

Advancing Multimodal Teacher Sentiment Analysis: The Large-Scale T-MED Dataset & The Effective AAM-TSA Model

Zhiyi Duan¹, Xiangren Wang¹, Hongyu Yuan¹, Qianli Xing^{2*}

¹Department of Computer Science, Inner Mongolia University, Hohhot, Inner Mongolia

²College of Computer Science and Technology, Jilin University, Changchun, Jilin

duanzy@imu.edu.cn, wangxr@mail.imu.edu.cn, yuanhongyu.1997@163.com, qianlixing@jlu.edu.cn

Abstract

Teachers' emotional states are critical in educational scenarios, profoundly impacting teaching efficacy, student engagement, and learning achievements. However, existing studies often fail to accurately capture teachers' emotions due to the performative nature and overlook the critical impact of instructional information on emotional expression. In this paper, we systematically investigate teacher sentiment analysis by building both the dataset and the model accordingly. We construct the first large-scale teacher multimodal sentiment analysis dataset, T-MED. To ensure labeling accuracy and efficiency, we employ a human-machine collaborative labeling process. The T-MED dataset includes 14,938 instances of teacher emotional data from 250 real classrooms across 11 subjects ranging from K-12 to higher education, integrating multimodal text, audio, video, and instructional information. Furthermore, we propose a novel asymmetric attention-based multimodal teacher sentiment analysis model, AAM-TSA. AAM-TSA introduces an asymmetric attention mechanism and hierarchical gating unit to enable differentiated cross-modal feature fusion and precise emotional classification. Experimental results demonstrate that AAM-TSA significantly outperforms existing state-of-the-art methods in terms of accuracy and interpretability on the T-MED dataset.

Introduction

In the intricate scenarios of modern education, the emotional states of teachers emerge as a pivotal factor, profoundly influencing not only the efficacy of pedagogical delivery but also the crucial aspects of student engagement, motivation, and ultimately, learning achievements (Han, Jin, and Yin 2023). Consequently, the systematic analysis of teacher emotions has garnered increasing attention within educational psychology and human-computer interaction research (Wang and Frenzel 2025).

Unlike general sentiment analysis, analyzing teacher emotions is distinguished by several unique characteristics. On the one hand, the performative nature of teaching requires educators to project specific emotions, such as enthusiasm or composure. Thus, teachers frequently exhibit hidden emotions, deliberately suppressing feelings like frustration or stress to uphold professionalism or shield students

from negative influences (King et al. 2024). On the other hand, teacher emotions are inherently affected by instructional information, including instructional subject and educational stage (Hargreaves 2000).

The current approaches to teachers' sentiment analysis have predominantly relied on single-modal data modeling, primarily focusing on textual transcripts or isolated audio cues (Zhao, Zhang, and Chu 2024). While these methods have offered valuable preliminary insights, their inherent limitations become apparent when confronted with the complexity and richness of real-world teaching environments. The inadequacy of single-modal data is further amplified by the distinctive nature of teachers' emotional expressions, characterized by deliberate deductiveness, e.g., emotive cues tailored to instructional goals, and strategic restraint, e.g., suppression of negative emotions to maintain classroom dynamics. In summary, rendering any single data modality is insufficient to disentangle performative displays from underlying emotion states.

Teacher sentiment analysis is inherently multimodal, manifested through a confluence of text, audio, video, and instructional information (Cheng et al. 2024). This integration can effectively capture environment-induced variations in emotional expression and decode implicit affective cues, which are critical for understanding teachers' true emotional experiences. Although current multimodal models try to integrate information from different modalities (Zhao, Zhang, and Chu 2024; Zheng et al. 2024), they often fail to effectively analyze teachers' emotions due to the performative nature. On the other hand, the existing multimodal datasets (McDuff and Soleymani 2017; Bagher Zadeh et al. 2018) are mostly designed for general sentiment analysis, which further interferes with the performance of the sentiment analysis models. Thus, there is an urgent need for both a high-quality multimodal dataset and specialized multimodal teacher sentiment analysis models.

To bridge this critical gap and facilitate a more holistic and accurate understanding of teacher emotional dynamics, this paper presents a systematic investigation into teacher sentiment analysis, encompassing both data construction and model development. We first present a novel human-machine collaboration annotation collaborative framework, and collect a large-scale dataset named T-MED, containing four modalities. We specially designed eight sentiment

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

labels, including three types of teacher-specific emotions: patience, enthusiasm, and expectation (Zhang et al. 2024; Frenzel et al. 2018; Rubie-Davies et al. 2023). Then, we propose a novel asymmetric attention-based multimodal teacher sentiment analysis model, AAM-TSA. We have crafted an audio-focused asymmetric attention mechanism to seamlessly unify multimodal features into a cohesive representation, while maintaining unique modality-specific traits and capturing inter-modal semantic relationships. In summary, our contributions are as follows:

- We construct a large-scale teacher multimodal sentiment analysis dataset, T-MED, for the first time. The T-MED dataset contains 14,938 samples of 4 modalities across 8 types of teacher emotional data, spanning over 17 hours in total and encompassing multidisciplinary scenarios from K-12 to higher education.
- We propose a novel asymmetric attention-based multimodal teacher sentiment analysis model, AAM-TSA. From the perspectives of multimodal information learning and fusion, AAM-TSA deeply alleviates the deductive problem in teacher’s emotional analysis.
- The experimental results show that the collected T-MED dataset and the proposed AAM-TSA are effective and reasonable. AAM-TSA is superior to the state-of-the-art multimodal sentiment analysis methods on the T-MED.

Related Work

Teachers’ Sentiment Analysis

Teacher sentiment analysis has undergone a significant evolution, transitioning from single-modal to multimodal approaches in response to the growing demand for a more comprehensive understanding of emotional expression in educational contexts (Li 2022; Shaikh and Doudpotta 2019; Lu et al. 2024). For instance, Zhang et al. extracted acoustic spectrum features, including Mel-frequency cepstral coefficient (MFCC) features and filter bank (FBank) features, utilizing parallel subnet training structures for feature fusion (Zhang and Li 2022).

Consequently, there has been a notable shift towards multimodal frameworks that integrate data from multiple sources. ProsodyBERT designed a parallel coding architecture to process pronunciation, MFCC and text features (Zhao, Zhang, and Chu 2024). AMTBA proposed a classroom event segmentation method based on voice activity detection (VAD), emphasizing the structured information processing of classrooms (Zheng et al. 2024). Despite these advancements, current multimodal methods are constrained by several critical limitations. A primary challenge lies in the effective fusion of heterogeneous data types, which necessitates sophisticated models capable of managing the complexity and variability inherent in multimodal inputs. Moreover, the development and evaluation of these models are significantly hindered by the scarcity of large-scale, annotated multimodal datasets, which are essential for training robust teachers’ sentiment analysis models.

Datasets for Teachers’ Sentiment Analysis

It is necessary to construct the domain-specific teacher sentiment analysis datasets as the general multimodal datasets for sentiment analysis (McDuff and Soleymani 2017; Bagher Zadeh et al. 2018) are hard to apply to teachers’ sentiment analysis due to the performative nature of teachers. He et al. constructed 1,825 pure audio clips, directly pointing to the core modality of teacher sentiment analysis (He et al. 2024). AMTBA dataset contained 2,109 voice emotion data for classroom voice emotion performance testing (Zheng et al. 2024). TESD dataset was constructed with a total of 3564 emotional voice samples from 6 male and 6 female teachers, which includes four emotional categories: anger, surprise, neutrality, and sadness (Cai et al. 2024). These teacher sentiment analysis datasets often face limitations such as small scale, limited modality, and restricted access due to privacy concerns, which hinder their broader research and application.

Human-Machine Collaboration Annotation Framework for T-MED Dataset

To address the critical need for high-quality datasets in teachers’ sentiment analysis, we present a novel human-machine collaboration annotation collaborative framework, yielding a high-quality dataset named T-MED with reliable sentiment labels. Our human-LLM annotation collaborative framework integrates iterative model refinement and expert oversight across six sequential steps as shown in Figure 1. A comprehensive explanation of each step is provided below. Note that, we specially annotate three distinct sentiment categories: patience, enthusiasm, and expectation to capture the characteristics of teacher emotions. The detailed information of T-MED is provided in Appendix A.

Step 1: Multimodal Data Acquisition from MOOC Resources. A total of 250 high-quality videos encompassing educational scenarios ranging from K-12 to higher education, are selected and downloaded from public MOOC platforms from multiple providers. To extract teacher-centric multimodal cues, open-source tools are employed for systematic processing: FFmpeg (Bellard and Developers 2000) for segmenting and preprocessing raw video files, Faster Whisper (Klein and Systran 2023) for generating accurate text transcripts from speech, and the Speaker Identification Model CAM++ (Alibaba DAMO Academy 2024) for isolating teacher-specific audio streams.

Finally, we obtain four core modalities: video clips, audio tracks, textual transcripts, and instructional information. Specifically, instructional information is a text-based modality encompassing two key dimensions: subject information and grade-level information, which are vital for contextualizing emotional expressions. Note that, eight types of emotional labels are finally identified: neutral, anger, joy, surprise, sadness, patience, enthusiasm, and expectation. The former five sentiments are common emotions in the general field, while patience, enthusiasm, and expectation are more reflective of professional characteristics specific to teachers.

Step 2: Pilot Annotation with Pre-trained Models. A subset of the multimodal data (10% of the total corpus) is

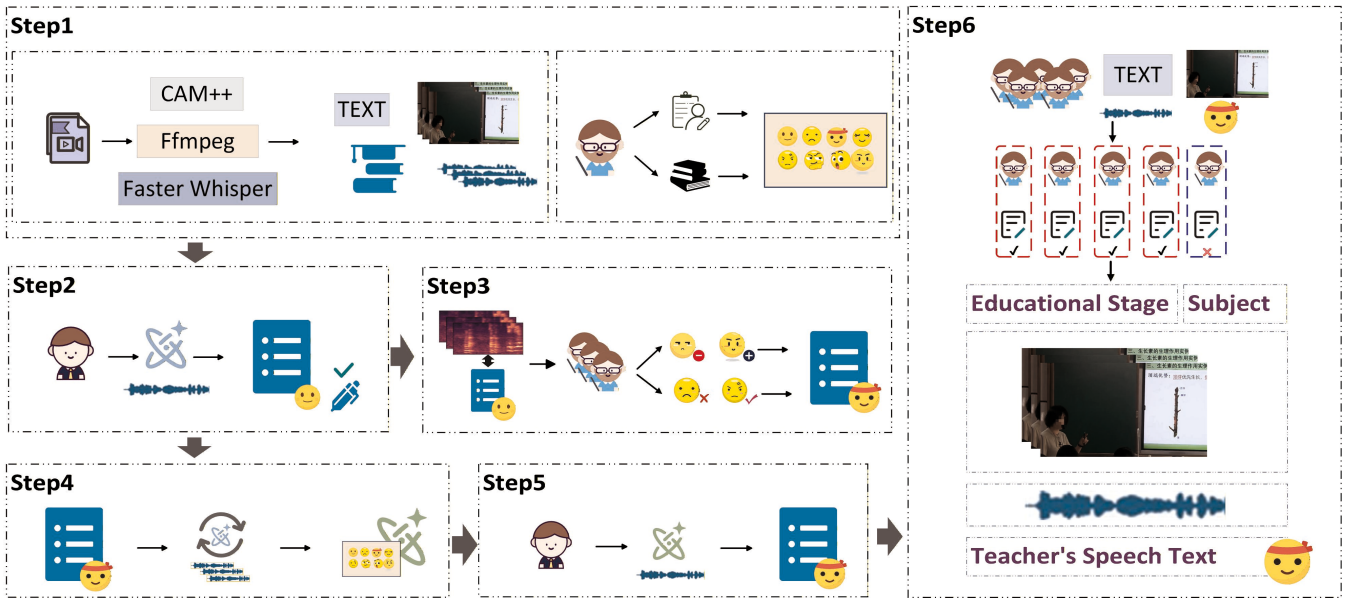


Figure 1: Diagram of T-MED dataset construction process.

subjected to initial emotion labeling via a pre-trained model emotion2vec-large (Ma et al. 2024). The purpose of this step is to generate preliminary emotional labels for subsequent expert review and model refinement while reducing the manual annotation burden for the larger corpus.

Step 3: Expert Validation and Correction of Pilot Annotations. A group of three experts is enlisted to conduct manual validation of the pre-labeling results generated in step 2. The experts collaboratively analyzed the initial pre-labeling outputs by cross-referencing audio clips with their corresponding Mel spectrograms. Using eight predefined emotional labels as guidance, they methodically refine inaccurate annotations and create a high-quality seed dataset for fine-tuning.

Step 4: Model Fine-tuning with Corrected Pilot Data. The pre-trained model is fine-tuned using the expert-corrected seed dataset, adapting its parameters to the unique emotional patterns of teachers. This refinement enhances the model’s ability to generalize to the broader dataset, minimizing domain-specific errors.

Step 5: Large-Scale Annotation with Fine-tuned Model. The entire multimodal dataset undergoes automated labeling using the fine-tuned model, leveraging its improved accuracy to process large volumes efficiently. This step scales the annotation workflow while maintaining consistency across the corpus.

Step 6: Final Expert Review of Full Dataset Annotations. A panel of five experts who are not involved in any prior annotation stages is formed to conduct the final review. These experts independently examine the predicted labels by integrating multimodal information. Adhering strictly to predefined consensus criteria, each annotation was validated through a rigorous voting mechanism: the data is included in the final dataset only when no fewer than four experts unanimously confirmed the correctness of a label; otherwise, the

data is discarded.

Asymmetric Attention-based Multimodal Teacher Sentiment Analysis Model

In this section, we propose a novel Asymmetric Attention-based Multimodal Teacher Sentiment Analysis Model, shorted as AAM-TSA, which consists of four modules to comprehensively utilize the information from four modalities, as shown in Figure 2. We specially design an audio-centric asymmetric attention mechanism to integrate the multimodal features into a unified representation, preserving both modality-specific characteristics and inter-modal semantic associations. The detailed descriptions of the four modules in AAM-TSA are provided in the following.

Multimodal Feature Extraction Module

Upon the T-MED dataset, we focus on processing four modalities of data: textual transcripts, audio tracks, video clips, and instructional information, all of which are crucial for mining the hidden emotional cues of teachers.

We extract the corresponding feature vectors for each modality. For textual transcripts, we utilize RoBERTa(Liu et al. 2020) for processing as shown in E.q. 1, where h_t is the text feature matrix, and T_{input} represents input token embeddings.

$$h_t = \text{RoBERTa}(T_{input}) \quad (1)$$

For audio tracks, HuBERT(Hsu et al. 2021) is employed as shown in E.q. 2, where h_a is the audio feature matrix and A_{input} is the raw waveform.

$$h_a = \text{HuBERT}(A_{input}) \quad (2)$$

Furthermore, we utilize a large language model to create text descriptions for video clips, using tailored prompts

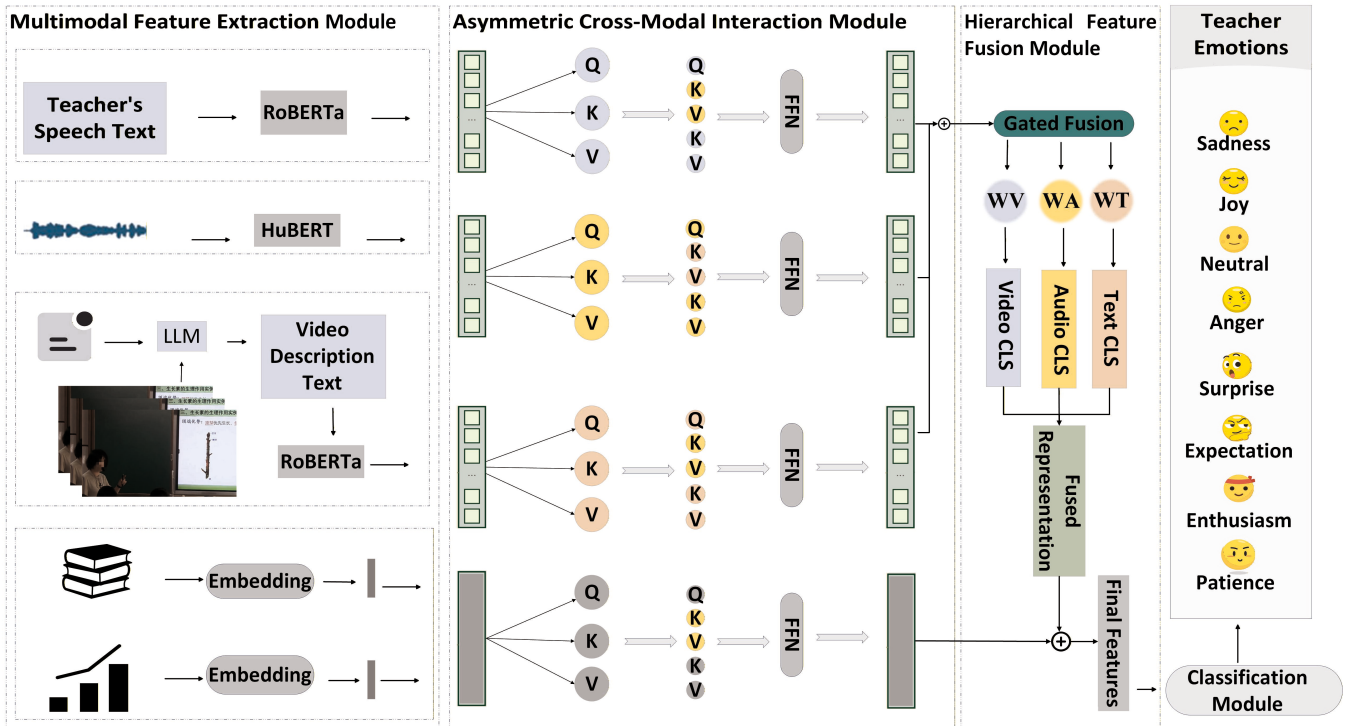


Figure 2: AAM-TSA model architecture diagram.

that emphasize four emotional dimensions: gestures, movements, eye contact, and facial expressions. The feature of video clips are extracted as follows:

$$V_{text} = \text{Video LLM}(V_{input}, P_{emo}), \quad (3)$$

$$h_v = \text{RoBERTa}(V_{text}), \quad (4)$$

where V_{input} is the video clip; P_{emo} is the emotion-focused prompt template, and its specific content is placed in the Appendix B. V_{text} is the generated description text. h_v is the video feature matrix used in our scheme.

To effectively integrate instructional information, we convert discrete educational stage and subject labels into continuous representations to a unified 768-dimensional vector space through a fully connected layer, which is consistent with the dimension of text, audio, and video features. The feature of instructional information are extracted as follows:

$$e_l = \text{Embedding}_{\text{stage}}(l), \quad (5)$$

$$e_s = \text{Embedding}_{\text{subject}}(s), \quad (6)$$

$$h_d = \text{LayerNorm}(\text{GELU}(W_c[e_l; e_s])), \quad (7)$$

where e_l and e_s are educational stage and subject embeddings. W_c projects to joint space. The output h_d is the instructional information feature which we need.

Despite the integration of textual transcripts, audio tracks, and video clips in existing multimodal models, these approaches often fail to effectively analyze teachers' emotions due to the obscure nature of their emotional expressions. By contrast, we convert the video clips to textual descriptions, which can alleviate the interference of a large

amount of emotionally irrelevant content in the original video clips. Furthermore, we introduce instructional information as a supplementary modality to enrich the contextual understanding of the teaching environment, thereby improving the overall efficacy of the emotional analysis.

Asymmetric Cross-Modal Interaction Module

In order to accurately fuse the features from different modalities, we design the asymmetric cross-modal fusion module with an audio-centric strategy.

For text modality, we first place the learnable classification token (CLS) at the first location of its input sequence. Then, we use the text modality as a query and actively retrieve relevant information from the key-value pairs of the audio modality (Vaswani et al. 2017) as follows:

$$h_t^{l+1} = \text{CrossAtt}(h_t^l, h_a^l) = \text{Softmax}\left(\frac{Q_t K_a^T}{\sqrt{d_k}}\right) V_a, \quad (8)$$

where Q_t and K_a, V_a are the query, key and value matrices. h_t^l represents the l -th layer features of text modality. h_a^l denotes audio features, and d_k is the dimension scaling factor. The resulting text modality is then enhanced by self-attention:

$$h_t^{(l)} = \text{SelfAtt}(h_t^{(l-1)}, h_t^{(l-1)}, h_t^{(l-1)}) + h_t^{(l-1)}, \quad (9)$$

where SelfAtt denotes multi-head self-attention within each modality. $h_t^{(l)}$ represents features of text modality at layer l . $h_t^{(l-1)}$ preserves original features. Finally, deeper feature representations are further extracted through feedforward

neural networks as follows:

$$h_t^{(l)} = \text{FFN}(h_t^{(l)}) + h_t^{(l)}, \quad (10)$$

where FFN is a two-layer perceptron with GELU activation. Applied position-wise to each token in $h_t^{(l)}$. For the video feature representation $h_t^{(l)}$ and the instructional information feature representation $h_d^{(l)}$, we employ the same process as text modality, and the detailed description is shown in Appendix C.

For audio modality, we also place CLS at the first location of its input sequence. Then, we use the audio modality as a query and actively retrieve relevant information from the key-value pairs of the text modality as follows:

$$h_a^{l+1} = \text{CrossAtt}(h_a^l, h_t^l) = \text{Softmax}\left(\frac{Q_a K_t^T}{\sqrt{d_k}}\right) V_t, \quad (11)$$

where Q_a and K_t , V_t are the query, key and value matrices respectively. h_a^l represents the l -th layer features of audio modality. h_t^l denotes text features. d_k is the dimension scaling factor. The resulting audio modality is then enhanced by self-attention:

$$h_a^{(l)} = \text{SelfAtt}(h_a^{(l-1)}, h_a^{(l-1)}, h_a^{(l-1)}) + h_a^{(l-1)}, \quad (12)$$

where SelfAtt denotes multi-head self-attention within each modality. $h_a^{(l)}$ represents features of audio modality at layer l . $h_a^{(l-1)}$ preserves original features. Finally, deeper feature representations are further extracted through feedforward neural networks as follows:

$$h_a^{(l)} = \text{FFN}(h_a^{(l)}) + h_a^{(l)}, \quad (13)$$

where FFN is a two-layer perceptron with GELU activation. Applied position-wise to each token in $h_a^{(l)}$. It is worth mentioning that in the cross-attention stage, unlike the other three modes that search related information from the key-value pairs of the audio mode, the audio modes retrieve relevant information from the key-value pairs of the text mode to achieve complementarity of the information. And we stack multiple layers of the above structures to achieve deep iterative fusion.

Hierarchical Feature Fusion Module

To alleviate information redundancy of different modalities, we design the hierarchical feature fusion module. First, the gated fusion mechanism is used to dynamically integrate the three-modal features. The CLS features of text, audio and video are spliced and inputted to the fully connected network, and normalized modal weights are generated through Softmax:

$$g_t, g_a, g_v = \text{Softmax}(\text{MLP}([h_t || h_a || h_v])), \quad (14)$$

where g_t, g_a, g_v are dynamic weights generated by softmax-normalized gates. $||$ denotes concatenation operation. MLP projects the concatenated features to weight space. Then, the feature representation after the fusion process of three modalities is obtained as follows:

$$h_{\text{fused}} = g_t \cdot h_t + g_a \cdot h_a + g_v \cdot h_v, \quad (15)$$

where h_{fused} is a feature vector after the fusion of three modalities' representations. Finally, the gating fusion result are spliced with instructional information feature as follows:

$$h_{\text{final}} = \text{MLP}([h_{\text{fused}} || h_d]), \quad (16)$$

where h_d is the instructional information feature. h_{final} is a feature after the fusion of all modalities.

The hierarchical fusion module preserves the dynamic traits of gated fusion while integrating semantic constraints of the teaching context, enabling the model to capture modal interactions and instructional information features simultaneously. Besides, it achieves a balance between modal complementarity and contextual relevance, which effectively addresses the multiple requirements of emotion recognition in educational scenarios.

Emotion Classification Module

After obtaining the fused feature embeddings, we design an emotion classification module to get the final prediction results. In this module, we employ a current emotion prediction classifier (Devlin et al. 2019), and the prediction process is formulated as follows:

$$p = \text{classifier}(h_{\text{final}}), \quad (17)$$

where p is the probability distribution of eight types of emotions. For classification tasks, we use standard cross entropy loss, and the specific formula is provided in the Appendix C.

Experiment

Experimental Setup

Dataset and Evaluation Metrics. We used the T-MED dataset to perform experiments, which is divided into 80% as the training set, 10% as the verification set and 10% as the test set. We select weighted accuracy (WA) and weighted F1 values (W-F1) as metrics, which can more effectively deal with the problem of uneven distribution of emotional categories in T-MED data. And the detailed settings are provided in the Appendix D.

Baseline Models. As there are few open-source multimodal teacher sentiment analysis methods, we selected nine general multimodal sentiment analysis methods as baseline methods. The nine baseline methods are: Tensor Fusion Network (TFN) (Zadeh et al. 2017), LMF (Liu et al. 2018), Memory Fusion Network (MFN) (Zadeh et al. 2018); Transformer-based MulT (Tsai et al. 2019), MISA (Hazarik, Zimmermann, and Poria 2020), ConFEDE (Yang et al. 2023), Semi-IIN (Lin et al. 2025), MPLMM (Guo, Jin, and Zhao 2024), and MFMB_Net (Tao et al. 2025). The MFMB-Net model adopts a multi-branch structure integrated feature fusion strategy to comprehensively capture the complexity of multimodal data. And detailed descriptions of other baseline models are provided in the Appendix D. All baseline models are implemented using the original authors' open-source code and meticulously reproduced under a unified experimental environment to ensure fair comparison.

Model	T	A	V	I	WA(%)	W-F1(%)
TFN (2017)	✓	✓	✓	✗	75.23	74.15
MFN (2018)	✓	✓	✓	✗	77.84	76.71
LMF(2018)	✓	✓	✓	✗	76.72	75.83
MuLT (2019)	✓	✓	✓	✗	79.20	78.54
MISA (2020)	✓	✓	✓	✗	78.93	78.26
ConFEDE(2023)	✓	✓	✓	✗	80.19	79.36
MPLMM (2024)	✓	✓	✓	✗	79.92	79.21
Semi-IIN (2025)	✓	✓	✓	✗	80.09	79.53
MFMB_Net (2025)	✓	✓	✓	✗	80.91	79.81
AAM-TSA	✓	✓	✓	✓	86.84	86.37

Table 1: Comparison results on T-MED.

Model	Expectation	Neutral	Surprise	Anger	Joy
MFMB_Net	89.59	92.26	56.86	42.55	65.49
AAM-TSA	91.16	95.62	69.09	64.75	81.68

Model	Enthusiasm	Patience	Sadness	M-F1	W-F1
MFMB_Net	47.44	68.93	47.76	63.86	79.81
AAM-TSA	60.00	75.76	55.00	74.13	86.37

Table 2: The Results of Fine-grained sentiment analysis

Performance Comparison on T-MED

Compared with all representative multimodal emotion recognition models, the proposed AAM-TSA achieves significant improvements in WA and W-F1 values as shown in Table 1. The T means that methods using text features, and the A means that methods use audio features. The V means that methods use video features, and the means that methods use features of instructional information. Compared with the current optimal baseline MFMB_Net, the accuracy rate is increased by 5.93%, and W-F1 is increased by 6.56%. We attribute it to our audio-centric design. General models often regard text modalities as core modalities rather than audio. For example, the ConFEDE architecture uses text as an anchor for comparison learning, which makes it suitable for teachers’ sentiment analysis tasks. Although MPLMM also introduced prompt engineering, the MPLMM aimed to deal with the problem of missing modalities rather than generating video description text.

Fine-grained Sentiment Analysis

In order to evaluate the performance more carefully, we compare the proposed AAM-TSA with an SOTA baseline MFMB_Net on eight emotions, as shown in Table 2 For five general emotions, AAM-TSA shows improvements in categories of neutral, surprise, anger, sadness and joy. Anger typically appears as a teaching management style rather than overt or intense anger. Its identification relies more on phonological and rhythmic features, such as tone and intensity, rather than explicit words that convey anger. For three instructional information emotions, AAM-TSA also shows improvements in categories of patience, enthusiasm, and expectation. Overall, the AAM-TSA improves 10.27% and 6.56% in M-F1 and W-F1, respectively. Therefore, we be-

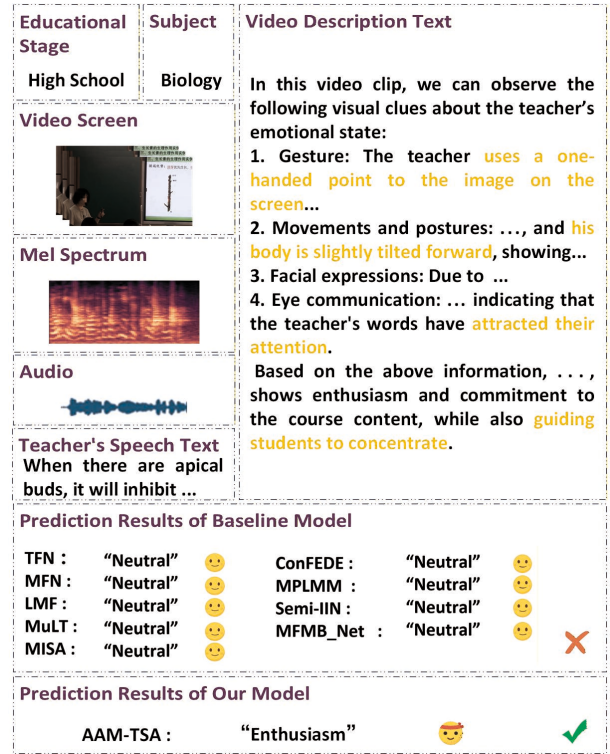


Figure 3: Analysis of predictive case.

lieve that our model significantly improves the recognition ability of various emotions by integrating domain knowledge, semantic visual cues, and enhanced audio analysis.

Case Study: Teacher Sentiment Analysis in a Biology Classroom From High School

We chose a general case to show the decision differences between our AAM-TSA and the baseline models. For all baseline models, their inputs are text, audio, and video, while our model replaces the input video with video description text and additional information for the subject and stage. The instruction information is high school(stage) biology (subject). The description text, teacher’s directive gestures and body forward teacher’s directive gestures and body forward, furthermore reflects the teacher’s positive emotional state when explaining key concepts. Combined with the characteristics of the fast speech speed and mid-frequency energy concentration of the dominant modal audio characteristics, AAM-TSA can be determined that the teacher’s emotion at this time is enthusiastic, while all baseline models are misjudged as neutral, as shown in Figure 3. In this case, text is a basic concept and has no emotional tendency. When the baseline model has no modal weight and additional information, it is easier for the baseline model to predict the case as neutral because neutrality is the majority class sample.

The Comparasions Among Different Modalities

To quantitatively evaluate the effectiveness of different modalities, we conduct systematic experiments on different

Variants	T	A	RV	VD	WA(%)	W-F1(%)
Var-T	✓	✗	✗	✗	69.41	64.08
Var-A	✗	✓	✗	✗	77.29	75.91
Var-VD	✗	✗	✗	✓	49.23	33.29
Var-TA	✓	✓	✗	✗	85.57	84.74
Var-TARV	✓	✓	✓	✗	84.97	84.08
Var-AVD	✗	✓	✗	✓	78.22	77.69
Var-TAVD	✓	✓	✗	✓	85.97	85.10

Table 3: The Comparasions Among Different Modalities

Variants	FE	CI	HF	WA(%)	W-F1(%)
Var-CIHF	✗	✓	✓	85.97	85.10
Var-FEHF	✓	✗	✓	86.11	85.43
Var-FEAC	✓	✓	✗	85.64	85.30
AAM-TSA	✓	✓	✓	86.84	86.37

Table 4: The Ablation Study of Different Modules

combinations of modalities. We design seven variants and the experimental results are shown in Table 3. The RV means that methods use raw video features directly, and the VD means that methods use features of video descriptions.

The Reason for Employing Audio-centric Strategy. We compare the performance among audio single-mode Var-A, text single-mode single-mode Var-T, and video description Var-VD. The experimental results show that the Var-A is superior to the Var-T, with the gap between the two being 7.88% in WA and 11.83% in W-F1. The Var-VD exhibits significantly poorer performance than the Var-T and the Var-A in both metrics. This discovery intuitively confirms that the core carrier of teachers’ emotional expression is audio information rather than text content.

The Damage of Using Raw Video Features. We use Var-TA to denote the model using text and audio, and use Var-TARV to denote the model using video features directly. Compared with the Var-TA model, the Var-TARV model shows a decrease of 0.60% in terms of WA and 0.66% in terms of W-F1. Furthermore, we use the proposed video description module to process the video clips and obtain two variant model denoted as Var-AVD Var-TAVD. The Var-AVD exhibits significantly poorer performance than the Var-TA in both metrics. Compared to Var-TARV The Var-TAVD model shows an increase by 0.40% interms of WA and by 0.36% in terms of W-F1. This shows that unprocessed original video features contain a large amount of emotion-independent noise, which will damage model performance.

The Ablation Study of Different Modules

To quantitatively evaluate the effectiveness of different modules in the AAM-TSA architecture and their contribution to final performance, we conducted systematic ablation experiments. We design three variants and the experimental results are shown in Table 4. The FE means that variant models keep the multimodal feature extraction module, and CI means that

variant models keep the asymmetric cross-modal interaction module. The HF means that variant models keep the hierarchical feature fusion module, and all the variant models keeps the emotion classification module

First, we construct a variant model denoted as Var-CIHF by removing the instructional information feature embedding module. Compared with the AAM-TSA, the WA and W-F1 decrease 0.87% and 1.27%, respectively. Thus, the introduction of school stage and subject as instructional information is crucial to understanding the situational dependence of teachers’ emotional expression.

Second, we construct a variant model denoted as Var-FEHF by removing the asymmetric cross-modal attention interaction module. Compared with the AAM-TSA, the WA and W-F1 decrease 0.73% and 0.94%. In the teacher’s sentiment analysis, the audio information carries the most significant emotional signals, allowing visual features to align the audio one-way (rather than two-way interaction), and allowing text and audio to complement each other in both directions. In addition, this module also allows instructional information features to actively query audio information, which can guide information flow and fusion effectively.

Fianlly, we construct a variant model denoted as Var-FEAC by removing the hierarchical feature fusion module. Compared with the AAM-TSA, the WA and W-F1 decrease 1.2% and 1.07%. This indicates that hierarchical feature fusion module can effectively utilize instructional information features and is critical to dealing with the importance differences and redundancy of multimodal information. And more experiments and analyses that demonstrate the superiority of our design are provided in the Appendix E.

Conclusion

In this work, we address the critical gap in teacher sentiment analysis by systematically advancing both dataset construction and model development. Firstly, we collect T-MED, the first large-scale multimodal dataset tailored for teacher sentiment analysis, through an efficient human-machine collaborative annotation framework. Secondly, we propose the AAM-TSA model, a novel asymmetric attention-based framework designed to address the performative nature of teacher emotional expressions. By leveraging an audio-centric asymmetric attention mechanism, AAM-TSA enables differentiated cross-modal interaction, allowing effective fusion of heterogeneous features while preserving modality-specific traits. Extensive experimental results on T-MED demonstrate that AAM-TSA outperforms nine multimodal sentiment analysis baselines by a significant margin in terms of WA and W-F1. This work not only provides the rigorous T-MED benchmark dataset to facilitate future research but also offers a specialized AAM-TSA model that advances the understanding of multimodal emotional expression in educational scenarios. By integrating pedagogical principles with technological advancements, our research paves the way for intelligent teaching tools, teacher development programs, and educational intervention systems, fostering improved teaching effectiveness and enriched student learning experiences.

Acknowledgments

This work was funded by the National Natural Science Foundation of China (Nos. 62567005 and 62406127), and Natural Science Foundation of Inner Mongolia Autonomous Region of China (No. 2025MS06004).

References

- Alibaba DAMO Academy. 2024. ClearerVoice-Studio.
- Bagher Zadeh, A.; Liang, P. P.; Poria, S.; Cambria, E.; and Morency, L.-P. 2018. Multimodal Language Analysis in the Wild: CMU-MOSEI Dataset and Interpretable Dynamic Fusion Graph. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2236–2246.
- Bellard, F.; and Developers, F. 2000. FFmpeg.
- Cai, T.; Wang, S.; Xiong, Y.; and Zhong, X. 2024. Exploiting Adaptive Adversarial Transfer Network for Cross Domain Teacher’s Speech Emotion Recognition. In *Computer Science and Education. Computer Science and Technology*, 202–213.
- Cheng, Z.; Cheng, Z.-Q.; He, J.-Y.; Sun, J.; Wang, K.; Lin, Y.; Lian, Z.; Peng, X.; and Hauptmann, A. G. 2024. Emotion-LLaMA: Multimodal Emotion Recognition and Reasoning with Instruction Tuning. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 37, 110805–110853.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *NAACL*.
- Frenzel, A. C.; Becker-Kurz, B.; Pekrun, R.; and Goetz, T. 2018. Emotion transmission in the classroom revisited: A reciprocal effects model of teacher and student enjoyment. *Journal of Educational Psychology*, 110(5): 628–639.
- Guo, Z.; Jin, T.; and Zhao, Z. 2024. Multimodal prompt learning with missing modalities for sentiment analysis and emotion recognition. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1726–1736.
- Han, J.; Jin, L.; and Yin, H. 2023. Mapping the research on language teacher emotion: A systematic literature review. *System*, 118: 103138.
- Hargreaves, A. 2000. Mixed emotions: teachers’ perceptions of their interactions with students. *Teaching and Teacher Education*, 16(8): 811–826.
- Hazarika, D.; Zimmermann, R.; and Poria, S. 2020. MISA: Modality-invariant and -specific representations for multimodal sentiment analysis. In *Proceedings of the 28th ACM International Conference on Multimedia*, 1122–1131.
- He, Y.; Lu, X.; Sun, D.; Pan, T.; Qiu, Y.; and Liu, J. 2024. Research on Teacher Classroom Teaching Speech Emotion Recognition Based on LSTM. In *2024 International Conference on Asian Language Processing (IALP)*, 326–331.
- Hsu, W.-N.; Bolte, B.; Tsai, Y.-H. H.; Lakhota, K.; Salakhutdinov, R.; and Mohamed, A. 2021. HuBERT: Self-Supervised Speech Representation Learning by Masked Prediction of Hidden Units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29: 3451–3460.
- King, J.; Almukhailid, H.; Mercer, S.; Babic, S.; Mairitsch, A.; and Sulis, G. 2024. Teacher emotions and the emotional labour of modern language (ML) teachers working in UK secondary schools. *International Review of Applied Linguistics in Language Teaching*, 62(3): 1213–1235.
- Klein, G.; and Systran. 2023. Faster Whisper Transcription.
- Li, L. 2022. Emotion Analysis Method of Teaching Evaluation Texts Based on Deep Learning in Big Data Environment. *Hindawi Computational Intelligence and Neuroscience*, 2022: 9909209.
- Lin, J.; Wang, Y.; Xu, Y.; and Liu, Q. 2025. Semi-IIN: Semi-Supervised Intra-Inter Modal Interaction Learning Network for Multimodal Sentiment Analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 1411–1419.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2020. Ro{BERT}: A Robustly Optimized {BERT} Pretraining Approach.
- Liu, Z.; Shen, Y.; Lakshminarasimhan, V. B.; Liang, P. P.; Bagher Zadeh, A.; and Morency, L.-P. 2018. Efficient Low-rank Multimodal Fusion With Modality-Specific Factors. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2247–2256.
- Lu, L.; Yuan, H.; Yang, S.; Feng, L.; and Ding, X. 2024. Review on Teacher’s Classroom Language Behavior Analysis Based on Clustering and Emotional Analysis. *Communications, Signal Processing, and Systems*, 487–496.
- Ma, Z.; Zheng, Z.; Ye, J.; Li, J.; Gao, Z.; Zhang, S.; and Chen, X. 2024. emotion2vec: Self-Supervised Pre-Training for Speech Emotion Representation. In *Findings of the Association for Computational Linguistics: ACL 2024*, 15747–15760.
- McDuff, D.; and Soleymani, M. 2017. Large-scale Affective Content Analysis: Combining Media Content Features and Facial Reactions. In *2017 12th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2017)*, 339–345.
- Rubie-Davies, C. M.; Weinstein, R. S.; Huang, F. L.; and et al. 2023. Beyond Pygmalion effects: Teacher expectancy as a driver of longitudinal student achievement growth. *Learning and Instruction*, 88: 101800.
- Shaikh, S.; and Doudpotta, S. M. 2019. Aspects Based Opinion Mining for Teacher and Course Evaluation. *Sukkur IBA Journal of Computing and Mathematical Science*, 3(1): 34–43.
- Tao, C.; Li, J.; Zang, T.; and Gao, P. 2025. A Multi-Focus-Driven Multi-Branch Network for Robust Multimodal Sentiment Analysis. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(2): 1547–1555.
- Tsai, Y.-H. H.; Bai, S.; Liang, P. P.; Kolter, J. Z.; Morency, L.-P.; and Salakhutdinov, R. 2019. Multimodal Transformer for Unaligned Multimodal Language Sequences. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 6558–6569.

- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, volume 30.
- Wang, H.; and Frenzel, A. C. 2025. “Exhaustive but effective”: A multi-site study investigating the profiles of teachers’ emotions and emotional labor. *Journal of School Psychology*, 110: 101456.
- Yang, J.; Yu, Y.; Niu, D.; Guo, W.; and Xu, Y. 2023. ConFEDE: Contrastive Feature Decomposition for Multimodal Sentiment Analysis. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 7617–7630.
- Zadeh, A.; Chen, M.; Poria, S.; Cambria, E.; and Morency, L.-P. 2017. Tensor Fusion Network for Multimodal Sentiment Analysis. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 1103–1114.
- Zadeh, A.; Chen, M.; Poria, S.; Cambria, E.; and Morency, L.-P. 2018. Memory Fusion Network for Multi-view Sequential Learning. In *Proceedings of ACL*, 22–32.
- Zhang, C.; Wang, Z.; Fang, Z.; and Xiao, X. 2024. Guiding student learning in video lectures: Effects of instructors’ emotional expressions and visual cues. *Computers & Education*, 216: 105062.
- Zhang, S.; and Li, C. 2022. Research on Feature Fusion Speech Emotion Recognition Technology for Smart Teaching. *Mobile Information Systems*, 2022: 7785929.
- Zhao, G.; Zhang, Y.; and Chu, J. 2024. A multimodal teacher speech emotion recognition method in the smart classroom. *Internet of Things*, 25: 101069.
- Zheng, Q.; Chen, Z.; Wang, M.; Shi, Y.; Chen, S.; and Liu, Z. 2024. Automated Multimodal Teaching Behavior Analysis: A Pipeline-Based Event Segmentation and Description. *IEEE Transactions on Learning Technologies*, 17: 1677–1693.