

# Hi-EF: Benchmarking Emotion Forecasting in Human-interaction

Haoran Wang<sup>1\*</sup>, Xinji Mai<sup>1\*</sup>, Zeng Tao<sup>1\*</sup>, Junxiong Lin<sup>1</sup>, Xuan Tong<sup>1</sup>, Ivy Pan<sup>2</sup>, Shaoqi Yan<sup>3</sup>,  
Yan Wang<sup>4†</sup>, Shuyong Gao<sup>5,6†</sup>

<sup>1</sup>College of Intelligent Robotics and Advanced Manufacturing, Fudan University

<sup>2</sup>The University of Hong Kong

<sup>3</sup>School of Electrical and Electronic Engineering, Shanghai Institute of Technology

<sup>4</sup>East China Normal University

<sup>5</sup>Carnegie Mellon University

<sup>6</sup>Keenon Robotics Co., Ltd

hrwang23@m.fudan.edu.cn, xjmai23@m.fudan.edu.cn, ztao19@fudan.edu.cn

## Abstract

Affective Forecasting is a psychology task that involves predicting an individual’s future emotional responses, often hampered by reliance on external factors leading to inaccuracies, and typically remains at a qualitative analysis stage. To address these challenges, we narrow the scope of Affective Forecasting by introducing the concept of Human-interaction-based Emotion Forecasting (EF). This task is set within the context of a two-party interaction, positing that an individual’s emotions are significantly influenced by their interaction partner’s emotional expressions and informational cues. This dynamic provides a structured perspective for exploring the patterns of emotional change, thereby enhancing the feasibility of emotion forecasting.

## Introduction

Affective computing is a branch of computer science and artificial intelligence that aims to enable computers to recognize, interpret, process, and simulate human emotions. Proposed by Rosalind Picard (Picard 2000), this field combines psychology, computer science, and cognitive science to imbue machines with emotional intelligence. Current research focuses primarily on calculating and analyzing an individual’s present emotional state using facial expressions, voice, and physiological signals, which is Emotion Recognition (ER) (Lo et al. 2020; Tran et al. 2015; Liu et al. 2020; Ma et al. 2019; Li et al. 2023; Zhang et al. 2022; Liu et al. 2022). However, despite significant progress in understanding and responding to current emotions, affective computing often overlooks the prediction of future emotions. Affective Forecasting (Wilson and Gilbert 2003) refers to the process by which individuals predict their future emotional reactions to various events. This cognitive mechanism plays a crucial role in decision-making, influencing choices by anticipating potential emotional outcomes. However, empirical studies (Wilson and Gilbert 2003; Gilbert and Wilson 2007; Wilson and Gilbert 2005) have shown that individuals often exhibit significant inaccuracies

in their affective forecasts. For instance, Gilbert and Wilson (Gilbert and Wilson 2007) highlight the prevalence of “impact bias”, where people tend to overestimate the intensity and duration of their emotional responses to future events. This bias is evident in both positive and negative contexts, such as overestimating happiness from future achievements or distress from anticipated setbacks. Moreover, factors such as focalism, where individuals focus too narrowly on a single event, and underestimating emotional adaptation, contribute to these forecasting errors (Wilson and Gilbert 2005). Understanding these predictive inaccuracies is critical, as they can lead to suboptimal decision-making in personal, professional, and social contexts. Recent advancements in this field aim to mitigate these biases by exploring the underlying cognitive and emotional mechanisms that drive affective forecasting (Wilson and Gilbert 2003).

Research on affective forecasting primarily focuses on psychology and economics. However, we aim to redefine the paradigm of affective forecasting to transform this problem into a deep learning challenge. We propose a new task: Emotion Forecasting (EF), predicting one party’s future emotion during a two-party interaction by using short-term contextual information and the current emotional state of the other party. First, emotional exchange is bidirectional: one person’s emotions can affect the other’s, allowing for more accurate predictions of emotional changes (Hatfield, Cacioppo, and Rapson 1993). Second, two-party interactions occur over short time frames, making the emotional predictions more reliable as responses to recent events are more direct and intense (Loewenstein and Prelec 1992). Lastly, the dynamic context of interactions provides rich background information, such as conversation content, tone, and emotional shifts, enhancing the precision of affective forecasting (Barsade 2002). By remodelling affective forecasting to focus on interpersonal emotional dynamics, we can better understand and predict emotional changes in relationships, providing new tools and methods for emotion management and support (Wilson and Gilbert 2005; Gilbert and Wilson 2007). And we illustrate the difference between ER task and EF task in Figure 1.

Based on these, we list these criteria to construct datasets for solving the EF task: (1) Multi-modal information: By

\*These authors contributed equally.

†Corresponding authors.

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




Figure 1: Comparing EF task with ER task.

analyzing facial expressions, conversational content, and vocal intonations, we can better assess a person’s emotions, thereby enabling more accurate predictions of the emotions the other party may experience. (2) Strong-Interactive Scenes: Videos should be sourced from highly interactive scenarios, enabling enough amount of interaction data for the dataset. (3) Multilayered-Contextual Information: The data format needs to include three layers of contextual information: contextual information, one party’s current emotional state, and the emotion of the other party that needs to be forecasted.

In adherence to these guidelines, we have developed the Human-interaction-based Emotion Forecasting (Hi-EF) dataset. Hi-EF boasts 3,069 Multilayered-Contextual Interaction Samples (MCIS), featuring in 5,242 video clips with abundant affect-related labels. To meet the stringent criteria for data collection, we have devised a comprehensive two-stage strategy to acquire a sufficient pool of candidate MCIS. Furthermore, we have implemented a reliable three-round annotation process to ensure high-quality labeling, given the complexity of the emotional information encapsulated in the labels to make sure the emotion is annotated with high quality. For the challenging EF task, we proposed the baseline method and made comprehensive experiments to demonstrate the feasibility of EF task. Besides that, more experiments of other affective computing tasks are also provided.

Our work makes the following significant contributions:

- We have transformed Affective Forecasting from a psychological and economic problem into a deep learning task and developed an Emotion Forecasting paradigm based on two-party interactions. This introduces a novel task in the field of affective computing: Emotion Forecasting (EF). Previously, affective computing primarily focused on recognizing current emotions. Our work provides a new perspective by aiming to predict future emotions.
- To address this novel EF task, we have constructed a Human-interaction-based Emotion Forecasting dataset (Hi-EF). This dataset introduces a unique data format, Multilayered-Contextual Interaction Samples (MCIS), specifically designed for the EF task. MCIS encompasses short-term contextual information, one party’s current emotional state, and the other party’s future emotional state. It includes three modalities and rich affect-related labels for both parties. Additionally, our dataset can also be utilized for other affective computing tasks.
- To generate MCIS, we developed a construction procedure including candidate MCIS generation, anomaly detection and multifaceted-assisted reliable annotation.

- We have designed a paradigm for the EF task and, based on this paradigm, established a baseline for the EF task. Extensive experiments and analyses were conducted on the Hi-EF dataset, demonstrating the feasibility of the EF task and the significance of the Hi-EF dataset. We will soon release our method’s code and the dataset for further research and validation.

## Significance of EF Task

Unlike traditional Affective Computing tasks that focus on recognizing current emotions, the EF task aims to predict future potential emotions based on interactional context. Accurate and reasonable predictions in this domain have a wide array of applications. We outline several future research directions and applications:

- **Individual Emotion Modeling.** By analyzing extensive data from a person’s interactions, we can model their emotional responses, identifying what types of interactions are likely to elicit specific emotions. This modeling is useful in the medical field for treating conditions such as depression, enabling personalized therapeutic approaches.
- **Anthropomorphic Emotion Generation.** In an interactive scenario, by evaluating the emotional state and contextual information of the other party, we can predict the appropriate emotional response a human should exhibit. This ability enables the generation of anthropomorphic emotions, which is advantageous for developing emotionally intelligent robots. Such robots can respond in a more human-like and empathetic manner, thereby enhancing their effectiveness in roles that require human interaction.

## Relevant Datasets of Emotion Forecasting: Emotion Recognition Datasets

Since the EF task is a novel task and there are no existing EF datasets, we will introduce the ER datasets that are most relevant to the EF task to aid in understanding our Hi-EF dataset. These datasets are presented in Table 1. By presenting these datasets, we aim to facilitate a better understanding of the distinctions between our Emotion Forecasting task, the Hi-EF dataset, and traditional Emotion Recognition tasks and datasets. This comparison will help clarify the unique aspects of our work. ER datasets have two main-stream types: video-driven ER datasets and multi-modal ER datasets. The primary classification is based on the dominant modalities that constitute the datasets and their corresponding related tasks, including Dynamic Facial Expression Recognition (DFER), Multi-modal Emotion Recognition (MER), and Emotion Recognition in Conversation (ERC).

**Video-driven Emotion Recognition Datasets.** Video-driven Emotion Recognition datasets are dedicated to identifying human emotions from dynamic visual information. These datasets includes DFER datasets, typically originating from controlled laboratory settings, as exemplified by the CK+ (Lucey et al. 2010) and Oulu-CASIA (Zhao et al. 2011) datasets. Additionally, there are DFER datasets in the wild, such as AFEW (Dhall et al. 2018), Aff-wild (Kollias et al. 2019), AFEW-VA (Kossaifi et al. 2017), CAER (Lee et al.

ID	Datasets	Class	Scene	Text Description		Modality	Affective Component				DFER	MER	ERC	EF
				Fac.	Scene		Polar.	Inten.	Emo.	Uncer.				
1	IEMOCAP (2008) (Busso et al. 2008)	5	1	-	-	V. &A. &T.	-	-	✓	-	-	✓	-	-
2	HOW (2011) (Morency, Mihalcea, and Doshi 2011)	3	1	-	-	V. &A. &T.	-	-	✓	-	-	✓	-	-
3	CMU-MOSEI (2018) (Zadeh et al. 2018)	6	1	-	-	V. &A. &T.	-	-	✓	-	-	✓	-	-
4	MELD (2018) (Poria et al. 2018)	7	1	-	-	V. &A. &T.	-	-	✓	-	-	✓	✓	-
5	CAER (2019) (Lee et al. 2019a)	7	1	-	-	V.	-	-	✓	-	✓	-	-	-
6	DFEW (2020) (Jiang et al. 2020)	7	1	-	-	V.	-	-	✓	-	✓	-	-	-
7	FERV39K (2022) (Wang et al. 2022a)	7	22	-	-	V.	-	-	✓	-	✓	-	-	-
8	MAFW (2022) (Liu et al. 2022)	11	1	✓	✓	V. &A. &T.	-	-	✓	-	✓	✓	-	-
9	Hi-EF (2025)	7	4	✓	✓	V. &A. &T.	✓	✓	✓	✓	✓	✓	✓	✓

Table 1: Summary of existing Affective Computing datasets (ID:1-4), ERC dataset (ID:4) and DFER datasets (ID:4-8) and our built Hi-EF. (Fac. = Facial movements; Polar. = Interaction polarity; Inten. = Interaction intensity; Emo. = Emotion; Uncer. = Uncertainty; V. = Video; A. = Audio; T. = Text.

2019b), DFEW (Jiang et al. 2020), FERV39k (Wang et al. 2022a). Datasets like MAFW (Liu et al. 2022), which offer diverse modalities and emotion labels while is still video-driven, since the annotation is mainly based on the visual modality, while it could also be used in MER task.

**Multi-modal Emotion Recognition Datasets.** MER datasets focus not just the facial expression. Instead, they pay more attention to fuse the different modalities to analyze the whole emotional state of a person. And the class is not only the expression, but also sentiment including positive, neutral and negative and Activation-Valence-Dominance like IEMOCAP (Busso et al. 2008), HOW (Morency, Mihalcea, and Doshi 2011), CMU-MOSEI (Zadeh et al. 2018) and MELD (Poria et al. 2018). Especially for MELD, it is text-based Emotion Recognition dataset designed for ERC.

The EF task differs from the ER task, although it does involve elements of ER. In the EF task, we first need to recognize one party’s emotions and then, using the contextual information, predict the potential emotions of the other party. This task structure results in the Hi-EF dataset’s MCIS format being entirely different from typical ER formats. For example, in DFER tasks, the dataset usually consists of single videos labeled with the expressions of the person in the video, typically using a seven-category classification for expressions. In contrast, MCIS includes information from both parties’ interactions. It comprises four consecutive videos: the first two videos provide contextual information, the third video presents one party’s emotional state, and the fourth video is where we need to predict the other party’s emotional state. More details about the EF task, Hi-EF, and MCIS will be elaborated in Section .

## The Hi-EF Dataset

### Paradigm of Emotion Forecasting

The EF task entails leveraging contextual information from the interaction between two parties and the emotions of one party to predict the potential future emotions of the other. Imagine a scenario where Party A shares a joyful experience with Party B, who shows significant interest. If A displays happiness during this interaction, it is highly probable that B will also exhibit happiness subsequently. This scenario can be divided into three layers. The first layer encompasses context-

tual information, serving as background and supplementary data. The second layer pertains to Party A’s emotional state, which directly influences Party B’s emotions. The third layer comprises Party B’s own emotions, which are what we aim to forecast. Therefore, we believe the EF task should incorporate the following three functions: understanding contextual information, recognizing Party A’s emotional state, and integrating these two layers of information for forecasting. These functions are defined as follows:

- Contextual Interaction Fusion Function  $I = f_{CI}(\cdot)$ : This function merges various layers of contextual information, providing a refined backdrop for emotion interpretation and forecasting.  $I$  indicates contextual information, such as non-verbal cues and the level of participation of the participants.
- Current Party Emotion Recognition Function  $E_A = f_{ER}(\cdot)$ : It recognizes Party A’s current emotional state from multi-modal cues.
- Future Emotion Forecasting Function  $E_B = f_{EF}(I, E_A)$ : Leveraging insights from  $f_{CI}$  and  $f_{ER}$ , this function forecasts the probable future emotion of Party B.  $E_B$  is the future emotion of Party B which will be forecasted.

It is worth noting that in our model architecture, the three functions are not explicitly defined as independent modules but are integrated into the overall design, which still follows the proposed paradigm.

### Design of Multilayered-Contextual Interaction Sample (MCIS)

To adapt to the EF task, we needed to define a data format. In the introduction, we also outlined our requirements for the dataset: Multi-modal, Strong-Interactive Scene, and Multilayered-Contextual information. Therefore, we designed the **Multilayered-Contextual Interaction Sample (MCIS)** as the data format for the Hi-EF dataset. Figure 2 illustrates MCIS. An MCIS primarily consists of the following components: Contextual Information: Video clip I and II; Current Information of Party A: Video clip III; Forecasted Information of Party B: Video clip IV. We provide three modalities for each video clip: video, text (utterance), and audio. Additionally, we provide annotation information for



Figure 2: Overview of MCIS: In order to realize the task of forecasting emotion during the interaction process, we design a new form of data MCIS with multilayered-contextual information which includes contextual information (clip I and II), current emotion state of Party A (clip III) and future emotion of Party B (clip IV). To better evaluate Party A’s emotional state and more accurately predict Party B’s emotions, we provide three modalities for MCIS: video, audio and text.

clips III and IV. It is important to note that for the Contextual Information part, we only use two clips. We did not choose more than two videos for contextual information mainly due to the following considerations: 1) In Affective Forecasting (Frederick, Loewenstein, and O’donoghue 2002; Loewenstein and Prelec 1992; Ainslie 1975; Laibson 1997), overly long contextual information might introduce noise, affecting the affective forecast. 2) Excessively long contextual information significantly increases computational costs. 3) Contextual information serves as supplementary data, while the primary data for EF is the current emotional state of Party A (Loewenstein and Prelec 1992).

### Dataset Construction Procedure

In this section, we present an overview of our database construction process, which consists of three stages: raw data collection and candidate clip generation, candidate MCIS creation, and MCIS annotation.

**Data Collection and Candidate MCIS Generation.** We source raw videos from shows rich in interactive scenes and extract subtitles to segment them into clips based on utterance timestamps. Since the initial clips contain noise, we refine them using facial detection and a state-of-the-art ER model (Tao et al. 2024) to isolate emotionally relevant segments for MCIS construction. Each MCIS is formed by combining three preceding clips with the emotionally significant one, as illustrated in Figure 3 a). This process yields a large number of candidate MCIS samples.

**Anomaly Detection Rules.** Our utterance-based MCIS generation strategy occasionally yields abnormal clips. To address this, we design an anomaly detection mechanism to ensure annotation quality. As illustrated in Figure 3 b), we define three anomaly types: (1) *Facial anomalies* — clips lacking a frontal face view of the speaker for at least half their duration; (2) *Scene transitions* — MCIS whose final two clips belong to different scenes; (3) *Interaction anomalies* — MCIS featuring the same individual in the last two clips. We implement an algorithm to pre-screen the first and third types, filtering out about half of the invalid MCIS. Given the limitations of

automatic detection under strict criteria, manual inspection remains indispensable. Therefore, we combine algorithmic and manual detection to enhance overall accuracy, which also aligns with the subsequent annotation step.

**Multifaceted-Assisted Reliable Annotation.** The Hi-EF database adopts a novel data format and undergoes a rigorous three-round manual annotation process (Figure 3 b–c) involving video, audio, and text modalities. In the first round, crowdsourced workers apply anomaly criteria to select valid MCIS. In the second round, professional annotators validate and refine these samples, ensuring consistency and annotating affective components and textual descriptions. The final round re-assesses the double-checked MCIS, focusing on the last two clips. Each MCIS is labeled by five independent annotators, and a professor reviews these annotations with multifaceted auxiliary information to determine or refine the final label. We provide comprehensive annotation details, including textual descriptions (facial actions, scene information) and emotional components (interaction polarity, intensity, emotion categories, and uncertainty). These annotations collectively support reliable emotion categorization and can facilitate other affective computing tasks. Detailed annotation guidelines are provided in the **Appendix**.

### Dataset Statistics

The Hi-EF dataset has been meticulously compiled from TV shows, amassing over 100 hours of content and dividing it into more than 60,000 video clips. Following preprocessing, 30,000 clips containing emotion have been retained, and through a series of data refinement, a half of the candidates are eliminated. Manual judgment has further refined the dataset, resulting in a selection of 3,069 MCIS and 5,242 video clips with high-quality annotations. Additional dataset statistics are provided in the **Appendix**

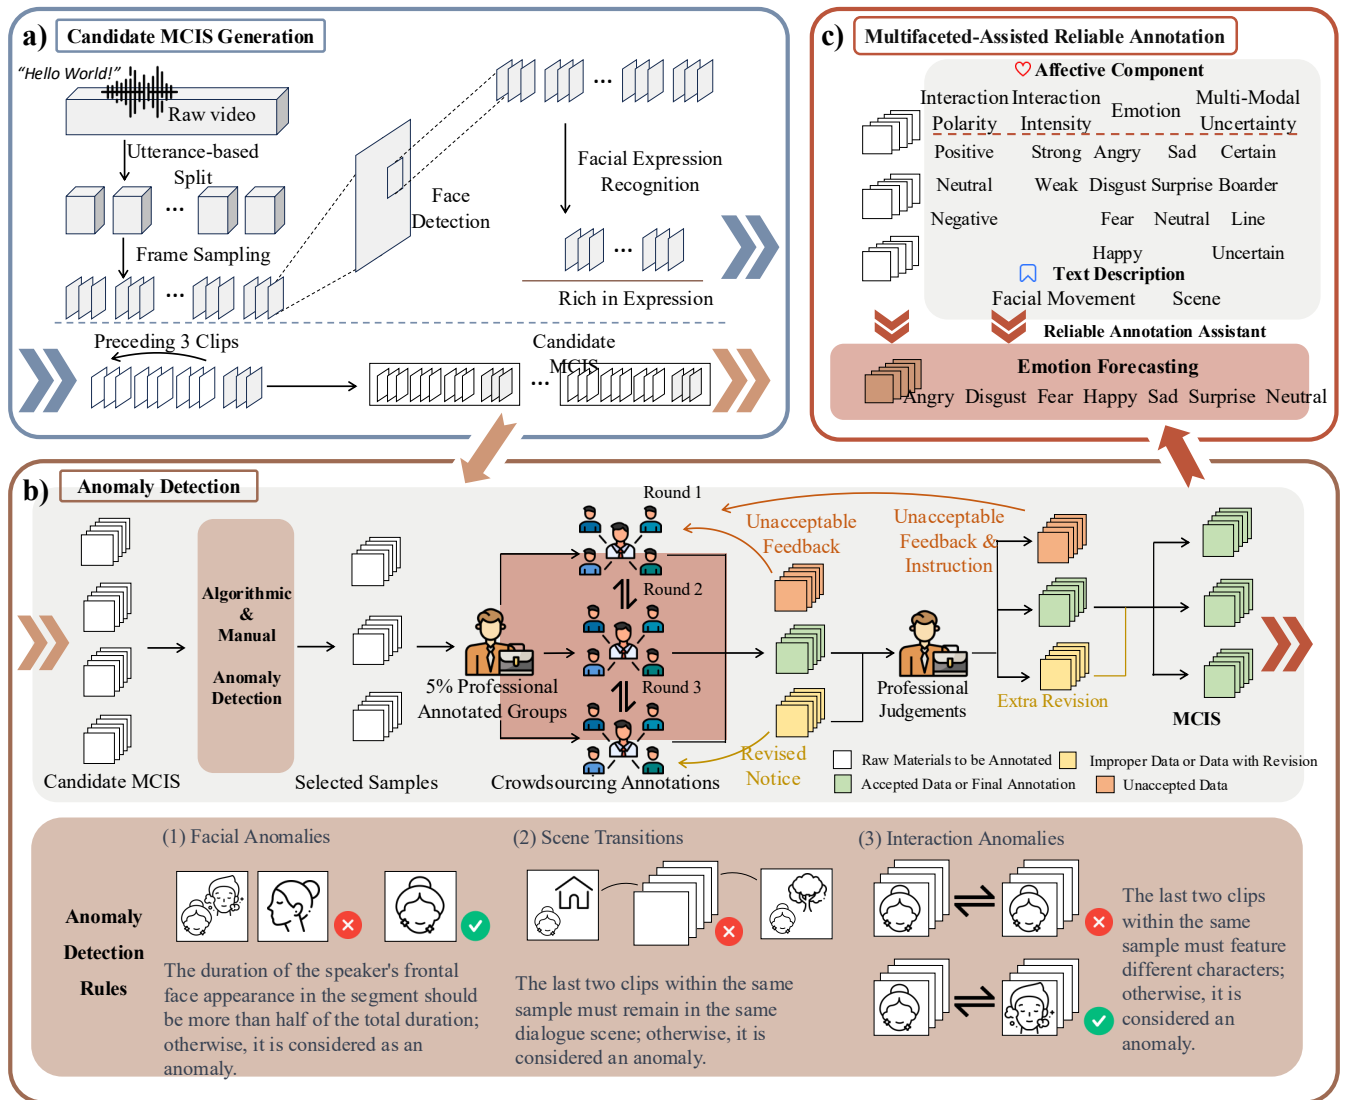


Figure 3: Overview of the Construction Procedure for the Hi-EF Dataset. a) Shows the process of generating candidate MCIS. b) Illustrates the anomaly detection and the rules and multi-facet assisted reliable annotation for candidate MCIS. c) Details the annotation process. All other annotation information is intended to support the more accurate labeling of emotions. Naturally, we will provide all the annotation details within the dataset.

## Experiment

### Experiment Setup

**Hi-EF Protocol.** To establish a robust benchmark for the EF task within the Hi-EF dataset, we partitioned the dataset, comprising 3,069 MCIS, into training (70%) and testing (30%) sets, with the training set further segmented into a validation subset. It is particularly important to note that, unlike traditional data formats, randomly splitting the training and test sets by proportion may result in overlapping clips between MCISs, leading to data leakage. Therefore, we implemented additional criteria to avoid this situation. For the EF task, our experiments unfolded as follows: We employed 2,150 MCIS for training. Subsequently, we use different encoders and fusion strategies for experiments.

**Implementation Details.** All experiments were implemented using PyTorch on Nvidia RTX 3090 GPUs. We set the learning rates to be in the range of  $1e-2$  to  $1e-5$ , weight decay to  $1e-6$ , and used a batch size of 8. All models were trained from scratch using optimizer (Adam) for 50 epochs, with uniform frame subsampling at an interval of 8 frames. Furthermore, due to the limited availability of data in the context of Hi-EF, we employed several data augmentation techniques, including random cropping, changes in lighting conditions, and image flipping. To reduce the computational demands, all cropped images were resized to  $112 \times 112$ , and the entire images were resized to  $112 \times 168$ .

**Evaluation Metrics.** Drawing inspiration from prevalent Affective Computing tasks (Schuller et al. 2010; Jiang et al.

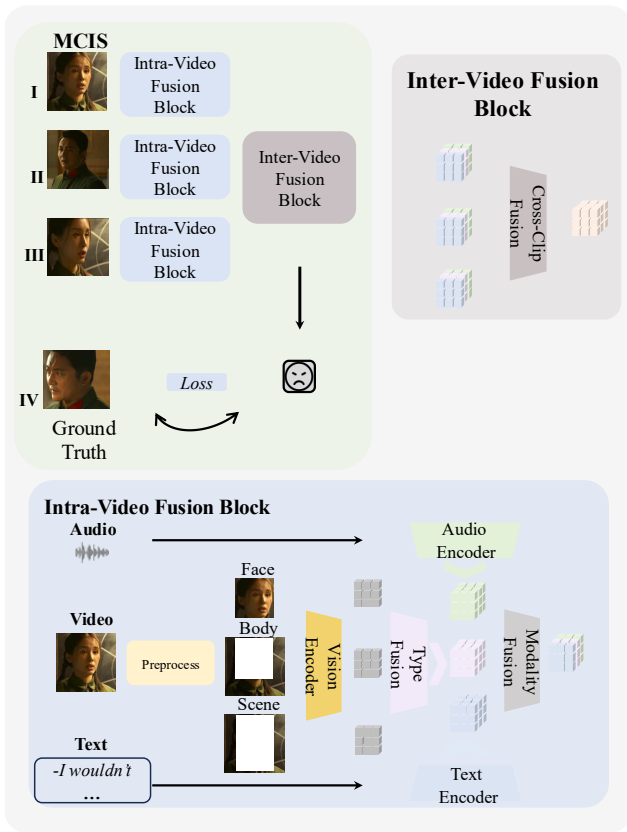


Figure 4: Overview of our model architecture. The first three clips in MCIS are respectively input into the intra-video fusion block, where the obtained features are then fed into the inter-video fusion block for prediction. The final prediction result is compared with the emotion label of clip IV in MCIS to calculate the loss. In the intra-video fusion block, video frames are divided into facial, posture, and scene parts, which are separately input into the vision encoder. These features are then fused to form the video modality information. Meanwhile, audio and text are input into the audio encoder and text encoder, respectively. After obtaining the features from the three modalities, they are fused to produce the feature representation of a single video. The inter-video fusion block primarily fuses the features of the three clips.

2020; Liu et al. 2022; Wang et al. 2022a,b; Lee et al. 2019b; Zhang et al. 2017), we selected two traditional metrics for evaluation: Weighted Average Recall (WAR) and Unweighted Average Recall (UAR).

**Model Architecture.** Our model architecture, as depicted in Figure 4, is divided into two main parts: intra-video fusion block and inter-video fusion block. Firstly, intra-video fusion block is further divided into two segments. One clip involves the alignment and fusion of the three modalities: video, audio, and text. The other clips focuses on the video modality, where we divide video frames into three parts: facial expressions, body posture, and scene information, as all three contribute to the overall emotional state of a person.

For the vision encoder, we use ResNet18(He et al. 2016), I3D(Carreira et al. 2019), and ViT(Dosovitskiy et al. 2020). For the audio encoder, we employ the commonly used Audio-CLIP(Guzhov et al. 2022) Encoder, and for the text encoder, we use the CLIP(Radford et al. 2021) text encoder. Regarding the intra-video fusion strategy, we adopt two approaches: one is the average method, and the other is fusion using a transformer. For inter-video fusion block, we primarily use average, transformer, and a combination of LSTM(Graves and Graves 2012) and Transformer(Vaswani 2017).

## Experimental Results

**Experimental Results of EF Task.** We present the experimental results of the EF task in Table 2. Our main goal is to study the impact of different vision encoders and intra-/inter-video fusion strategies on EF performance. We first compare the vision encoders. Since ResNet18 and I3D cannot be aligned with the other two modalities, we only report video-only results for them, using simple fusion strategies; consequently, they do not achieve competitive performance. We then focus on ViT-based models and analyze the effect of modalities. The results show that the video modality contributes most to emotion analysis. Adding text alone degrades performance, as the text is treated as noise, whereas adding audio enables the model to better interpret the textual information. Using all three modalities improves WAR by 3.82% and UAR by 1.33% over the video-only setting, indicating that our modality choice is appropriate for the EF task.

**Effect of Intra-Video Fusion.** During two-party interactions, emotional influence arises not only from facial expressions but also from body posture and the surrounding scene. We therefore compare using only facial information with using the combination of face, body, and scene. The combined setting yields higher prediction accuracy, confirming that these factors jointly affect emotional changes. We also study different fusion strategies for these three types of information. The results show that transformer-based fusion clearly outperforms simple averaging, as it can focus on the most informative emotional cues.

**Effect of Inter-Video Fusion.** Since MCIS contains multiple layers of information, an appropriate inter-video fusion strategy is crucial. We compare three strategies: average, Transformer, and LSTM+Transformer. LSTM+Transformer performs best, where LSTM captures temporal dependencies among the three clips and the transformer highlights key signals that drive emotional changes. Their combination leads to the most accurate prediction of the other party’s potential emotions.

**Effect of MCIS.** We conduct ablation experiments on MCIS to study the impact of different input configurations, as reported in Table 3. With a single input clip, clip III achieves the best results. With two clips, the combination of clips II and III performs best, indicating that more recent information is more beneficial for EF. When three clips are used, the performance surpasses that of the II+III setting, with WAR and UAR improvements of 2.64% and 2.33%, respectively, demonstrating the effectiveness of the MCIS design.

**Derived Tasks of Hi-EF.** Beyond EF, the Hi-EF dataset can also support other affective computing tasks, such as DFER,

	Video			Text	Audio	Inter Fusion	WAR	UAR
	Face	Body	Scene					
<b>Resnet18 Backbone</b>								
R18 (He et al. 2016)	$\times$	$\times$	$\times$	$\times$	$\times$	Average	22.22	17.02
R18	R18	R18	Average	$\times$	$\times$	Average	23.15	18.58
R18	R18	R18	Transformer	$\times$	$\times$	Average	32.41	19.75
R18	R18	R18	Average	$\times$	$\times$	Transformer	28.70	21.09
<b>I3D Backbone</b>								
I3D(Carreira et al. 2019)	$\times$	$\times$	$\times$	$\times$	$\times$	Average	28.57	21.73
I3D	I3D	I3D	Average	$\times$	$\times$	Average	29.62	21.16
I3D	I3D	I3D	Transformer	$\times$	$\times$	Average	29.63	21.19
I3D	I3D	I3D	Average	$\times$	$\times$	Transformer	29.62	22.99
<b>ViT Backbone</b>								
ViT (Dosovitskiy et al. 2020)	ViT	ViT	Transformer	$\times$	$\times$	LSTM+Transformer	31.37	22.39
ViT	ViT	ViT	Transformer	$\times$	AudioCLIP	LSTM+Transformer	33.34	21.27
ViT	ViT	ViT	Transformer	CLIP	$\times$	LSTM+Transformer	31.03	21.80
ViT	ViT	ViT	Average	CLIP	AudioCLIP	Average	31.44	18.55
ViT	ViT	ViT	Transformer	CLIP	AudioCLIP	Average	33.38	21.71
ViT	ViT	ViT	Average	CLIP	AudioCLIP	LSTM+Transformer	34.49	23.19
ViT	ViT	ViT	Transformer	CLIP	AudioCLIP	LSTM+Transformer	<b>35.19</b>	<b>23.72</b>

Table 2: The comparison of EF experiments result(%) with different modalities, backbones, and intra and inter fusion strategies. **Bold** indicates the best performance, while underlining indicates the second-best performance.

Input	Video	Intra Fusion	Text	Audio	Inter Fusion	WAR	UAR
I						29.50	19.38
II	ViT	Transformer	CLIP	AudioCLIP	\	30.89	20.55
III						31.44	20.80
I + II						31.86	20.93
I + III	ViT	Transformer	CLIP	AudioCLIP	LSTM + Transformer	32.41	21.13
II + III						<u>32.55</u>	<u>21.39</u>
I + II + III	ViT	Transformer	CLIP	AudioCLIP	LSTM + Transformer	<b>35.19</b>	<b>23.72</b>

Table 3: Different input strategies of EF experiments results(%) on Hi-EF. **Bold** indicates the best performance, while underlining indicates the second-best performance.

MER, and ERC. The corresponding experimental results are reported in the **Appendix**.

## Conclusions and Discussion

In this paper, we have transformed the task of Affective Forecasting from a psychological and economic problem into a deep learning challenge. We define a new task called Emotion Forecasting, which focuses on forecasting emotions during two-party interactions. To facilitate this task, we propose the Hi-EF dataset and a baseline. The dataset features a novel data format, MCIS, which consists of multilayered-contextual interactive information between two parties. We provide rich emotion-related annotations and three modalities for each MCIS. Our primary contribution lies in designing the EF task and creating a suitable dataset and data format to accommodate this task, ultimately demonstrating its feasibility. We believe that the EF task and the Hi-EF dataset will provide a new perspective for future research in affective computing.

**Limitations.** Our baseline model does not employ a wide

variety of model architectures, and the fusion strategies used are simple and commonly adopted methods without much innovation. Our primary contribution lies in designing the EF task and creating a suitable dataset and data format to accommodate this task, ultimately demonstrating its feasibility. We believe that the EF task and the Hi-EF dataset will provide a new perspective for future research in affective computing. In the future, we plan to expand the Hi-EF dataset, test more models, and continuously provide experimental results.

## Acknowledgements

This work was supported by National Natural Science Foundation of China (No.62576109, 62072112), National Key Research and Development Program of China (2023YFC3604802) and National Natural Science Foundation of China (No. 62406075)

## References

- Ainslie, G. 1975. Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological bulletin*, 82(4): 463.
- Barsade, S. G. 2002. The ripple effect: Emotional contagion and its influence on group behavior. *Administrative science quarterly*, 47(4): 644–675.
- Busso, C.; Bulut, M.; Lee, C.-C.; Kazemzadeh, A.; Mower, E.; Kim, S.; Chang, J. N.; Lee, S.; and Narayanan, S. S. 2008. IEMOCAP: Interactive emotional dyadic motion capture database. *Language resources and evaluation*, 42: 335–359.
- Carreira, J.; Noland, E.; Hillier, C.; and Zisserman, A. 2019. A short note on the kinetics-700 human action dataset. *arXiv preprint arXiv:1907.06987*.
- Dhall, A.; Kaur, A.; Goecke, R.; and Gedeon, T. 2018. Emotiv 2018: Audio-video, student engagement and group-level affect prediction. In *Proceedings of the 20th ACM International Conference on Multimodal Interaction*, 653–656.
- 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*.
- Frederick, S.; Loewenstein, G.; and O’donoghue, T. 2002. Time discounting and time preference: A critical review. *Journal of economic literature*, 40(2): 351–401.
- Gilbert, D. T.; and Wilson, T. D. 2007. Prospect: Experiencing the future. *Science*, 317(5843): 1351–1354.
- Graves, A.; and Graves, A. 2012. Long short-term memory. *Supervised sequence labelling with recurrent neural networks*, 37–45.
- Guzhov, A.; Raue, F.; Hees, J.; and Dengel, A. 2022. Audio-clip: Extending clip to image, text and audio. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 976–980. IEEE.
- Hatfield, E.; Cacioppo, J. T.; and Rapson, R. L. 1993. Emotional contagion. *Current directions in psychological science*, 2(3): 96–100.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Jiang, X.; Zong, Y.; Zheng, W.; Tang, C.; Xia, W.; Lu, C.; and Liu, J. 2020. Dfew: A large-scale database for recognizing dynamic facial expressions in the wild. In *Proceedings of the 28th ACM international conference on multimedia*, 2881–2889.
- Kollias, D.; Tzirakis, P.; Nicolaou, M. A.; Papaioannou, A.; Zhao, G.; Schuller, B.; Kotsia, I.; and Zafeiriou, S. 2019. Deep affect prediction in-the-wild: Aff-wild database and challenge, deep architectures, and beyond. *International Journal of Computer Vision*, 127(6-7): 907–929.
- Kossaifi, J.; Tzimiropoulos, G.; Todorovic, S.; and Pantic, M. 2017. AFEW-VA database for valence and arousal estimation in-the-wild. *Image and Vision Computing*, 65: 23–36.
- Laibson, D. 1997. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2): 443–478.
- Lee, J.; Kim, S.; Kim, S.; Park, J.; and Sohn, K. 2019a. Context-Aware Emotion Recognition Networks. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Lee, J.; Kim, S.; Kim, S.; Park, J.; and Sohn, K. 2019b. Context-aware emotion recognition networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10143–10152.
- Li, H.; Niu, H.; Zhu, Z.; and Zhao, F. 2023. CLIPER: A Unified Vision-Language Framework for In-the-Wild Facial Expression Recognition. *arXiv preprint arXiv:2303.00193*.
- Liu, D.; Ouyang, X.; Xu, S.; Zhou, P.; He, K.; and Wen, S. 2020. SAANet: Siamese action-units attention network for improving dynamic facial expression recognition. *Neurocomputing*, 413: 145–157.
- Liu, Y.; Dai, W.; Feng, C.; Wang, W.; Yin, G.; Zeng, J.; and Shan, S. 2022. MAFW: A Large-scale, Multi-modal, Compound Affective Database for Dynamic Facial Expression Recognition in the Wild. In *Proceedings of the 30th ACM International Conference on Multimedia*, 24–32.
- Lo, L.; Xie, H.-X.; Shuai, H.-H.; and Cheng, W.-H. 2020. MER-GCN: Micro-expression recognition based on relation modeling with graph convolutional networks. In *2020 IEEE conference on multimedia information processing and retrieval (MIPR)*, 79–84. IEEE.
- Loewenstein, G.; and Prelec, D. 1992. Anomalies in intertemporal choice: Evidence and an interpretation. *The Quarterly Journal of Economics*, 107(2): 573–597.
- Lucey, P.; Cohn, J. F.; Kanade, T.; Saragih, J.; Ambadar, Z.; and Matthews, I. 2010. The extended cohn-kanade dataset (ck+): A complete dataset for action unit and emotion-specified expression. In *2010 IEEE computer society conference on computer vision and pattern recognition-workshops*, 94–101. IEEE.
- Ma, J.; Tang, H.; Zheng, W.-L.; and Lu, B.-L. 2019. Emotion recognition using multimodal residual LSTM network. In *Proceedings of the 27th ACM international conference on multimedia*, 176–183.
- Morency, L.-P.; Mihalcea, R.; and Doshi, P. 2011. Towards multimodal sentiment analysis: Harvesting opinions from the web. In *Proceedings of the 13th international conference on multimodal interfaces*, 169–176.
- Picard, R. W. 2000. *Affective computing*. MIT press.
- Poria, S.; Hazarika, D.; Majumder, N.; Naik, G.; Cambria, E.; and Mihalcea, R. 2018. Meld: A multimodal multi-party dataset for emotion recognition in conversations. *arXiv preprint arXiv:1810.02508*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.

- Schuller, B.; Vlasenko, B.; Eyben, F.; Wöllmer, M.; Stuhlsatz, A.; Wendemuth, A.; and Rigoll, G. 2010. Cross-corpus acoustic emotion recognition: Variances and strategies. *IEEE Transactions on Affective Computing*, 1(2): 119–131.
- Tao, Z.; Wang, Y.; Lin, J.; Wang, H.; Mai, X.; Yu, J.; Tong, X.; Zhou, Z.; Yan, S.; Zhao, Q.; et al. 2024. A3lign-DFER: Pioneering Comprehensive Dynamic Affective Alignment for Dynamic Facial Expression Recognition with CLIP. *arXiv preprint arXiv:2403.04294*.
- Tran, D.; Bourdev, L.; Fergus, R.; Torresani, L.; and Paluri, M. 2015. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, 4489–4497.
- Vaswani, A. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.
- Wang, Y.; Sun, Y.; Huang, Y.; Liu, Z.; Gao, S.; Zhang, W.; Ge, W.; and Zhang, W. 2022a. Ferv39k: A large-scale multi-scene dataset for facial expression recognition in videos. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 20922–20931.
- Wang, Y.; Sun, Y.; Song, W.; Gao, S.; Huang, Y.; Chen, Z.; Ge, W.; and Zhang, W. 2022b. Dpcnet: Dual path multi-excitation collaborative network for facial expression representation learning in videos. In *Proceedings of the 30th ACM International Conference on Multimedia*, 101–110.
- Wilson, T. D.; and Gilbert, D. T. 2003. Affective forecasting. *Advances in experimental social psychology*, 35(35): 345–411.
- Wilson, T. D.; and Gilbert, D. T. 2005. Affective forecasting: Knowing what to want. *Current directions in psychological science*, 14(3): 131–134.
- Zadeh, A. B.; 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.
- Zhang, K.; Huang, Y.; Du, Y.; and Wang, L. 2017. Facial expression recognition based on deep evolutionary spatial-temporal networks. *IEEE Transactions on Image Processing*, 26(9): 4193–4203.
- Zhang, Y.; Chen, M.; Shen, J.; and Wang, C. 2022. Tailor versatile multi-modal learning for multi-label emotion recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 9100–9108.
- Zhao, G.; Huang, X.; Taini, M.; Li, S. Z.; and Pietikäinen, M. 2011. Facial expression recognition from near-infrared videos. *Image and vision computing*, 29(9): 607–619.