[x_1, \dots, x_N]$, where N is the number of frames in the utterance. First, the sequence is fed into a CNN layer, which consists of convolution and max-pooling operations, to obtain a sequence of pooled vectors $[p_1, \dots, p_K]$. Then, a bidirectional LSTM layer follows the CNN layer to obtain a vector s_i for each $i \in \{1, \dots, K\}$, which is concatenated from the forward and backward LSTMs. Thus, the obtained sequence of $[s_1, \dots, s_K]$ keeps the temporal ordering for each $i \in \{1, \dots, K\}$, and will be used as the speech representation for alignment with the text representation. In practice, multiple CNN layers and bidirectional LSTM layers are adopted to encode the low-level speech signals.

$$p_i = CNN([x_1, \dots, x_N]), i \in \{1, \dots, K\} \quad (1)$$

$$\vec{s}_i = \overrightarrow{LSTM}(p_i), i \in \{1, \dots, K\} \quad (2)$$

$$\overleftarrow{s}_i = \overleftarrow{LSTM}(p_i), i \in \{1, \dots, K\} \quad (3)$$

$$s_i = [\vec{s}_i, \overleftarrow{s}_{K-i+1}], i \in \{1, \dots, K\} \quad (4)$$

Text Representation. For text encoding, each sentence is represented as a sequence of vectors $[e_1, \dots, e_M]$, where M is the number of words in a sentence, and e_j is the word embeddings of the j th word. Similar to the process in speech representation, a CNN layer is first used to obtain a sequence of pooled vectors $[q_1, \dots, q_L]$. For each $j \in \{1, \dots, L\}$, a bidirectional LSTM layer follows to obtain the concatenation vector t_j from the forward and backward LSTMs. The sequence of $[t_1, \dots, t_L]$, which are essentially text features transformed from the sequence of M words, keeps the temporal ordering for each $j \in \{1, \dots, L\}$, and will be used as the text representation for alignment with the speech representation of $[s_1, \dots, s_K]$.

$$q_j = CNN([e_1, \dots, e_M]), j \in \{1, \dots, L\} \quad (5)$$

$$\vec{t}_j = \overrightarrow{LSTM}(q_j), j \in \{1, \dots, L\} \quad (6)$$

$$\overleftarrow{t}_j = \overleftarrow{LSTM}(q_j), j \in \{1, \dots, L\} \quad (7)$$

$$t_j = [\vec{t}_j, \overleftarrow{t}_{L-j+1}], j \in \{1, \dots, L\} \quad (8)$$

2.2. Bidirectional Alignment Network

An utterance is made up of a speech signal and a text sequence (transcribed by a person or a speech recognizer), where the speech and the text are sequentially correlated over time. To capture their temporal correlations, we propose a bidirectional alignment network to learn the representation of one modality with the help of the other based on attention models.

Attention-Aligned Speech Representation. Given the speech representation $[s_1, \dots, s_K]$ and the text representation $[t_1, \dots, t_L]$ obtained from the previous CNN-LSTM encoders for an utterance, we calculate the attention weight between the i th speech vector and the j th text vector as follows:

$$a_{j,i} = \tanh(s_i^\top t_j) \quad (9)$$

$$\alpha_{j,i} = \frac{\exp^{a_{j,i}}}{\sum_{k=1}^K \exp^{a_{j,k}}} \quad (10)$$

$$\tilde{s}_j = \sum_{k=1}^K \alpha_{j,k} s_k \quad (11)$$

where \tilde{s}_j is the weighted summation of the speech vectors, which is considered as an aligned speech vector corresponding to the j th text vector. Finally, we apply an average-pooling operation on the sequence of aligned vectors to obtain the text-aligned speech representation $a_s = \text{average-pooling}([\tilde{s}_1, \dots, \tilde{s}_L])$.

Attention-Aligned Text Representation. We apply a similar process to learn the attention-aligned text representation. Given the text representation $[t_1, \dots, t_L]$ and the speech representation $[s_1, \dots, s_K]$ from the CNN-LSTM encoders for an utterance, we calculate the attention weight between the j th text vector and the i th speech vector as follows:

$$b_{i,j} = \tanh(t_j^\top s_i) \quad (12)$$

$$\beta_{i,j} = \frac{\exp^{b_{i,j}}}{\sum_{l=1}^L \exp^{b_{i,l}}} \quad (13)$$

$$\tilde{t}_i = \sum_{l=1}^L \beta_{i,l} t_l \quad (14)$$

where \tilde{t}_i is the weighted summation of the text vectors, which is considered as an aligned text vector corresponding to the

i th speech vector. Finally, we place an average-pooling layer to obtain the speech-aligned text representation $a_t = \text{average-pooling}(\widehat{[t_1, \dots, t_K]})$.

2.3. Group Gated Fusion Layer

For the speech modality, we have obtained two representations: a_s as the text-aligned speech representation and h_s as the last hidden state of the BiLSTM layer. For the text modality, we have another two similar representations: a_t and h_t . To exploit these representations, we design a novel group gated fusion (GGF) layer to learn the contribution of each representation automatically. Considering different grouped learning processes, we make a_s and a_t as the first group, h_s and h_t as the second group and design the two separate gates corresponding to each group. The first gate controls contributions of the aligned representations a_s and a_t , while the second gate controls contributions of the last hidden states h_s and h_t . The equations governing the GGF layer are as follows:

$$p_s = \tanh(W_s^a a_s) \quad (15)$$

$$p_t = \tanh(W_t^a a_t) \quad (16)$$

$$z_p = \sigma(W_z^p [a_s, a_t]) \quad (17)$$

$$q_s = \tanh(W_s^h h_s) \quad (18)$$

$$q_t = \tanh(W_t^h h_t) \quad (19)$$

$$z_q = \sigma(W_z^q [h_s, h_t]) \quad (20)$$

$$h = \underbrace{z_p \odot p_s + (1 - z_p) \odot p_t}_{\text{Group 1}} + \underbrace{z_q \odot q_s + (1 - z_q) \odot q_t}_{\text{Group 2}} \quad (21)$$

Here, $W_s^a, W_t^a, W_s^h, W_t^h$ are the weights for the non-linear (i.e., \tanh) transformations from a_s, a_t, h_s, h_t respectively. W_z^p, W_z^q are used to learn the contribution of each modality within each group. σ is the sigmoid function and \odot means element-wise product. h is obtained by summing the gated representations of the two groups, and as the final representation for the following emotion classification layer.

2.4. Emotion Classification Layer

Given the obtained representation h , we first apply a fully-connected layer with rectified linear units (ReLUs) for non-linear transformation g , and use a softmax output layer to get \tilde{y} for emotion classification of an utterance. The training objective \mathcal{L} is to minimize the negative log-likelihood, where N is the total samples in training, C is the total number of emotion classes and $y_{i,c} = 1$ if the ground-truth label is c else 0.

$$g = \text{ReLU}(W_g h) \quad (22)$$

$$\tilde{y} = \text{softmax}(W_e g + b) \quad (23)$$

$$\mathcal{L} = -\log \prod_{i=1}^N \sum_{c=1}^C y_{i,c} \log(\tilde{y}_{i,c}) \quad (24)$$

3. Experiments

For model evaluation, we conducted 5-fold cross validation on the Interactive Emotional Dyadic Motion Capture (IEMOCAP) dataset [20], which consists of five sessions with one male and one female speaker each. We used 4 sessions as training set¹

¹We randomly select 5% of the utterances as the validation set.

and the remaining session as testing set to ensure speaker independence. To stay consistent with most previous investigations on IEMOCAP, we use the subset covering the four emotional categories of *happy, angry, sad* and *neutral*.

3.1. Speech and Text Features

We extract both mel-spectrogram and MFCCs as the acoustic features from the speech modality. Each frame in the speech utterance corresponds to a feature vector consisting of 26-dimensional mel-spectrograms, 13-dimensional MFCCs and their first-order deltas, leading to a 52-dimensional vector. Following [21], we set the maximal length of a speech utterance to 7.5s, with longer utterances cut at 7.5s and shorter ones padded with zeros.

For each utterance in the IEMOCAP dataset, there is a corresponding human transcription, which can also be obtained by an automatic speech recognizer [14, 17] in the deployed emotion recognition systems. Both word-level and character-level embeddings can be used to represent the textual transcriptions. We adopt word embeddings to represent each word within an utterance in the IEMOCAP dataset, and initialize the embeddings with the pre-trained 300-dimensional Glove vectors [22].

3.2. Settings and Metrics

We initialize all the network weights in the GBAN model with Xavier normal initializer [23], and use the Adam [24] optimizer by setting the learning rate as 0.0001. To alleviate overfitting, we put a dropout layer [25] with a rate of 0.5 in the GGF layer and before the output softmax layer, and set the coefficient of $L2$ regularization over the network weights as 0.01.

We adopt two widely used metrics for evaluation: weighted accuracy (WA), which is the overall classification accuracy and unweighted accuracy (UA), which is the average recall over the emotion categories. We first compute the metrics for each fold and then present the average accuracy over all the folds.

3.3. Comparison of Representations

Using the CNN-LSTM encoders, the last hidden state of the bidirectional LSTM layer can be used to represent a speech utterance as h_s and the corresponding text as h_t . Adopting the bidirectional alignment network, we obtained the attention-aligned representations a_s for speech and a_t for text. We evaluate their discriminative power in emotion classification using the same classification layer (See Section 2.4). The accuracy (WA) comparison among the representations on the five folds is shown in Table 1.

Table 1: Comparison across representations in emotion classification accuracy (WA) on the IEMOCAP dataset (5 folds).

Fold	h_s	h_t	a_s	a_t
1	0.5860	0.6285	0.6476	0.6762
2	0.6384	0.7060	0.6728	0.7368
3	0.5920	0.5860	0.6150	0.6420
4	0.6936	0.6494	0.6948	0.7289
5	0.6030	0.6497	0.6561	0.6815
Avg	0.6226	0.6439	0.6573	0.6931

We observe that the attention-aligned representations outperform the last-hidden-state representations significantly on all folds for both speech and text. Since both a_s and a_t take additional information from the other modality, this may explain

why they outperform their counterparts. Another interesting observation is that a_t outperforms all other representations. This verifies our hypothesis that the alignment direction is important for the learned representation. The reason that a_t outperforms a_s may be attributed to the different lengths of speech frames and text sequence for the same utterance, where the number of frames can go up to 750 frames while the text sequence may consist of around 20 words. Since the speech sequence is much longer, the attention-aligned speech representation a_s is not as effective as its counterpart a_t . A similar observation is also reported in [21], where the attention mechanism brings slight improvements for speech emotion recognition on the improvised subset of IEMOCAP.

3.4. Comparison of Information Fusion Methods

Given different representations obtained from speech and text, the next interesting question to ask is: “How to make use of all the representations to achieve better performance?” Various information fusion methods can be adopted, such as simple concatenation, tensor fusion network [12] (TFL), gated multimodal units [19] (GMU) and the proposed group gated fusion (GGF) layer. Performance comparisons among these fusion methods on the IEMOCAP dataset are shown in Table 2.

On average, GGF obtains the best accuracy and outperforms all other methods on the folds of 2, 4 and 5. An interesting observation is that both Concat-1 and Concat-2 are strong baselines, although outperformed by TFL and GMU with (a_s, a_t) as the input, and further outperformed by GGF given all the four representations (a_s, a_t, h_s, h_t) . Since the IEMOCAP dataset is relatively small, the simple *concatenation* method may even have enough capacity to learn the underlying patterns. GGF computes a sigmoid weight over the non-linear transformation of the representations within each group, which may be the key to its superior performance to all other methods.

Table 2: Comparison across fusion methods in emotion classification accuracy (WA) on the IEMOCAP dataset (5 folds).

Fold	Input (a_s, a_t)			Input (a_s, a_t, h_s, h_t)	
	Concat-1	TFL	GMU	Concat-2	GGF
1	0.6805	0.6900	0.6985	0.6996	0.6911
2	0.7417	0.7405	0.7552	0.7540	0.7737
3	0.6600	0.6720	0.6890	0.6720	0.6770
4	0.7440	0.7440	0.7591	0.7478	0.7654
5	0.6868	0.7038	0.6975	0.7017	0.7123
Avg	0.7026	0.7101	0.7199	0.7150	0.7239

3.5. Comparison with Existing Approaches

We compare the proposed GBAN model with the existing published approaches in Table 3. They all perform 5-fold cross validation in a speaker-independent manner by using four sessions as training set and the whole remaining session as test set. All the three speech-only models (i.e., CNN-Att [21], LSTM-Att [26] and Self-Att [27]) adopt attention mechanism based on CNN, LSTM and self-attention, respectively. In multimodal emotion recognition, we implement all the models marked with * for fair comparisons using the same experimental settings.

We observe that the multimodal approaches using both speech and text generally outperform the speech-only approaches. Note that the *Att-LSTM* model, proposed by Xu et al. [17] is outperformed by all other models based on bidirectional aligned representations (BiAtt), since it adopts only a unidirectional alignment between speech and text. Finally, the pro-

Table 3: Comparison based on the IEMOCAP dataset (S:Speech, T:Text). All the experiments use 5-fold cross validation and leave one session out as the test set.

Input	Model	WA	UA
S	CNN-Att (Neumann et al. [21])	0.5610	-
	LSTM-Att (Ramet et al. [26])	0.6250	0.5960
	Self-Att (Tarantino et al. [27])	0.6810	0.6380
S+T	Att-LSTM (Xu et al. [17])*	0.6974	0.6401
	BiAtt-Concat*	0.7026	0.6561
	BiAtt-TFL (Zadeh et al. [12])*	0.7101	0.6627
	BiAtt-GMU (Arevalo et al. [19])*	0.7199	0.6584
	GBAN (This Paper)	0.7239	0.7008

posed GBAN model obtains the best performance in terms of both WA and UA on the IEMOCAP dataset.

3.6. Weights Analysis

We further investigate the behavior of the proposed GBAN model by analyzing the weights of z_p and z_q in the GGF layer, which determine the contribution of each representation (i.e., a_s , a_t , h_s and h_t) in emotion classification. Figure 2 shows the average weight of each representation for emotion classification on the 5-fold experiments of the GBAN model. We observe that both a_t and h_t have consistently higher weights within their respective groups (i.e., Group-1 and Group-2) on all the five folds, indicating that the text modality contribute more in emotion classification. This observation is in line with Section 3.3, which shows that both h_t and a_t are more discriminative in emotion classification. Besides, this analysis also illustrates the interpretable benefit of the proposed GGF layer, which allows to present the contribution of a particular modality in multimodal emotion recognition.

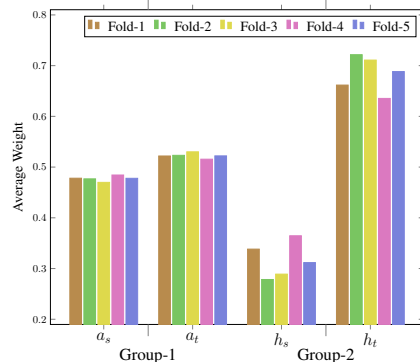


Figure 2: Average weight of each representation in emotion classification on the 5-fold experiments of the GBAN model.

4. Conclusion

This paper presents a model named Gated Bidirectional Alignment Network (GBAN) for multimodal emotion recognition, which consists of a novel attention-based bidirectional alignment network to exploit the alignment information between speech and text explicitly, and a new group gated fusion layer to learn the contribution of each representation from both modalities automatically. We show empirically that the bidirectional alignment network leads to more discriminative representations for emotion classification, and the group gated fusion layer fuses multiple representations effectively in an interpretable manner. GBAN outperforms existing state-of-the-art approaches in emotion classification on the IEMOCAP dataset.

5. 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] K. Han, D. Yu, and I. Tashev, "Speech emotion recognition using deep neural network and extreme learning machine," in *Proc. Interspeech 2014*, 2014, pp. 223–227.
- [3] A. Satt, S. Rozenberg, and R. Hoory, "Efficient emotion recognition from speech using deep learning on spectrograms," in *Proc. Interspeech 2017*, 2017, pp. 1089–1093.
- [4] X. Ma, Z. Wu, J. Jia, M. Xu, H. Meng, and L. Cai, "Emotion recognition from variable-length speech segments using deep learning on spectrograms," in *Proc. Interspeech 2018*, 2018, pp. 3683–3687.
- [5] X. Wu, S. Liu, Y. Cao, X. Li, J. Yu, D. Dai, X. Ma, S. Hu, Z. Wu, X. Liu *et al.*, "Speech emotion recognition using capsule networks," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6695–6699.
- [6] M. Tahon and L. Devillers, "Towards a small set of robust acoustic features for emotion recognition: challenges," *IEEE/ACM transactions on Audio, Speech, and Language Processing*, vol. 24, no. 1, pp. 16–28, 2015.
- [7] A. Barbulescu, R. Ronfard, and G. Bailly, "Which prosodic features contribute to the recognition of dramatic attitudes?" *Speech Communication*, vol. 95, pp. 78–86, 2017.
- [8] G. Aguilar, V. Rozgic, W. Wang, and C. Wang, "Multimodal and multi-view models for emotion recognition," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 991–1002.
- [9] T. Bänziger, D. Grandjean, and K. R. Scherer, "Emotion recognition from expressions in face, voice, and body: the Multimodal Emotion Recognition Test (MERT)," *Emotion*, vol. 9, no. 5, p. 691, 2009.
- [10] 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.
- [11] V. Heusser, N. Freymuth, S. Constantin, and A. Waibel, "Bimodal speech emotion recognition using pre-trained language models," *arXiv preprint arXiv:1912.02610*, 2019.
- [12] A. Zadeh, M. Chen, S. Poria, E. Cambria, and L.-P. Morency, "Tensor fusion network for multimodal sentiment analysis," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Sep. 2017, pp. 1103–1114.
- [13] Z. Aldeneh, S. Khorram, D. Dimitriadis, and E. M. Provost, "Pooling acoustic and lexical features for the prediction of valence," in *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, 2017, pp. 68–72.
- [14] 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.
- [15] S. Poria, N. Majumder, D. Hazarika, E. Cambria, A. Gelbukh, and A. Hussain, "Multimodal sentiment analysis: Addressing key issues and setting up the baselines," *IEEE Intelligent Systems*, vol. 33, no. 6, pp. 17–25, 2018.
- [16] N. Majumder, D. Hazarika, A. Gelbukh, E. Cambria, and S. Poria, "Multimodal sentiment analysis using hierarchical fusion with context modeling," *Knowledge-Based Systems*, vol. 161, pp. 124–133, 2018.
- [17] H. Xu, H. Zhang, K. Han, Y. Wang, Y. Peng, and X. Li, "Learning alignment for multimodal emotion recognition from speech," in *Proc. Interspeech 2019*, 2019, pp. 3569–3573.
- [18] D. Hazarika, S. Poria, A. Zadeh, E. Cambria, L.-P. Morency, and R. Zimmermann, "Conversational memory network for emotion recognition in dyadic dialogue videos," in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 2018, pp. 2122–2132.
- [19] J. Arevalo, T. Solorio, M. Montes-y Gómez, and F. A. González, "Gated multimodal units for information fusion," *arXiv preprint arXiv:1702.01992*, 2017.
- [20] 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.
- [21] 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," *Proc. Interspeech 2017*, pp. 1263–1267, 2017.
- [22] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 2014, pp. 1532–1543.
- [23] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics*, 2010, pp. 249–256.
- [24] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015.
- [25] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, no. 56, pp. 1929–1958, 2014.
- [26] 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.
- [27] L. Tarantino, P. N. Garner, and A. Lazaridis, "Self-attention for speech emotion recognition," in *Proc. Interspeech 2019*, 2019, pp. 2578–2582.