

Enhancing Multimodal Affective Analysis with Learned Live Comment Features

Zhaoyuan Deng, Amith Ananthram, Kathleen McKeown

Columbia University

zd2286@columbia.edu, amith@cs.columbia.edu, kathy@cs.columbia.edu

Abstract

Live comments, also known as *Danmaku*, are user-generated messages that are synchronized with video content. These comments overlay directly onto streaming videos, capturing viewer emotions and reactions in real-time. While prior work has leveraged live comments in affective analysis, their use has been limited due to the relative rarity of live comments across different video platforms. To address this, we first construct the Live Comment for Affective Analysis (LCAffect) dataset which contains live comments for English and Chinese videos spanning diverse genres that elicit a wide spectrum of emotions. Then, using this dataset, we use contrastive learning to train a video encoder to produce synthetic live comment features for enhanced multimodal affective content analysis. Through comprehensive experimentation on a wide range of affective analysis tasks (sentiment, emotion recognition, and sarcasm detection) in both English and Chinese, we demonstrate that these synthetic live comment features significantly improve performance over state-of-the-art methods.

Code, Dataset, and Appendix —

<https://github.com/dzy49/AffectiveLiveComm>

Introduction

Live comments, or *Danmaku*, are a subtitle system originating from Japan and popularized in China. This system overlays real-time user-generated messages directly onto online video streams. These comments, synchronized with a video’s timeline, offer a rich layer of viewer interaction, providing insights into user emotions, user reactions, and the video’s broader context (see Figure 1). The richness of these comments makes them an ideal resource for enhancing multimodal affective analysis, which seeks to interpret and categorize the emotional content and dynamics of videos. Although prior research, such as the study by Niu et al. (2018), has demonstrated how live comments can improve performance in affective analysis tasks like emotion prediction, their broader application remains constrained. Many video platforms do not support live comments, rendering these insights impossible for videos on those platforms. Consequently, the rarity of live comments poses significant challenges in fully exploiting their analytical potential.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 1: An example video frame and dialogue from the TV show *The Big Bang Theory* with accompanying live comments overlaid. Below each original Chinese comment, we include its English translation. By training a video encoder to produce representations similar to these live comments, we can produce multimodal features that adapt well to affective analysis tasks like emotion recognition.

In this work, we bridge this gap by proposing a novel method for producing synthetic live comment features to improve the performance of multimodal affective analysis. We demonstrate the effectiveness of these features in English and Chinese across several video understanding tasks: sentiment analysis, emotion recognition, and sarcasm detection.

One major challenge in affective learning from live comments stems from the nature of existing large-scale live comment datasets which are primarily tailored to the live comment generation task. These datasets typically feature a large percentage of popular content genres such as video games, sports, and music videos. However, these categories typically lack the rich emotional context required for detailed affective analysis. This mismatch complicates the training of models capable of accurately understanding and interpreting nuanced emotions. Additionally, the majority of these

datasets are collected from Bilibili, which features amateur-created videos of everyday life. These videos often provoke a more limited and homogeneous range of emotional responses (Rosenbusch, Evans, and Zeelenberg 2019), especially when compared to the diverse and intense emotions elicited by content like dramas or documentaries. Moreover, most datasets comprise exclusively Chinese content, which further restricts their multilingual applicability.

A second significant challenge is the dependency of existing methods on the availability of live comments. In the real world, the presence of live comments can be limited and not uniformly distributed across all video genres. This reliance restricts the practical application of these methods, as they cannot work in scenarios where comments are non-existent.

To address these challenges, we first construct a large-scale, diverse dataset comprising everyday life videos in Chinese as well as bilingual TV shows, movies, and documentaries, enabling our system to learn a broad spectrum of emotional experiences across various contexts. Then, using our dataset, we develop a multimodal encoder that effectively learns from our collected videos and comments by contrasting different contexts. This encoder is capable of producing synthetic live comment features from videos, thereby enabling the inference of emotional context across modalities, even in the absence of live comments.

Our experimental analysis demonstrates that adding our synthetic live comment features improves over state-of-the-art (SOTA) approaches that leverage text, acoustic, and visual modalities directly. Our system achieves new SOTA performance on multiple multimodal datasets, particularly enhancing accuracy in detecting sentiment, emotions, and sarcasm. This not only underscores the utility of live comments in affective video analysis but also opens new avenues for multimodal learning applications.

Our main contributions are:

- 1) Live Comments for Affective Analysis dataset (LCAffect), an expansive corpus of over 11 million live comments. It includes bilingual video content (Chinese and English) spanning a diverse range of topics and genres, curated to capture a wide spectrum of emotions.
- 2) A contrastive encoder that learns to project videos into live comment representation space, allowing the inference of synthetic live comment features for any video.
- 3) A downstream multimodal fusion model for affective analysis tasks that utilizes not only three modalities but also our synthetic live comment feature.
- 4) New SOTA across three affective analysis tasks, including a 3.18-point increase in accuracy on CH-SIMS v2 (sentiment analysis), a 2.89-point increase in F1 on MELD (emotion recognition), and a 3.0-point increase in F1 on MuSTARD (sarcasm detection).

Related Work

Learning from Comments Previous research has highlighted the value of user comments in improving performance across various multimodal tasks. Fu et al. (2017) show that live comments on Twitch help predict highlights in e-sport games. Similarly, Yang, Ai, and Hirschberg (2019)

show that live comments help identify humor in online videos. Additionally, Niu et al. (2018) show that features derived from live comments improve affective content analysis. Moreover, Hanu et al. (2022) show that user comments can help learn more contextualized representations for image, video, and audio, thus improving video-text retrieval.

These studies depend on the availability of user comments which are not always present in the real world. In contrast, our work seeks to develop a system that is pre-trained to align video segments with corresponding live comments. This approach aims to learn a multimodal representation space that effectively augments various downstream understanding tasks. Consequently, our system can be directly applied to videos lacking associated live comments, producing synthetic live comment features for such videos. This capability directly addresses the limitations of prior work, enhancing the analysis and interpretation of videos that would be excluded due to the absence of live comments.

Live Comment Datasets Several datasets featuring live comments have been established in the literature. LiveBot (Ma et al. 2018) and VideoIC (Wang, Chen, and Jin 2020) are large-scale datasets drawn from Bilibili, a site which primarily features short, user-generated videos focused on everyday life. In contrast, Lalanne, Bournet, and Yu (2023) collected live-streamed video and comments from Twitch, a platform centered around video games. Finally, MovieLC (Chen et al. 2023) is a compilation of famous movies with accompanying comments from Tencent Video.

Our work builds on these efforts by creating a new dataset that includes videos from a broader range of platforms. By integrating multiple video hosting platforms, our dataset includes both a variety of short, user-generated content and longer formats like TV shows, movies, and documentaries in Chinese and English. This expansive collection allows the study of many genres and enables the training of a multimodal encoder that can learn live comment features effectively for affective analysis from diverse videos.

Multimodal Affective Analysis Previously, multimodal affective computing often relied on hand-crafted algorithms to perform initial feature extraction. Recently, however, there has been a shift towards using pre-trained modality encoders with end-to-end tuning to improve performance. Specifically, Yi et al. (2024) use CLIP (Radford et al. 2021)’s spatial encoder and TimeSformer (Bertasius, Wang, and Torresani 2021) to encode visual features, while Wu et al. (2024) employ pre-trained audio encoders such as Data2Vec (Baeovski et al. 2022) and Hubert (Hsu et al. 2021) to process acoustic features. Although these studies demonstrate the efficacy of an end-to-end framework, they are limited to two modalities. Our work extends this approach by integrating all three modalities (text, acoustic, visual), achieving more robust video analysis. Another promising research direction in affective analysis involves the integration of external knowledge. Ghosal et al. (2020) employ features produced by a common sense knowledge model to enhance emotion recognition, while Hu et al. (2022) exploit the complementary knowledge underlying sentiment analysis and emotion recognition to build a knowledge-sharing framework. Our work parallels these efforts by leveraging external knowl-

edge derived from live comments.

The Live Comment Dataset

Collection of Videos

We collect videos and their accompanying live comments from several popular Chinese video streaming websites: Bilibili, Tencent Video, iQIYI, and Youku. Our dataset includes both Chinese and English content and is partitioned into three distinct subsets based on content type:

User-Generated Content This subset includes 2,464 videos from Bilibili featuring naturally occurring conversations from everyday life. The initial set of 682 videos is available from the Linguistic Data Consortium (UPenn) and was manually inspected to ensure content quality and relevance. Each video includes interactions involving at least two people with observable emotions. The remaining 1,782 videos are sourced from Bilibili’s recommendation API, which selects content similar to the manually inspected set.

TV shows and documentaries We collect 443 episodes from 8 TV shows and documentaries. This subset includes content from multiple platforms, including dramas, sitcoms, and documentaries focused on social life. Their inclusion aims to provide the model with insights into scripted content that elicits a wide range of emotions in different contexts.

Movies Sourced from Tencent Video, this subset features 59 popular comedy, drama, and action films. They feature the longest videos, enabling the analysis of viewer interactions over extended contexts through live comments.

In addition, we collect text transcripts for the videos. For user-generated content, transcripts are obtained using Bilibili’s API, which uses automatic speech recognition. For movies and TV shows, we extract transcripts from hard-coded subtitles via optical character recognition (OCR). We extract both English and Chinese subtitles when available.

Dataset Statistics

Our LCAffect dataset contains 11,356,212 live comments aligned with 829 hours of video. Table 1 provides a detailed breakdown of the dataset, including subset-specific statistics. As shown in Table 2, our dataset establishes a new standard in live comment analysis by offering nearly twice the video duration found in the largest existing datasets, encompassing a broader variety of content types, and featuring over twice as many live comments as any comparable dataset.

Contrastive Pre-training

Although multiple studies have demonstrated the effectiveness of live comments in enhancing video understanding, the availability of live comments is largely limited outside of China and Japan. This restriction significantly limits the practical application of these studies. To overcome this challenge, we propose a video encoder pre-trained on our diverse live comment dataset. This model leverages a standard contrastive learning framework akin to CLIP (Radford et al. 2021), where positive examples consist of matching live comments and videos, and negative examples involve non-matching pairs. The aim is to produce a synthetic live

Category	User G Content	TV Shows	Movies
# Videos	2,464	443	59
# w. Eng. Subtitle	-	281	33
# Comm. (k)	2,367	4,272	4,717
Dur. (h)	493.4	212.7	122.9
Avg. Dur. (s)	720.9	1,728.5	7,500.0
Avg. # Comm	960.6	9,643.4	79,956.6
Avg. # Char	8.8	12.1	9.9

Table 1: Subset-specific statistics of our LCAffect dataset.

Dataset	LiveBot	VideoIC	MovieLC	LCAffect
# Videos	2,361	4,951	85	2,966
# Comm. (k)	895	5,330	1,406	11,356
Dur. (h)	114	557	175	829
Content Lang.	zh	zh	zh	zh & en
Composition	User G. Content	User G. Content	Mov.	User G.C. TV, Mov.

Table 2: Comparison with other live comment datasets.

comment feature from a video span that can be used to enhance performance in various affective analysis tasks.

Pre-training Model Architecture

We adopt a CLIP-style contrastive pre-training approach. During the training phase, we partition each video into segments s of σ seconds. For each segment s_i , we gather the associated text transcript t_i , video frames f_i , and live comments c_i , and treat the segment as an individual training sample. For this stage, we use only the text and visual modalities. We create our Video-to-Live Comment (V2LC) encoder by employing a pre-trained text encoder to encode the video text transcript into embedding \mathbf{S}_T and a Video Vision Transformer (Arnab et al. 2021) to encode the video frames into embedding \mathbf{S}_F . These encoded outputs are then integrated using a cross-modality encoder to produce the final output video segment embedding \mathbf{S} . The cross-modality encoder employs the Cross-Attention to Concatenation strategy, where we first perform two streams of cross-attention, and then the output is concatenated and processed by another Transformer to model the global context.

Let \mathbf{S}_T and \mathbf{S}_F denote the transcript text and video frame embeddings respectively. Let \mathbf{S} denote the V2LC output segment embedding produced by the multimodal interactions:

$$\begin{cases} \mathbf{S}_T \leftarrow \text{Attention}(Q_t, K_f, V_f), \\ \mathbf{S}_F \leftarrow \text{Attention}(Q_f, K_t, V_t), \\ \mathbf{S} \leftarrow \text{Mean}(\text{Transformer}(\text{Concat}(\mathbf{S}_T, \mathbf{S}_F))) \end{cases} \quad (1)$$

Then, we create the live comment feature embedding \mathbf{C} by encoding a segment s_i ’s corresponding live comments c_i with a second text encoder. We perform contrastive learning by maximizing the cosine similarity of the video segment embedding \mathbf{S} and live comment embedding \mathbf{C} .

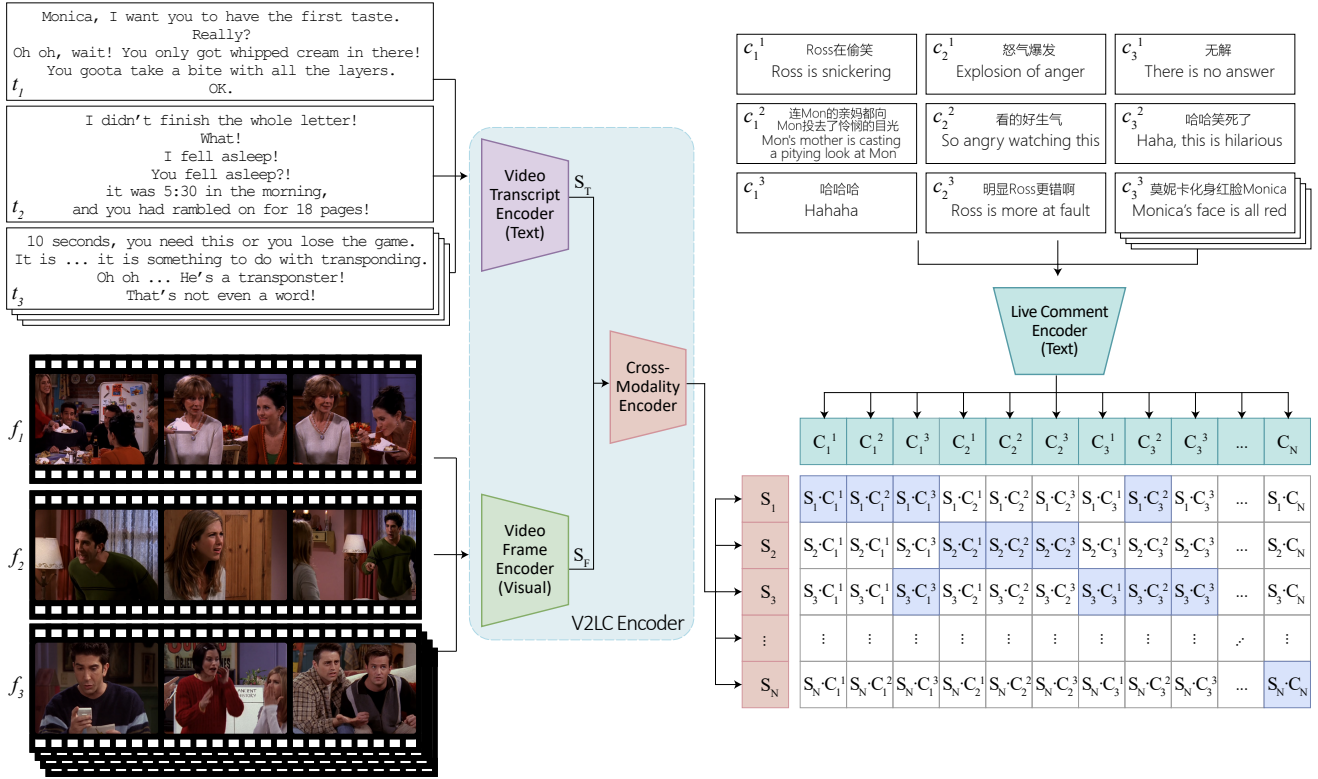


Figure 2: Our contrastive pre-training approach. We train our V2LC encoder to predict the correct pairings of a batch of (video, live comment) training examples. Correct pairings of video segment embedding S and live comment embedding C are highlighted in blue. Comments c_1^3 and c_3^2 have high similarity; thus they are correct matches for both Segment 1 and 3.

The goal of the V2LC encoder is to project each input video segment s_i into the live comment representation space, so we freeze the live comment encoder and optimize only the video segment-to-live comment loss, training the V2LC encoder to match the representation of s_i 's accompanying live comments. The resulting training objective is:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N \left[\log \frac{e^{(S_i \cdot C_i) * \tau}}{\sum_{j=1}^N e^{(S_i \cdot C_j) * \tau}} \right] \quad (2)$$

where N is the number of video segment-live comment pairs in each batch, $(S_i \cdot C_i)$ is the similarity (dot product) between video segment embedding S_i and live comment embedding C_i , and τ is a learned temperature parameter.

We use the V2LC encoder to produce synthetic live comment features to augment models for downstream tasks.

Multi-label Objective

The original CLIP architecture was designed primarily for single-label classification tasks. In our setting, there exists a many-to-many relationship, where a single video segment usually contains multiple live comments, and identical comments may appear in multiple video segments. To reconcile this discrepancy, we have modified the CLIP training framework to accommodate multiple labels: Given a batch of N videos and K live comments, the model is trained to predict

which of the $N \times K$ possible (video, live comment) pairings across a batch actually co-occur.

An additional challenge arises when multiple similar comments from different video segments appear within the same batch. In our dataset, similar comments naturally occur more frequently than the more distinct content pairs seen in the image-text datasets used in CLIP. During pre-training, if similar comments from different video segments are sampled within the same batch, it might lead to confusion in the contrastive learning setup. To mitigate this issue, we refine the training targets such that for a given video segment s , any live comment c in the batch with a vector similarity exceeding a predefined threshold θ to any actual live comment of s is also included in the target label set. Such comments are deemed correct, thereby allowing the model to learn a good live comment representation more efficiently.

Downstream Fine-tuning

Multimodal Fusion Encoder

Currently, in multimodal affective analysis, the text modality is commonly encoded with pre-trained language models like BERT (Devlin et al. 2019). In contrast, the acoustic and visual modalities have traditionally relied on hand-crafted feature extractors such as OpenSmile (Eyben, Wöllmer, and Schuller 2010) for audio and OpenFace (Baltrušaitis, Robin-

son, and Morency 2016) for facial expressions (Liu et al. 2022). These tools manually define the features to be extracted, which might not capture the full complexity of the data. However, recent work has introduced large-scale pre-trained encoders for these modalities which have yielded significant improvements in multimodal affective analysis. Therefore, we propose a downstream Multimodal Fusion Encoder that utilizes large-scale pre-trained encoders across all three modalities—text, acoustic, and visual—to allow more comprehensive modeling of affective tasks.

We first extract features f_m from each modality with pre-trained encoders, then we use a cross-modality encoder to fuse them. To accommodate three modalities, we construct the following cross-modality encoder:

$$\begin{cases} \mathbf{Z}_t \leftarrow \text{SelfAttn.}(\text{Attention}(Q_t, K_{a+v}, V_{a+v})), \\ \mathbf{Z}_a \leftarrow \text{SelfAttn.}(\text{Attention}(Q_a, K_{v+t}, V_{v+t})), \\ \mathbf{Z}_v \leftarrow \text{SelfAttn.}(\text{Attention}(Q_v, K_{t+a}, V_{t+a})), \\ \mathbf{Z} \leftarrow \text{Concat}(\text{Mean}(\mathbf{Z}_t), \text{Mean}(\mathbf{Z}_a), \text{Mean}(\mathbf{Z}_v)), \end{cases} \quad (3)$$

where

$$\begin{aligned} \mathbf{Q}_{m1} &= W_q \cdot f_{m1} \\ \mathbf{K}_{m2+m3} &= W_k \cdot \text{Concat}(f_{m2}, f_{m3}) \\ \mathbf{V}_{m2+m3} &= W_v \cdot \text{Concat}(f_{m2}, f_{m3}) \end{aligned} \quad (4)$$

Finally, we pass the concatenated features \mathbf{Z} through feed-forward layers to produce our final prediction.

Augmenting with Live Comment Features

To enhance performance on downstream tasks, first, we apply self-attention to the synthetic live comment features produced by the V2LC encoder before mean-pooling. Then, we take the mean of these attentionally-pooled synthetic live comment features and concatenate it with \mathbf{Z} , our multimodal representation from our fusion encoder. These combined features are then processed through feed-forward layers. We hypothesize that this joint training enables rich integration of the emotional context captured by the synthetic live comment features from our V2LC encoder, improving their efficacy in downstream affective analysis.

Experiments

Evaluation Datasets

We evaluated our work on six widely used affective analysis datasets: Chinese Multimodal Sentiment Analysis Dataset (CH-SIMS) (Yu et al. 2020), Chinese Multimodal Sentiment Analysis Dataset v2.0 (CH-SIMS v2) (Liu et al. 2022), Multimodal Opinion Sentiment and Emotion Intensity (MOSI) (Zadeh et al. 2016), Multimodal Sentiment Analysis (MOSEI) (Bagher Zadeh et al. 2018) for sentiment analysis, Multimodal EmotionLines Dataset (MELD) (Poria et al. 2019) for emotion recognition, and Multimodal Sarcasm Detection Dataset (MUSARD) (Castro et al. 2019) for sarcasm detection. Summaries of the datasets can be found in Table 3.

CH-SIMS: This dataset contains short video segments from Chinese TV shows and movies, annotated with modality-specific and multimodal sentiment. We predict the multimodal sentiment labels.

Dataset	Lang	#Total segments	Avg duration (s)	Avg word count
CH-SIMS	zh	2,282	3.67	15.8
CH-SIMS v2	zh	4,402	3.63	17
MOSI	en	2,199	4.2	12
MOSEI	en	22,856	7.28	19.7
MELD	en	13,708	3.59	8.0
MUSARD	en	690	19.17	42.4

Table 3: Statistics of our evaluation datasets.

CH-SIMS v2: An extension of the original CH-SIMS dataset, CH-SIMS v2 adds an additional 2,121 video segments and improves the balance and diversity of the dataset.

MOSI: This dataset contains English movie review monologues from YouTube labeled with multimodal sentiment.

MOSEI: An extension of MOSI, MOSEI expands its coverage to include more videos and topics.

MELD: Drawn from the TV show *Friends*, this dataset contains utterance-level labels chosen from 7 emotions: anger, sadness, joy, neutral, fear, surprise, or disgust.

MUSARD: This dataset contains video clips from *Friends*, *The Golden Girls*, *Sarcasmaholics Anonymous*, and *The Big Bang Theory*. These audiovisual utterances, with their surrounding context, are annotated with sarcasm labels.

Implementation

Data Processing For pre-training, we use a segment length σ of 8 seconds, balancing between required context and the length of downstream datasets. We sample 8 frames uniformly from each segment. To reduce noise in the dataset, we employ multiple filtering strategies. First, we exclude comments lacking substantial content, specifically those shorter than 2 characters or those without Chinese characters. Second, we compile a list of low-signal terms and exclude comments containing these words. For user-generated videos, we trim the first and last 15 seconds as they tend to include repetitive comments such as greetings and farewells. For movies, we trim the first and last 5 minutes, and for TV shows, we trim the start and end of each show.

Segments containing fewer than 5 live comments are excluded from pre-training to allow efficient GPU batching. For each epoch, we randomly select 5 live comments per segment so that a batch with N samples has $K = 5N$ comments. We sample 10% of our data for validation.

Implementation Details We pre-train two variants of our V2LC encoder: one that uses only Chinese text transcripts and another that uses English text transcripts when available. For the Chinese-only variant, we use `Chinese-ROBERTa-wwm-ext` (Cui et al. 2020) as our transcript encoder; for the bilingual variant, we use `XLM-ROBERTa-base` (Conneau et al. 2019). For encoding live comments, We use `Chinese-ROBERTa-wwm-ext`.

For downstream fine-tuning, we use pre-trained encoders tailored to each modality and language. For the text

Models	Modality	CH-SIMS v2						CH-SIMS			
		Acc ₂	Acc ₂ ^w	F1	Corr	R ²	MAE	Acc ₂	F1	Corr	MAE
MuT	T+A+V	79.50	69.61	79.59	70.32	47.15	31.7	78.56	79.66	56.41	45.3
AV-MC	T+A+V	83.46	74.54	83.52	76.04	57.37	28.6	-	-	-	-
VLP2MSA	T+V	-	-	-	-	-	-	79.43	79.26	60.38	41.29
MMML	T+A	-	-	-	-	-	-	82.93	82.9	73.26	33.2
Ours	T+A+V	86.94	76.70	86.88	82.32	66.71	26.6	83.16	83.16	74.32	32.8
+ LC	T+A+V	90.12*	80.24*	90.10*	86.14*	70.82*	25.9*	84.88*	84.89*	76.12*	31.9*

Table 4: Results on CH-SIMS and CH-SIMS v2 for Chinese sentiment analysis. The performances of baseline models are shared by their authors. Our experimental results are averages across three random seeds. T, A, and V indicate text, acoustic, and visual modalities, respectively. ‘‘Ours’’ denotes our vanilla Multimodal Fusion Encoder and ‘‘+ LC’’ denotes our model augmented with synthetic live comment features. Asterisks indicate t-test $p < 0.05$ when compared with our vanilla Multimodal Fusion Encoder.

Models	Modality	MOSI					MOSEI				
		Acc ₂	Acc ₇	F1	Corr	MAE	Acc ₂	Acc ₇	F1	Corr	MAE
MuT	T+A+V	84.10	-	83.90	71.1	86.1	82.5	-	82.3	71.3	58.0
UniMSE	T+A+V	86.9	48.68	86.42	80.9	69.1	87.50	54.39	87.46	77.3	52.3
VLP2MSA	T+V	86.28	-	86.26	81.3	69.6	85.97	-	85.89	77.0	53.5
MMML	T+A	88.16	48.25	88.15	83.8	64.29	86.73	54.95	86.49	79.08	51.74
Ours	T+A+V	88.26	49.62	88.17	84.10	63.4	86.24	54.95	86.21	79.10	51.6
+ LC	T+A+V	89.02*	50.14*	89.06*	84.50	63.2	87.15*	55.15	87.12*	80.22*	51.3

Table 5: Results on MOSI and MOSEI for English sentiment analysis.

Models	Modality	MELD	MUSTARD		
		F1 _w	P _w	R _w	F1 _w
SVM	T+A	-	65.2	62.9	63.0
UniMSE	T+A+V	65.51	-	-	-
Ours	T+A+V	66.73	68.2	67.4	67.5
+ LC	T+A+V	69.62*	70.6*	70.5*	70.5*

Table 6: Results on MELD for emotion recognition and MUSTARD for sarcasm detection.

modality, we use Chinese-RoBERTa-wmm-ext for Chinese datasets and RoBERTa-base for English datasets (Liu et al. 2019). For the acoustic modality, we adopt Chinese-HuBERT-base to encode Chinese data and Data2Vec-audio-base for English data. For visual features, we process all data using TimeSformer-base.

Other implementation details are specified in the Appendix.

Baseline Models

We compare our model against multiple competitive baselines:

Multimodal Transformer (MuT) (Tsai et al. 2019): Chosen for its pioneering approach of applying attention across modalities, MuT enables dynamic adaptation between modalities at different time steps, addressing the challenges posed by unaligned multimodal data.

Support Vector Machine (SVM): Known for its robust

performance on small-sized datasets, SVM can, at times, surpass neural models. Following Castro et al. (2019), features are concatenated via early fusion and fed into an SVM classifier.

Acoustic Visual Mixup Consistent framework (AV-MC) (Liu et al. 2022): The current SOTA for CH-SIMS v2, AV-MC introduces a modality mixup module as data augmentation, which mixes the acoustic and visual modalities from different videos to enhance performance.

Multimodal Sentiment Knowledge-sharing Framework (UniMSE) (Hu et al. 2022): Chosen for its strong performance on MOSEI and MELD, this model proposes a knowledge-sharing framework that unifies two affective analysis tasks: multimodal sentiment analysis and emotion recognition in conversation, demonstrating the effectiveness of utilizing complementary knowledge in affective tasks.

Vision-Language Pre-Training To Multimodal Sentiment Analysis (VLP2MSA) (Yi et al. 2024): Chosen for its use of pre-trained vision models, VLP2MSA extracts spatio-temporal features from sparsely sampled video frames, offering advantages over traditional visual feature extraction. It captures not only facial expressions but also body movements, providing a more comprehensive analysis of visual information.

Multimodal Multi-loss Fusion Network (MMML) (Wu et al. 2024): The current SOTA for CH-SIMS and MOSI, MMML utilizes pre-trained acoustic models as feature extractors for the acoustic modality.

Features	Logistic Reg.		w/ Multi. Fusion Enc.		
	Acc ₂	F1	Acc ₂	Corr	MAE
Erlangshen	72.89	72.83	86.28	82.21	27.3
COLD.	75.63	75.62	87.27	83.72	26.4
VIT-face	78.70	78.55	86.39	82.69	26.6
Ours (LC)	73.66	73.68	90.12	86.14	25.9

Table 7: Results of our feature comparison study.

Metrics

We evaluate model performance using standard metrics for each task as established in the literature. Further details can be found in the Appendix.

Results

We present our experimental results on sentiment analysis in Tables 4 and 5 and emotion recognition and sarcasm detection in Table 6. First, we note that our vanilla Multimodal Fusion Encoder before live comment feature augmentation (as indicated by “Ours”) outperforms all of our baselines in nearly every metric across all tasks. This demonstrates the strength of utilizing pre-trained encoders for all modalities in affective analysis and establishes a new SOTA on its own.

When we treat this vanilla Multimodal Fusion Encoder as a strong baseline, we see that the addition of synthetic live comment features (as indicated by “+ LC”) yields additional gains, outperforming the vanilla Multimodal Fusion Encoder (and our other baselines) by significant margins.

For Chinese sentiment analysis, we observe a 3.18-point increase in accuracy and a 3.82-point improvement in Pearson Correlation (Corr) on CH-SIMS v2, and a 0.9-point decrease in Mean Absolute Error (MAE) on CH-SIMS.

For English, in emotion recognition, we see a 2.89-point increase in weighted F1-score on MELD. In sarcasm detection, we see a 3.0-point increase in weighted F1-score on MuSTARD. And in sentiment analysis, we also see improvements, though more modest. This is likely due to the genre shift from our pre-training dataset to the English sentiment analysis benchmarks. As noted in the dataset section, MOSI and MOSEI primarily contain YouTube monologues featuring one speaker which differs significantly from the conversational content in our collected pre-training corpus.

Affective Features Comparison Study

Affective Feature Baseline Models We compare our synthetic live comment features with features from models trained for related affective analysis tasks:

Text-based Chinese sentiment features from Erlangshen-Roberta-sentiment (Erlangshen) (Zhang et al. 2022): representations from a text sentiment analysis model (Chinese-RoBERTa-wwm-ext fine-tuned on 227,347 sentiment-labeled texts from 8 datasets).

Text-based Chinese offensive language features from COLDETECTOR (Deng et al. 2022): representations from an offensive language detection model (Chinese-RoBERTa-wwm-ext fine-tuned on 37,480

sentences annotated with binary offense labels from the COLDDataset).

Image-based facial emotion recognition features from Vit-face-expression (VIT-face) (Todor Pakov 2024): representations from a facial emotion recognition model (Vision Transformer (Dosovitskiy et al. 2020) fine-tuned on 35,887 faces from the FER2013 dataset, annotated with one of seven emotions: Anger, Disgust, Fear, Happiness, Sadness, Surprise, and Neutral).

Affective Feature Experiments We evaluate the effectiveness of our synthetic live comment features through controlled experimentation in two settings. In the first, we train a logistic regression classifier on individual features to measure how well each feature discriminates our evaluation data. In the second, we augment our vanilla Multimodal Fusion Encoder with individual features to evaluate how well each feature can be leveraged by more expressive multimodal models. For these experiments, we focus on Chinese sentiment analysis with CH-SIMS v2 due to the diversity of its content and the amount of data available for fine-tuning. Additional experimental details can be found in the Appendix. We present the results of these experiments in Table 7.

Affective Feature Comparison Results When using a simple logistic regression classifier, our synthetic live comment features outperform text-based Chinese sentiment from the Erlangshen even though that model was trained on a large quantity of task-specific sentiment data. Additionally, though less performant, our features are competitive with our other baseline features from models that benefit from supervised training on related affective tasks.

The strength of our synthetic live comment features becomes apparent when we move to our Multimodal Fusion Encoder. In this setting, they outperform the baselines by wide margins across all metrics considered. We hypothesize that though our synthetic live comment features are less discriminative in a linear model for a specific affective task (like sentiment analysis), they encode rich information that can be adapted to any affective task by the more expressive multimodal models that have become the standard in AI today. Our experimental results demonstrate this effectiveness.

Conclusion

In this work, we demonstrate the efficacy of synthetic live comment features in multimodal affective analysis. By constructing the LCAffect dataset and training a multimodal encoder on it, we enable the inference of these synthetic features for any video content. Our enhanced multimodal model, augmented with synthetic live comment features, achieves new SOTA results in sentiment analysis, emotion recognition, and sarcasm detection across both English and Chinese content. This advance confirms that synthetic live comment features can effectively capture a broad spectrum of affective states, opening new avenues across diverse dimensions of affective analysis.

Acknowledgments

This research is being developed with funding from the Defense Advanced Research Projects Agency (DARPA), via the CCU Program contract HR001122C0034. The views, opinions, and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

References

- Arnab, A.; Deghani, M.; Heigold, G.; Sun, C.; Lučić, M.; and Schmid, C. 2021. ViViT: A Video Vision Transformer. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 6816–6826.
- Baevski, A.; Hsu, W.-N.; Xu, Q.; Babu, A.; Gu, J.; and Auli, M. 2022. data2vec: A General Framework for Self-supervised Learning in Speech, Vision and Language. *arXiv:2202.03555*.
- 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 Gurevych, I.; and Miyao, Y., eds., *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2236–2246. Melbourne, Australia: Association for Computational Linguistics.
- Baltrušaitis, T.; Robinson, P.; and Morency, L.-P. 2016. OpenFace: An open source facial behavior analysis toolkit. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 1–10.
- Bertasius, G.; Wang, H.; and Torresani, L. 2021. Is Space-Time Attention All You Need for Video Understanding? In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 813–824. PMLR.
- Castro, S.; Hazarika, D.; Pérez-Rosas, V.; Zimmermann, R.; Mihalcea, R.; and Poria, S. 2019. Towards Multimodal Sarcasm Detection (An ‘Obviously’ Perfect Paper). In Korhonen, A.; Traum, D.; and Márquez, L., eds., *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4619–4629. Florence, Italy: Association for Computational Linguistics.
- Chen, J.; Ding, J.; Chen, W.; and Jin, Q. 2023. Knowledge Enhanced Model for Live Video Comment Generation. In *2023 IEEE International Conference on Multimedia and Expo (ICME)*, 2267–2272.
- Conneau, A.; Khandelwal, K.; Goyal, N.; Chaudhary, V.; Wenzek, G.; Guzmán, F.; Grave, E.; Ott, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- Cui, Y.; Che, W.; Liu, T.; Qin, B.; Wang, S.; and Hu, G. 2020. Revisiting Pre-Trained Models for Chinese Natural Language Processing. In Cohn, T.; He, Y.; and Liu, Y., eds., *Findings of the Association for Computational Linguistics: EMNLP 2020*, 657–668. Online: Association for Computational Linguistics.
- Deng, J.; Zhou, J.; Sun, H.; Zheng, C.; Mi, F.; Meng, H.; and Huang, M. 2022. COLA: A Benchmark for Chinese Offensive Language Detection. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 11580–11599. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Burstein, J.; Doran, C.; and Solorio, T., eds., *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Eyben, F.; Wöllmer, M.; and Schuller, B. 2010. Opensmile: the munich versatile and fast open-source audio feature extractor. In *Proceedings of the 18th ACM International Conference on Multimedia*, MM ’10, 1459–1462. New York, NY, USA: Association for Computing Machinery. ISBN 9781605589336.
- Fu, C.-Y.; Lee, J.; Bansal, M.; and Berg, A. 2017. Video Highlight Prediction Using Audience Chat Reactions. In Palmer, M.; Hwa, R.; and Riedel, S., eds., *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 972–978. Copenhagen, Denmark: Association for Computational Linguistics.
- Ghosal, D.; Majumder, N.; Gelbukh, A.; Mihalcea, R.; and Poria, S. 2020. COSMIC: CommonSense knowledge for eMotion Identification in Conversations. In Cohn, T.; He, Y.; and Liu, Y., eds., *Findings of the Association for Computational Linguistics: EMNLP 2020*, 2470–2481. Online: Association for Computational Linguistics.
- Hanu, L.; Thewlis, J.; Asano, Y. M.; and Rupprecht, C. 2022. VTC: Improving Video-Text Retrieval with User Comments. In *ECCV*.
- Hsu, W.-N.; Bolte, B.; Tsai, Y.-H. H.; Lakhotia, K.; Salakhutdinov, R.; and Mohamed, A. 2021. HuBERT: Self-Supervised Speech Representation Learning by Masked Prediction of Hidden Units. *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, 29: 3451–3460.
- Hu, G.; Lin, T.-E.; Zhao, Y.; Lu, G.; Wu, Y.; and Li, Y. 2022. UniMSE: Towards Unified Multimodal Sentiment Analysis and Emotion Recognition. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 7837–7851. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Lalanne, J.; Bournet, R.; and Yu, Y. 2023. LiveChat: Video

- Comment Generation from Audio-Visual Multimodal Contexts. *ArXiv*, abs/2311.12826.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Liu, Y.; Yuan, Z.; Mao, H.; Liang, Z.; Yang, W.; Qiu, Y.; Cheng, T.; Li, X.; Xu, H.; and Gao, K. 2022. Make Acoustic and Visual Cues Matter: CH-SIMS v2.0 Dataset and AV-Mixup Consistent Module. In *Proceedings of the 2022 International Conference on Multimodal Interaction, ICMi '22*, 247–258. New York, NY, USA: Association for Computing Machinery. ISBN 9781450393904.
- Ma, S.; Cui, L.; Dai, D.; Wei, F.; and Sun, X. 2018. LiveBot: Generating Live Video Comments Based on Visual and Textual Contexts. In *AAAI Conference on Artificial Intelligence*.
- Niu, J.; Li, S.; Mo, S.; Yang, S.; and Fan, B. 2018. Affective Content Analysis of Online Video Clips with Live Comments in Chinese. In *2018 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computing, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBD-Com/IOP/SCI)*, 849–856.
- Poria, S.; Hazarika, D.; Majumder, N.; Naik, G.; Cambria, E.; and Mihalcea, R. 2019. MELD: A Multimodal Multi-Party Dataset for Emotion Recognition in Conversations. In Korhonen, A.; Traum, D.; and Màrquez, L., eds., *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 527–536. Florence, Italy: Association for Computational Linguistics.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. *arXiv:2103.00020*.
- Rosenbusch, H.; Evans, A. M.; and Zeelenberg, M. 2019. Multilevel Emotion Transfer on YouTube: Disentangling the Effects of Emotional Contagion and Homophily on Video Audiences. *Social Psychological and Personality Science*, 10(8): 1028–1035.
- Todor Pakov. 2024. vit-face-expression (Revision 78ed8d3).
- 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 Korhonen, A.; Traum, D.; and Màrquez, L., eds., *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 6558–6569. Florence, Italy: Association for Computational Linguistics.
- Wang, W.; Chen, J.; and Jin, Q. 2020. VideoIC: A Video Interactive Comments Dataset and Multimodal Multitask Learning for Comments Generation. In *Proceedings of the 28th ACM International Conference on Multimedia, MM '20*, 2599–2607. New York, NY, USA: Association for Computing Machinery. ISBN 9781450379885.
- Wu, Z.; Gong, Z.; Koo, J.; and Hirschberg, J. 2024. Multimodal Multi-loss Fusion Network for Sentiment Analysis. *arXiv:2308.00264*.
- Yang, Z.; Ai, L.; and Hirschberg, J. 2019. Multimodal Indicators of Humor in Videos. In *2019 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*, 538–543.
- Yi, G.; Fan, C.; Zhu, K.; Lv, Z.; Liang, S.; Wen, Z.; Pei, G.; Li, T.; and Tao, J. 2024. VLP2MSA: Expanding vision-language pre-training to multimodal sentiment analysis. *Knowledge-Based Systems*, 283: 111136.
- Yu, W.; Xu, H.; Meng, F.; Zhu, Y.; Ma, Y.; Wu, J.; Zou, J.; and Yang, K. 2020. CH-SIMS: A Chinese Multimodal Sentiment Analysis Dataset with Fine-grained Annotation of Modality. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 3718–3727. Online: Association for Computational Linguistics.
- Zadeh, A.; Zellers, R.; Pincus, E.; and Morency, L.-P. 2016. MOSI: Multimodal Corpus of Sentiment Intensity and Subjectivity Analysis in Online Opinion Videos. *arXiv:1606.06259*.
- Zhang, J.; Gan, R.; Wang, J.; Zhang, Y.; Zhang, L.; Yang, P.; Gao, X.; Wu, Z.; Dong, X.; He, J.; Zhuo, J.; Yang, Q.; Huang, Y.; Li, X.; Wu, Y.; Lu, J.; Zhu, X.; Chen, W.; Han, T.; Pan, K.; Wang, R.; Wang, H.; Wu, X.; Zeng, Z.; and Chen, C. 2022. Fengshenbang 1.0: Being the Foundation of Chinese Cognitive Intelligence. *CoRR*, abs/2209.02970.