

# KNOWLEDGE-AWARE BAYESIAN CO-ATTENTION FOR MULTIMODAL EMOTION RECOGNITION

Zihan Zhao<sup>1</sup>, Yu Wang<sup>1,2\*</sup>, Yanfeng Wang<sup>1,2\*</sup>

<sup>1</sup>Cooperative Medianet Innovation Center, Shanghai Jiao Tong University <sup>2</sup>Shanghai AI Laboratory  
{zihanzhao, yuwangsjtu, wangyanfeng}@sjtu.edu.cn

## ABSTRACT

Multimodal emotion recognition is a challenging research area that aims to fuse different modalities to predict human emotion. However, most existing models that are based on attention mechanisms have difficulty in learning emotionally relevant parts on their own. To solve this problem, we propose to incorporate external emotion-related knowledge in the co-attention based fusion of pre-trained models. To effectively incorporate this knowledge, we enhance the co-attention model with a Bayesian attention module (BAM) where a prior distribution is estimated using the emotion-related knowledge. Experimental results on the IEMOCAP dataset show that the proposed approach can outperform several state-of-the-art approaches by at least 0.7% unweighted accuracy (UA).

**Index Terms**— multimodal emotion recognition, transfer learning, Bayesian attention, knowledge injection

## 1. INTRODUCTION

Emotion recognition is the process of identifying human emotion. Multi-modality for emotion recognition has recently received great interest because different modalities can provide complementary clues and the fusion of them can effectively improve the performance [1, 2, 3]. For example, the meanings of words and their relation express one's emotion in text modality, while intonation and pitch can also be useful for recognizing emotions in speech modality. In this paper, we focus our research on text and speech modalities.

Emotion datasets usually face the problem of insufficient data [4], mainly due to the difficulties caused by the subjective nature of labeling emotions. The common solution to solve this problem is to leverage pre-trained models. For the text modality, BERT [5] is the most commonly used model and is well suited for feature extraction. Recently, speech-based pre-trained models have also emerged, and the most state-of-the-art of them is the wav2vec 2.0 [6], which has shown

promising performance in ASR [7], speaker verification [8] and emotion recognition [1].

Attention mechanisms enable models to learn which parts are important and which parts are irrelevant [9]. So it is well suited for handling sequential inputs, such as text and speech. Self-attention is an attention mechanism relating different positions of a single sequence to compute a representation of the sequence [9]. Based on attention mechanisms, co-attention is a fusion method in which the representation of the sequence of a modality is computed by relating the sequence of the other modality [10]. Several works have proposed to leverage pre-trained models of both speech and text modalities and attention mechanisms for multimodal emotion recognition and achieved promising results [1, 2, 11, 12].

However, not all parts in the data sequences are related to emotion, and without the guidance of external knowledge, it is challenging for the attention to learn emotionally relevant parts accurately on its own [13]. To solve this problem, we propose to incorporate external emotion-related knowledge. More specifically, an emotion lexicon is leveraged with the attention mechanisms to improve its interpretability and robustness for the emotion recognition task. The emotion lexicon has shown its effectiveness in the previous methods. The work in [14] used commonsense knowledge to build a graph and retrieved emotion-related information for the nodes in it using an emotion lexicon. Some other works proposed to integrate an emotion lexicon with contextual information for emotion recognition in conversations [15]. However, there is no current work that combines the emotion lexicon with attention mechanisms. In this paper, we propose a Bayesian framework [16] with an emotion knowledge-related attention map as the prior distribution. The posterior distribution of the attention weights can then be learned. The reason that a Bayesian attention module (BAM) is introduced instead of simply adding the knowledge-related attention map on the original one is that the knowledge we introduce is not collected on the dataset we use, which may introduce bias. If a rather "hard" way is used for incorporating the knowledge, it may affect the learning of the attention model. The approach we propose allows the introduction of prior knowledge while achieving an integrated training of the attention weights, thus preventing this problem. We also show that the combination

\* Corresponding authors

This work was supported by National Key R&D Program of China (No.2022ZD0162100), National Natural Science Foundation of China (No.62106140), Shanghai Science and Technology Committee (No.21511101100) and Shanghai Key Lab of Digital Media Processing and Transmission (No.18DZ2270700)



**Fig. 1.** The proposed model. Word-level force alignment is applied for using a word-level emotion lexicon. Knowledge-aware Bayesian co-attention is used to fuse two modalities and also inject knowledge.

with the late fusion model further improves the performance.

We summarize our major contributions as follows:

- We introduce emotion-related knowledge in the attention model to help the model focus more on relevant information for emotion recognition.
- To effectively incorporate knowledge, we further propose to use BAM, which not only brings randomness in the model and thus can help in modeling complex dependencies but also incorporates knowledge as prior in a "soft" way. We also show that it's complementary to the late fusion model.
- We evaluate the proposed models on the popular IEMOCAP dataset [17]. Experimental results show that they can outperform existing multimodal approaches using both speech and text modalities.

## 2. PROPOSED METHODS


The proposed model is shown in Figure 1. First, we use BERT to extract text embeddings and wav2vec 2.0 to extract frame-level speech embeddings. Then we use word-level force alignment to get word-level speech embeddings in Section 2.1. Next, we apply self-attention to both embeddings to process input sequences and soften knowledge in Section 2.2. We introduce the general co-attention module in Section 2.3 and the proposed knowledge-aware co-attention module to fuse two modalities in Section 2.4. Average pooling layers and linear layers are then used to output the probability distribution for emotion classification.

### 2.1. Word-level Force Alignment

The knowledge we adopt is of word-level, and we want to introduce knowledge for both text and speech modalities, so we use the word-level force alignment to cope with frame-level speech embeddings:

$$\mathbf{u}_k^W = \frac{\sum_{f=s_k}^{e_k} \mathbf{u}_f}{e_k - s_k}, \quad (1)$$

where  $\mathbf{u}_k^W$  is the  $k^{\text{th}}$  segmentation of the word-level speech embedding  $\mathbf{U}^W$ ,  $\mathbf{u}_f$  is the  $f^{\text{th}}$  frame of the frame-level speech embedding,  $s$  is the starting frame,  $e$  is the end frame. Segment-level speech embeddings like word-level speech embeddings are closely related to the prosody, and prosody can convey characteristics of the utterance like emotional state because it contains the information of the cadence of speech signals [1].



**Fig. 2.** Co-attention module. Multi-head attention is used to fuse two modalities.

### 2.2. Self-attention Module

Self-attention is an attention mechanism relating different positions of a single sequence to compute a representation of the sequence [9]. Without loss of generality, the computation of one-head attention is given here:

$$\{\mathbf{Q}_s, \mathbf{K}_s, \mathbf{V}_s\} = \mathbf{W}_{\{1,2,3\}}^i \mathbf{U}^i + b_{\{1,2,3\}}^i, \quad (2)$$

where  $\mathbf{U}^i$  is the embedding for modality  $i \in \{W, T\}$ ,  $W$  is for word-level speech embeddings and  $T$  is for text embeddings.  $\mathbf{W}$  is the weight and  $b$  is the bias. In the self-attention module,  $i$  in Equation 2 are the same for  $\mathbf{Q}_s$ ,  $\mathbf{K}_s$  and  $\mathbf{V}_s$ . Thus, we can obtain the processed embeddings as:

$$\mathbf{M} = \text{softmax} \left( \frac{\mathbf{Q}_s \mathbf{K}_s^\top}{\sqrt{d}} \right) \quad (3)$$


$$\text{Attention} = \mathbf{M} \mathbf{V}, \quad (4)$$

where  $d$  represents the dimension of the embeddings. Multi-head attention performs this process multiple times to learn information from various representation subspaces. We also use the self-attention module to soften the knowledge from the NRC VAD lexicon [18], which is the emotion lexicon we use. Because the knowledge is relatively sparse, without this step, the model will have difficulty learning useful information from the knowledge. The process of softening the knowledge is as follows: Given a sequence, we first look up the lexicon to find words that appeared in the lexicon. The VAD values we use range from -1 to 1. Then we compute the  $L^2$ -norm of each word to get its intensity:  $i_k = \sqrt{v_k^2 + a_k^2 + d_k^2}$ , where  $i_k$  is the  $k^{\text{th}}$  of vector  $\mathbf{i}$  which is the intensity for the sequence and  $v, a, d$  are for valence, arousal, dominance respectively. Intensity is assigned to 0 for words that do not appear in the lexicon. Now we get the knowledge  $\mathbf{i}$  for the sequence, but it faces the problem of sparsity. So we dot product the  $k^{\text{th}}$  row of the attention map in Equation 3 with the  $\mathbf{i}$  to get the softening intensity  $i_k^{\text{soft}}$  which is the  $k^{\text{th}}$  of vector  $\mathbf{i}^{\text{soft}}$ :

$$i_k^{\text{soft}} = \mathbf{M}_k \mathbf{i}. \quad (5)$$

### 2.3. Co-attention Module

The co-attention module [10] is depicted in Figure 2. There are two branches in the co-attention module, each of which has the same structure as the Transformer encoder [9]. First embeddings from two modalities are sent to multi-head attention respectively. The computation of  $\mathbf{Q}_c$ ,  $\mathbf{K}_c$  and  $\mathbf{V}_c$  is



**Fig. 3.** Knowledge aware Bayesian attention during training.  $i_{Q_c/K_c}^{soft}$  - intensity for the sequence which computes  $Q_c/K_c$  and is used to compute prior distribution. Reparameterization is used for the backpropagation of gradients.

the same as Equation 2, but here each branch has different modalities as  $Q_c$  or  $K_c, V_c$ . Then attention map is computed using Equation 3. The output of multi-head attention is computed using Equation 4 and then added with its input which is used to compute its  $Q_c$ . Layernorm layers are added to help the model converge [19]. The feed-forward layer is then applied, and it consists of two linear layers with a nonlinear function. The output of the feed-forward layer is added with its input and then Layernorm is applied. Co-attention module can be applied several times. The output is the mean of the two branches of the final co-attention module.

#### 2.4. Knowledge-aware Bayesian Co-attention Module

The proposed Knowledge-aware Bayesian Co-attention Module has the same structure as the co-attention module in Section 2.3 except the computation of the attention map in Equation 3 which is shown in Figure 3. In the BAM [16], attention weights  $W$  are not deterministic but data-dependent local random variables from a variational distribution. It was shown that this variational distribution can approximate the posterior of attention weights  $W$  using reparameterizable attention distributions and also a prior as regularization. Following [16], we use a Weibull distribution as the variational distribution and Gamma distribution as the prior distribution, then the regularization KL divergence can be computed as:

$$\begin{aligned} \text{KL}(\text{Weibull}(k, \lambda) \parallel \text{Gamma}(\alpha, \beta)) &= \frac{\gamma\alpha}{k} - \alpha \log \lambda \\ &+ \log k + \beta\lambda\Gamma\left(1 + \frac{1}{k}\right) - \gamma - 1 - \alpha \log \beta + \log \Gamma(\alpha), \end{aligned} \quad (6)$$

where  $k, \lambda$  are the parameters of Weibull distribution,  $\alpha, \beta$  are the parameters of Gamma distribution,  $\gamma$  is the Euler's constant, and  $\Gamma$  is the gamma function. The original attention map of the co-attention module computed using Equation 3 is used to compute  $\lambda$  of Weibull distribution. The sample of the Weibull distribution after normalization is the attention map of the BAM during training. During inference, the posterior expectations are used to obtain the point esti-

mates [20]. However, sampling directly from the distribution will fail backpropagation of gradients. Following [16], the Weibull distribution is a reparameterizable distribution, so this problem can be solved by sampling from the standard uniform distribution  $\varepsilon \sim \text{Uniform}(0, 1)$ . Then the sample from Weibull distribution is equivalent to drawing:  $S := \lambda(-\log(1 - \varepsilon))^{1/k}$ , where  $\lambda, k$  are the parameter of Weibull distribution. The sample  $S$  is multiplied with  $V_c$  to get the output of the knowledge-aware Bayesian attention, similar to Equation 4.

Unlike [16], which uses  $K_c$  to compute parameters of the prior distribution in Equation 6, we use the knowledge vector  $i_{Q_c}^{soft}, i_{K_c}^{soft}$  in Equation 5 to compute them where  $i_{Q_c/K_c}^{soft}$  is the intensity for the sequence which computes  $Q_c/K_c$ :  $P = \text{softmax}(i_{Q_c}^{soft \top} i_{K_c}^{soft})$ . We use  $P$  to compute  $\alpha$  of the Gamma distribution in Equation 6. In this way, prior knowledge can be introduced while still leaving the attention weights to be learned by the model itself, which alleviates the problem of bias in the emotion lexicon. In this way the  $k^{\text{th}}$  of the intensity  $i_{Q_c, k}^{soft}$  works like temperature for  $i_{K_c}^{soft}$  in every row. If  $i_{Q_c, k}^{soft}$  is high, then this word is highly related to emotion, and the temperature will strengthen its connection with other emotional words. If  $i_{Q_c, k}^{soft}$  is low, a preposition for example, then this word may not be related to emotion, and the temperature will weaken its connection with other emotional words. The  $\beta$  of the Gamma distribution and the  $k$  of the Weibull distribution remain as hyperparameters.

### 3. EXPERIMENTS AND DISCUSSION

#### 3.1. Datasets

The IEMOCAP dataset is an acted, multimodal and multi-speaker dataset with five sessions [17]. We use 4 emotional classes: anger, happiness, sadness, and neutral. The number of utterances representing them was 1103, 1636, 1084, and 1708 respectively. Alignment information of IEMOCAP is provided in their datasets. Other datasets can also get alignment information using force alignment tools such as Prosody Aligner [21] and Montreal Forced Aligner [22]. We use the NRC VAD lexicon [18] as the emotion lexicon. It includes a list of more than 20,000 English words and their valence (positiveness-negativeness), arousal (active-passive), and dominance (dominant-submissive) scores.

#### 3.2. Settings and Metrics

Cross-entropy loss plus KL loss in Equation 6 was used as the loss function, and an Adam optimizer [23] was applied using a learning rate of 1e-4. The models were trained using a batch size of 32 and early stopping was applied with a patience of 6. Dropout [20] was applied with a probability of 0.1 after every feed-forward layer except the output layer to prevent overfitting. We use wav2vec 2.0-base and BERT-base un-cased models, both of which have 768-dimensional embeddings. Following [6], we extract embeddings from all 12 transformer encoder layers in wav2vec 2.0, and we also

**Table 1.** Results of the single modality and multimodality. T - text, F - frame-level speech, W - word-level speech, Self-att - self-attention, Co-att - co-attention, BAM - Bayesian attention module

Modalities	Models	UA	WA
T	Self-att	66.6%	65.2%
F	Self-att	63.8%	62.2%
W	Self-att	63.4%	61.7%
T+F	Co-att	74.8%	73.5%
T+W	Co-att	74.9%	73.6%
T+W	Co-att+knowledge	74.6%	73.3%
T+W	Co-att+BAM	74.8%	73.7%
<b>T+W</b>	<b>Co-att+BAM+knowledge</b>	<b>75.6%</b>	<b>74.0%</b>

apply the weighted average for different layers to get the final word-level speech embeddings. Similarly, we also apply the weighted average for the text modality. All speech samples are normalized by global normalization which is a frequently used setting for this dataset.  $k$  of Weibull distribution was set to 1 and  $\beta$  of Gamma distribution was set to 10. Five-fold cross-validation was repeated 5 times and the results were the average of them. A frame length of 400 is used. We observed that more layers of the self-attention module and co-attention module didn't bring improvement in all models, so they were set to 1. Unweighted accuracy (UA) and weighted accuracy (WA) were used as our evaluation metrics.

### 3.3. Results of the Single Modality

The first part of Table 1 shows the results of the single modality. They have the same structure as one branch of Figure 1 except that there is no co-attention module. Word-level speech embeddings are computed from frame-level speech embeddings using Equation 1. Text modality has the best performance. This is because in IEMOCAP there are not many complex situations where speech is more effective, such as sarcasm. We can also observe that using the word-level alignment information still has similar results compared to the complete frame-level speech embeddings, but with a much smaller size.

### 3.4. Results of the Multimodality

The second part of Table 1 shows the results of multimodality. Compared to the results of the single modality, we can see that the multimodality can yield good performance improvements which are consistent with that observed in [1, 2]. It can also be seen that the word-level speech embedding-based co-attention fusion gives a similar performance as the frame-level one. Next, the results of word-level co-attention with a knowledge map directly added to the original attention map (Co-att+knowledge) are shown. It's a rather "hard" way which is shown to degrade the performance slightly. We show the results of word-level co-attention with BAM (Co-att+BAM) next, in which we use the prior [16] computed using  $\mathbf{K}_s$  in Equation 2. We can observe that incorporating the BAM alone without knowledge cannot improve the

**Table 2.** Results of the combination with late fusion

Models	UA	WA
Late fusion w/o Co-att	75.7%	74.2%
Late fusion+Co-att	76.3%	75.2%
<b>Late fusion+Co-att+BAM+knowledge</b>	<b>77.0%</b>	<b>75.5%</b>

**Table 3.** Comparison with existing methods

Models	UA	WA
Chen et al. [3], 2020	72.1%	71.1%
Makiuchi et al. [24], 2021	73.0%	73.5%
Kumar et al. [25], 2021	75.0%	71.7%
Chen et al. [2], 2022	75.3%	74.3%
Zhao et al. [1], 2022	76.3%	-
<b>Ours</b>	<b>77.0%</b>	<b>75.5%</b>

model performance. This may be because it's challenging to train BAM to fuse two modalities for multimodal emotion recognition without external guidance. Finally, with the BAM and prior estimated using the emotion-related knowledge (Co-att+BAM+knowledge), it yields the best results among all the co-attention-based models, which validates the effectiveness of the proposed knowledge-enhance BAM model.

### 3.5. Results of the Combination with Late Fusion

In [1], they show that the co-attention model can combine with the late fusion model simply with score fusion to boost the performance. Table 2 shows that this combination can also give performance gains for our knowledge-aware co-attention. The late fusion model that fuses text and word-level speech embeddings has the same structure as that in [1]. The late fusion model that fuses text and frame-level speech embeddings has a similar performance and is omitted here. We can see that late fusion with the knowledge-aware Bayesian co-attention model has the best performance, yielding 77% UA and 75.5% WA, which demonstrates the complementarity of the late fusion and knowledge-aware Bayesian co-attention.

### 3.6. Comparison with Existing Methods

We compare the proposed method with state-of-the-art multimodal emotion recognition methods in Table 3. For a fair comparison, all the experiments use the text and speech data from IEMOCAP, 5-fold cross-validation, and four classes. It shows that our proposed model achieves the best performance and surpasses other methods by at least 0.7% UA.

## 4. CONCLUSION

In this paper, we have proposed to incorporate emotion-related knowledge in Bayesian co-attention modules for multimodal emotion recognition. Experimental results have shown that the proposed Bayesian co-attention model can outperform the baseline multimodal emotion recognition methods and achieves an accuracy of 77.0% UA and 75.5% WA on the 5-fold CV on IEMOCAP.

## 5. REFERENCES

- [1] Z. Zhao, Y. Wang, and Y. Wang, “Multi-level fusion of wav2vec 2.0 and BERT for multimodal emotion recognition,” in *Proc. Interspeech*, 2022, pp. 4725–4729.
- [2] W. Chen, X. Xing, X. Xu, and et al., “Key-sparse transformer for multimodal speech emotion recognition,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 6897–6901.
- [3] M. Chen and X. Zhao, “A multi-scale fusion framework for bimodal speech emotion recognition,” in *Proc. Interspeech*, 2020, pp. 374–378.
- [4] L. Pepino, P. Riera, and L. Ferrer, “Emotion Recognition from Speech Using Wav2vec 2.0 Embeddings,” in *Proc. Interspeech*, 2021, pp. 3400–3404.
- [5] J. Devlin, M. W. Chang, K. Lee, and et al., “BERT: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [6] A. Baevski, Y. Zhou, A. Mohamed, and et al., “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 12449–12460, 2020.
- [7] Y. Zhang, J. Qin, D. S. Park, and et al., “Pushing the limits of semi-supervised learning for automatic speech recognition,” *arXiv preprint arXiv:2010.10504*, 2020.
- [8] Z. Fan, M. Li, S. Zhou, and et al., “Exploring wav2vec 2.0 on speaker verification and language identification,” *arXiv preprint arXiv:2012.06185*, 2020.
- [9] A. Vaswani, N. Shazeer, N. Parmar, and et al., “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [10] J. Lu, D. Batra, D. Parikh, and et al., “ViLBERT: Pre-training task-agnostic visiolinguistic representations for vision-and-language tasks,” *Advances in neural information processing systems*, vol. 32, 2019.
- [11] S. Schneider, A. Baevski, R. Collobert, and et al., “wav2vec: Unsupervised Pre-Training for Speech Recognition,” in *Proc. Interspeech*, 2019, pp. 3465–3469.
- [12] Y. Liu, M. Ott, N. Goyal, and et al., “RoBERTa: A robustly optimized BERT pretraining approach,” *arXiv preprint arXiv:1907.11692*, 2019.
- [13] J. Bai, Y. Wang, H. Sun, and et al., “Enhancing self-attention with knowledge-assisted attention maps,” in *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2022, pp. 107–115.
- [14] Y. Fu, S. Okada, L. Wang, and et al., “CONSK-GCN: conversational semantic-and knowledge-oriented graph convolutional network for multimodal emotion recognition,” in *IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 2021, pp. 1–6.
- [15] G. Tu, J. Wen, C. Liu, and et al., “Context- and sentiment-aware networks for emotion recognition in conversation,” *IEEE Transactions on Artificial Intelligence*, vol. 3, no. 5, pp. 699–708, 2022.
- [16] X. Fan, S. Zhang, B. Chen, and et al., “Bayesian attention modules,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 16362–16376, 2020.
- [17] C. Busso, M. Bulut, C. C. Lee, and et al., “IEMOCAP: Interactive emotional dyadic motion capture database,” *Language resources and evaluation*, vol. 42, no. 4, pp. 335–359, 2008.
- [18] S. Mohammad, “Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 english words,” in *Proceedings of the 56th annual meeting of the association for computational linguistics (volume 1: Long papers)*, 2018, pp. 174–184.
- [19] J. L. Ba, J. R. Kiros, and G. E. Hinton, “Layer normalization,” *arXiv preprint arXiv:1607.06450*, 2016.
- [20] N. Srivastava, G. Hinton, A. Krizhevsky, and et al., “Dropout: a simple way to prevent neural networks from overfitting,” *The journal of machine learning research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [21] K. Gorman, J. Howell, and M. Wagner, “Prosodylab-aligner: A tool for forced alignment of laboratory speech,” *Canadian Acoustics*, vol. 39, pp. 192, 2011.
- [22] M. McAuliffe, M. Socolof, S. Mihuc, and et al., “Montreal forced aligner: Trainable text-speech alignment using kald,” in *Proc. Interspeech*, 2017, pp. 498–502.
- [23] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [24] M. R. Makiuchi, K. Uto, and K. Shinoda, “Multimodal emotion recognition with high-level speech and text features,” in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 350–357.
- [25] P. Kumar, V. Kaushik, and B. Raman, “Towards the explainability of multimodal speech emotion recognition,” in *Proc. Interspeech*, 2021, pp. 1748–1752.