

# Friends-MMC: A Dataset for Multi-modal Multi-party Conversation Understanding

Yueqian Wang<sup>1</sup>, Xiaojun Meng<sup>2</sup>, Yuxuan Wang<sup>3</sup>, Jianxin Liang<sup>1</sup>, Qun Liu<sup>2</sup>, Dongyan Zhao<sup>1,4\*</sup>

<sup>1</sup>Wangxuan Institute of Computer Technology, Peking University

<sup>2</sup>Huawei Noah’s Ark Lab

<sup>3</sup>Beijing Institute for General Artificial Intelligence

<sup>4</sup>National Key Laboratory of General Artificial Intelligence

{wangyueqian,liangjx,zhaodongyan}@pku.edu.cn, {xiaojun.meng, qun.liu}@huawei.com, wangyuxuan1@bigai.ai

## Abstract

Multi-modal multi-party conversation (MMC) is a less studied yet important topic of research due to that it well fits real-world scenarios and thus potentially has more widely-used applications. Compared with the traditional multi-modal conversations, MMC requires stronger character-centered understanding abilities as there are many interlocutors appearing in both the visual and textual context. To facilitate the study of this problem, we present Friends-MMC in this paper, an MMC dataset that contains 24,000+ unique utterances paired with video context. To explore the character-centered understanding of the dialogue, we also annotate the speaker of each utterance, the names and bounding bboxes of faces that appear in the video. Based on this Friends-MMC dataset, we further study two fundamental MMC tasks: conversation speaker identification and conversation response prediction, both of which have the multi-party nature with the video or image as visual context. For conversation speaker identification, we demonstrate the inefficiencies of existing methods such as pre-trained models, and propose a simple yet effective baseline method that leverages an optimization solver to utilize the context of two modalities to achieve better performance. For conversation response prediction, we fine-tune generative dialogue models on Friend-MMC, and analyze the benefits of speaker information. The code and dataset will be publicly available, and thus we call for more attention on modelling speaker information when understanding conversations.

## Datasets —

<https://github.com/yellow-binary-tree/Friends-MMC>

## Introduction

Multi-modal dialogue systems have attracted extensive attention in recent studies (Zang et al. 2021; Zheng et al. 2022; Feng et al. 2023; Zhu et al. 2023; Liu et al. 2023; Li et al. 2023a; Maaz et al. 2023; Li et al. 2023b), especially with the rapid development of multi-modal large language models. However, there are two main deficiencies of existing works: (1) As most multi-modal datasets are collected from human annotations, LLM generated contents or chat history

from social media, these dialogues are mostly presented in question-answer format (AlAmri et al. 2019) between a human and a system, instead of among several human interlocutors; (2) the speakers are bystanders (Das et al. 2016) and discuss the given visual content such as an image or a video, instead of really being situated in the context.

In contrast, in real-world face-to-face conversations, the interlocutors are situated in the visual & audio context<sup>1</sup>, which means that the conversation and the visual context influence each other rather than interlocutors discussing about a given fixed visual context. Also, there can be more than 2 interlocutors involved, which means that modelling speaker information is crucial for comprehending the conversation. Understanding face-to-face conversations is an important foundation for achieving embodied artificial intelligence (Zhou et al. 2023; Xu et al. 2023; Zhang et al. 2023). Therefore, we emphasize that multi-modal multi-party conversation is a very important field for real industrial applications yet a less studied topic.

To overcome the shortcomings of existing work on multi-modal conversation, in this paper we extend its scopes and propose a new field of research: multi-modal multi-party conversation (MMC). Compared with traditional multi-modal conversations, MMC differs in the following aspects: (1) The conversation is not only between a user and an assistant, but among an arbitrary number of interlocutors, therefore the information of the speaker should be explicitly modeled; (2) Instead of chatting on images / videos as bystanders, the interlocutors are situated within the visual context, or in other words, the visual context provides rich information about the interlocutors, such as their identities, expressions and actions.

To foster the study of MMC, we build Friends-MMC, a multi-modal multi-party conversation dataset collected from the famous TV series *Friends*.<sup>2</sup> An overview of Friends-

<sup>1</sup>In this work, either images or videos can be used as the visual context. When videos are used, the audios are also included. To avoid cluttering, we keep using “visual context” to refer to videos and omit “& audio” in the rest of this paper.

<sup>2</sup>We have investigated the copyright issue of releasing our dataset, it can be released and used for non-commercial purposes. In fact, previous work (Poria et al. 2018) also created a dataset collected from TV Series *Friends*, and the authors have already re-

\*Corresponding Author

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

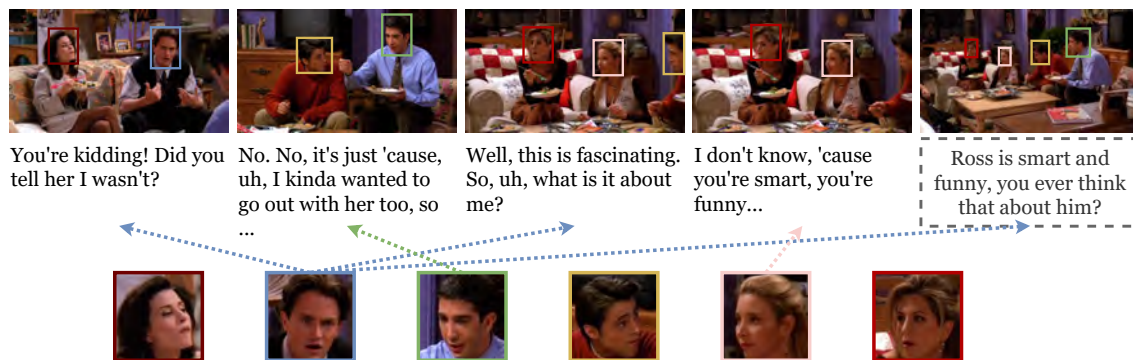


Figure 1: An example of multi-modal multi-party conversation. The task of conversation speaker identification is to infer the dotted arrows pointing from characters to utterances, and the task of conversation response prediction is to infer the last utterance in the dotted rectangular. Only one frame of the video is shown as the visual context to avoid clutter.

MMC is shown in Figure 1. A session in Friends-MMC consists of several turns, each paired with a video clip as visual context, in which the face tracks are detected and classified by character<sup>4</sup>. Compared to the existing multi-modal dialogue and multi-party conversation datasets, our proposed Friends-MMC have some traits worth emphasizing:

**a) Modalities** of available data are multiple, including: textual information of each utterance<sup>3</sup>, visual & audio context in forms of frames and videos, face bounding boxes and face character names. Utilizing all of these modalities can be challenging for existing models; **b) Conversations** are taken from daily life such as TV series, which are more natural and diverse compared to existing multi-party conversation datasets (Ouchi and Tsuboi 2016; Hu et al. 2019) that are mostly collected from chats on the topic of computers. **c) Reasoning** can be very complex. For example, in terms of conversation speaker identification, the speaker may not appear in its corresponding video or frame. Therefore, the preceding or succeeding textual and visual context, as well as their temporal relations, should be taken into account, which is quite difficult to solve for even humans in our experiments.

We focus on two fundamental MMC tasks: **Conversation Speaker Identification** and **Conversation Response Prediction**. The goal of conversation speaker identification is to link the speaker of every turn to the faces in the corresponding visual context. It requires a system to not only infer the speaker of each turn from the textual content, but also understand the visual context in which the dialogue happens. Existing works (Gu et al. 2021; Meng, Mou, and Jin 2018) of speaker identification mostly focus on text-only multi-party conversation, where only the speaker of the last one or few turns should be predicted given the speaker label of previous turns. Our paper introduces the multi-modal information of the interlocutors to make this task more challeng-

leaded raw videos in their homepage.

<sup>3</sup>In this paper we use the term “turn” to denote all contexts of a single turn of speaking, including utterance (*i.e.*, textual content), visual & audio context, face information and speaker information. We use the term “utterance” to denote the textual content only.

ing and better conform to real-world scenarios where all speaker labels are not available in the dialogue history. To accomplish this task, we present our task-specific baseline method including three modality-specific modules which is further introduced later in the “Conversation Speaker Identification” Section. With our task-specific designs, this system can achieve competitive accuracy on conversation speaker identification, far exceeding recent multi-modal pre-trained models, while takes a small amount of computation and enjoys a high degree of flexibility.

The other task is conversation response prediction, one of the most popular tasks for dialogue modelling. It requires a system to predict the last utterance with respect to the context. Compared with existing response generation task in multi-modal dialogues, the context of MMC is very heterogeneous in modality. We fine-tune text-only and multi-modal dialogue models on Friends-MMC with different sources of speaker information, to validate that the speaker information is critical to the response prediction of MMCs. In other words, the ability of correctly identifying the conversation speaker can benefit response prediction, which is an issue that has been ignored in existing multi-modal dialogue researches where two speakers simply take turns to speak to each other.

In summary, our contributions are three-fold: (1) We formally propose multi-modal multi-party conversation, a new and valuable field of research, and study two sub-tasks of it: conversation speaker identification and conversation response prediction; (2) We build and release Friends-MMC, a dataset to facilitate the study of multi-modal multi-party conversation; (3) We design a baseline for conversation speaker identification, validate its performance on Friends-MMC, and analyze the benefits of speaker information on conversation response prediction.

## Friends-MMC Dataset

In this section, we describe the dataset collection and annotation procedure for constructing the Friends-MMC dataset, which covers all the 220 episodes from 10 seasons of the TV show *Friends*. The reasons we use *Friends* are: (1) it is

a sitcom series, which has numerous conversations that contain diverse topics of daily life; (2) Though having as many as 220 episodes, it has a relatively small number of main characters, which is convenient for automatic face labelling and data cleaning; and (3) It's easy to get publicly available resources like high-quality subtitles that are often manually revised and paired perfectly with the video by a large group of TV fans, which greatly reduces manual labour during the data construction process as well as guarantees the data quality.

The content and speaker of each turn are extracted from transcripts and subtitles<sup>4</sup>. Faces and their character names in each frame are detected and labelled automatically for the train set (Season 1, 2, and from 4 to 10), and manually for the test set (Season 3) to ensure its accuracy.

## Construction Process

Figure 2 shows the overall construction process of the dataset. Now we introduce every step in details:

**Video Preprocessing.** We crop a clip from the video according to the start and end timestamp in the subtitle. We use an off-the-shelf face detector (Zhang et al. 2017) to detect faces for each frame in the clip. Following (Kalogeiton and Zisserman 2020), we merge the faces in adjacent frames into face tracks and thus remove the faces that are not in any track to clean out false positive faces.

**Face Prototype Construction.** C1C (Kalogeiton and Zisserman 2020) is a dataset with human-labelled face tracks for season 3 of *Friends*. We choose a set of 18 main characters, manually select 20 faces in different viewing angles per character, and encode them using Facenet-512 (Schroff, Kalenichenko, and Philbin 2015) to get face feature vectors in the embedding space of Facenet-512 as prototypes for each character.

**Automatic Face Labelling.** We automatically label the detected faces tracks with character names by finding their nearest neighbour in the mentioned embedding space. In particular, for each detected face track, we encode each frame in the track with Facenet-512 and calculate the cosine similarity between their feature and all prototypes. If the mean value of the largest 5 cosine similarities is greater than a threshold  $t = 0.6$  (which is set to maximize the validation accuracy described in the following paragraph), we label this face track with the corresponding character name, otherwise we think this face does not belong to any of the main characters and discard it.

To validate the accuracy of automatic face labelling, we use the same process to detect and label faces for season 3 and compare the results with human-annotated ones from C1C. The verification follows the rule that if the IoU of bounding boxes of an automatically labelled face and a human-annotated face is greater than 0.5, we identify them as a pair of identical faces. Given this threshold, 95% of all pairs of identical faces are labelled with correct names,

<sup>4</sup><https://my-subtitles.co/showlists/subtitles-610-friends>  
<https://fangj.github.io/friends>

which verifies the effectiveness of our automatic face labelling method.

**Test & Test-noisy Set.** For the test set, we directly use the human-annotated faces in C1C to guarantee the accuracy of face labelling, thus serving as high-quality ground-truths for this test set. Moreover, in order to align with the fact of imperfect face recognition in real-world scenarios and be consistent with the train set, we also create a more challenging test-noisy set by randomly removing 20% labelled face tracks.

**Image Frame Selection.** As many recent multi-modal pre-trained models accept images instead of videos as input, we also provide an alternative option of visual context where each turn is paired with only 1 frame. In such case, we select the frame with the most detected faces from each video clip, as well as its face bounding boxes and face names, as the paired visual context of this turn.

**Session Selection with Sliding Windows.** We use a sliding window of size  $m$  to select  $m$  adjacent turns as a dialogue session if the following conditions are met: (1) all speakers are from the main characters; (2) the time intervals between all adjacent turns are shorter than 8 seconds, which is a heuristic rule to prevent selecting turns from different scenes. Therefore, we use  $m = \{5, 8\}$  to create two types of sessions with different context lengths (5 turns and 8 turns). Note that different dialogue sessions may contain the same turn, as it belongs to different contexts and thus the preceding or succeeding textual and visual contents differ.

## Dataset Statistics

Dataset statistics are shown in Table 3. Apart from the basic statistics, we also count the proportion of speakers whose faces are not detected in the current clip or frame, or even not appear in all the  $m$  clips or frames within this entire session.

As table 3 shows, the test-noisy set includes a significantly larger number of speakers not in the current clips (or frames for image as input) than the test set. Therefore, this case is more challenging for speaker identification task, as the candidate model needs to really understand the conversation and find out more clues from the context rather than only the current clips or frames to infer who is the real speaker.

## Conversation Speaker Identification

### Task Introduction

Conversation speaker identification requires models to identify the speaker of each turn given the textual and visual context. Existing works on speaker identification mostly focuses on the text-only multi-party conversation. It often asks models to predict the speaker of the last few turns, given the dialogue history and the speaker of previous turns. However, in real-world scenarios, speaker labels are usually either available for all turns or not available for any turn. To simulate it, when we extend the task to multi-modality, we provide the basic visual information of interlocutors so that the speaker of each turn can be predicted.

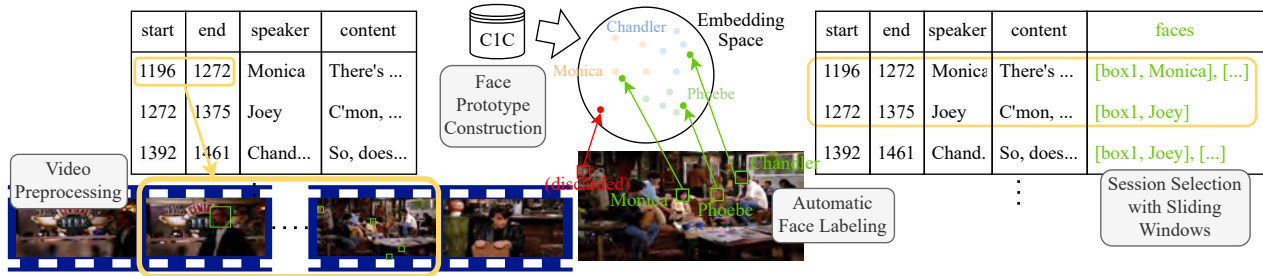


Figure 2: An overview of the construction process of Friends-MMC dataset.

	5 turns			8 turns		
	train	test	test-noisy	train	test	test-noisy
# sessions	13584	2017	2017	8730	1325	1325
# unique turns	21092	3069	3069	16990	2480	2480
# words in utterance	18.87	20.28	20.28	18.71	20.42	20.42
# speakers in each session	2.83	2.85	2.85	3.43	3.47	3.47
# face tracks per clip	2.41	3.12	2.50	2.39	3.14	2.52
avg. secs per face track	2.31	2.71	2.72	2.30	2.74	2.73
% speakers not in current clip	13.43	1.03	19.26	13.51	1.10	18.93
% speakers not in all clips	6.13	0.17	1.13	5.57	0.14	0.44
# faces per frame	1.61	2.20	1.76	1.60	2.21	1.78
% speakers not in current frame	24.05	6.52	25.64	24.15	6.42	25.30
% speakers not in all frames	9.53	1.01	3.32	7.45	0.42	1.37

Table 1: Dataset Statistics of Friends-MMC. We provide a train set, a test set and a more challenging test-noisy set.

However, too many decisive clues for identifying the speaker, such as voice characteristics and facial movements, can also be provided by the video and audio. This will cause models to tend to only rely on the rich video information and ignore the dialogue context, which deviates from our original motivation of promoting researches on conversation systems. Therefore, we also propose an alternative setting: providing only one frame and no audio as the visual context. In this setting, the basic information of interlocutors such as identities, expressions, and the scene they are in, can also be provided by this frame, and it leaks less decisive clues and shortcuts for directly identifying the speaker.

### Baseline Method

To effectively utilize all the modalities of the proposed benchmark dataset, including visual, audio, textual and face tracks, we propose a baseline method, which consists of three modules: 1) a visual model  $M_1$  to recognize speaking faces, 2) a text model  $M_2$  to analyse multi-speaker relations based on dialogue contexts, and 3) a quadratic binary optimization problem solver to combine their results and thus identify the speaker of each turn. This modular design makes our system enjoy a high degree of flexibility, as one can easily change the visual model or the text model with alternative ones when different contextual information (e.g., using image or video as visual input) is provided. Figure 3 shows the overview of our proposed method, and we introduce each module in the following paragraphs.

**Visual Model** We use a visual model to predict the probability of each face belonging to the current speaker individually:  $p_{face} = M_1(face) \in (0, 1)$ . There are two dif-

ferent visual contexts: 1) If the provided visual context is an image frame, we use an inception model (Szegedy et al. 2014) pre-trained on VGGFace2 (Cao et al. 2017) and fine-tuned on the train set of Friends-MMC as the visual model  $M_1$ , and  $face$  is an image of the face cropped using bounding box  $b$ . When fine-tuning on Friends-MMC, the speaking label of a face  $y_{face}$  is set to 1 if the character name  $c$  of this face is identical to the speaker name  $y$ , and 0 otherwise:  $y_{face} = \mathbf{1}[c = y]$ . We use the cross-entropy classification loss as the training objective. 2) If the provided visual context is a video, we use TalkNet (Tao et al. 2021), a state-of-the-art active speaker detection model, as the visual model  $M_1$ , and  $face$  is a video of the face cropped using bounding box sequence  $b$ , accompanied by the audio of the same period. Note that using TalkNet in a zero-shot manner is good enough, we try to fine-tune it on the train set of Friends-MMC as like the CNN model, but it does not result in better performance.

**Text Model** We use a DeBERTa-v3 (He, Gao, and Chen 2021) first trained on Ubuntu Dialogue Corpus (Hu et al. 2019) and then fine-tuned on the train set of Friends-MMC as  $M_2$  to predict whether every two utterances in a dialogue session are spoken by the same speaker. The intuitive reason behind it is that for some utterances, it is hard to identify the speaker solely by feeding the video or frame into  $M_1$ . We thus try to conjecture its speaker by finding whether it likely shares the same speaker with another utterance, for which we have confidences or prior knowledge to infer its speaker. In addition,  $M_2$  helps to achieve a relatively global optimum speaker inference by considering all turns together.

Given a dialogue session consists of  $m$  utterances, we

prepend an  $\langle \text{bos} \rangle$  token to each utterance as the input of  $M_2$  as like:  $\langle \text{bos} \rangle u_1 \dots \langle \text{bos} \rangle u_m$ . We use the last hidden state of each  $\langle \text{bos} \rangle h_i$  as the representation of each utterance, and use a head layer to calculate the similarity of every two representations:

$$p_{sim}^{ij} = \sigma(W_2 \text{GeLU}(W_1[h_i; h_j; |h_i - h_j|] + b_1) + b_2)$$

where  $i, j = 1, \dots, m$ , and  $(W_1, b_1, W_2, b_2)$  are learnable parameters.  $\sigma$  is the sigmoid activation function, and  $p_{sim}^{ij} \in (0, 1)$  is a scalar that denotes the probability of two utterances spoken by the same person. The loss function is defined as:

$$\mathcal{L}_{M_2} = \text{MSE}(p_{sim}, y_{sim}) + \text{MSE}(p_{sim}, p_{sim}^T)$$

where  $y_{sim} \in \{0, 1\}^{m \times m}$  is the ground truth label of whether any two utterances are from the same speaker, and  $\text{MSE}$  denotes mean squared error loss.

**Quadratic Binary Optimization Problem Solver** This module aims to integrate the outputs of both visual ( $M_1$ ) and textual ( $M_2$ ) models for speaker identification. In particular, we design two matrices using the outputs of these two models, and address it as the quadratic binary optimization.

For each dialogue session, we first obtain a candidate speaker set by recording all faces appeared in every frame / video:  $\mathbf{C} = \{c_1, \dots, c_l\}$ . We construct a vision reward matrix  $\mathbf{B} \in R^{l \times m}$  of selecting a character  $c_i$  as the speaker of the turn  $u_j$ . If the face of  $c_i$  appears in the frame / video  $v_j$ ,  $b_{ij}$  is set to the probability of that face as a speaking face predicted by  $M_1$ , otherwise  $b_{ij} = 0$ . Apparently,  $\mathbf{B}$  can only express those situations that the speaker appears in the corresponding frame / video. To address those problems, we construct another text reward matrix  $\mathbf{A} \in R^{m \times m}$  to reveal the reward for assigning the same speaker to two turns  $u_i$  and  $u_j$ . We first use  $M_2$  described in the previous subsection and get a similarity matrix  $p_{sim}$ . However, if we simply use this similarity matrix  $p_{sim}$  as the reward matrix  $\mathbf{A}$ , since all elements in  $p_{sim}$  are larger than 0, the optimization solver tends to assign the same speaker to every turn to get the maximum rewards. To avoid which, we subtract the similarity matrix with the mean value of its elements as the reward matrix, i.e.,  $\mathbf{A} = p_{sim} - \text{mean}(p_{sim})$

As Figure 3 shows, with  $\mathbf{A}$  and  $\mathbf{B}$  in hand, the task of multi-modal multi-party speaker identification can be represented by a quadratic binary optimization problem:

$$\begin{aligned} \text{Maximize} \quad & f(X) = (1 - \alpha)X^T A X + \alpha X B \\ \text{s.t.} \quad & X \in \{0, 1\}^{m \times l}, \\ & \sum_{j=1}^l X_{ij} = 1, \quad i = 1, 2, \dots, m \end{aligned}$$

where  $\alpha$  is a hyperparameter to control the weight of two rewards and is selected according to the performance on a validation set held-out from the train set. We use 0.8 for frame as visual context, 0.7 for video as visual context, and 0.2 when ground truth labels of the text model are provided

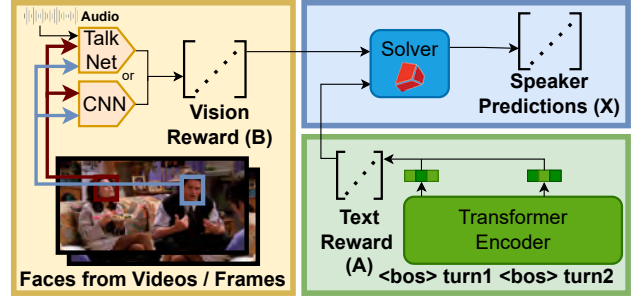


Figure 3: Model overview of the three modules in different colors: the visual ( $M_1$ ) is yellow, the textual ( $M_2$ ) is green, and the optimization solver taking vision and text reward matrix as input is blue.

( $M_2^\dagger$ ). By now, this problem can be easily solved using optimization problem solvers like (Gurobi Optimization, LLC 2023), which adaptively makes decisions based on the output of  $M_1$  and  $M_2$ . The reason we use an optimization solver instead of an end-to-end pre-trained model is that this task of identifying speakers still remains challenging to use the general attention mechanism of pre-trained models to fuse different modalities.

## Experiment Results

**Implementations** We conduct experiments in five settings with information of different modalities used: (1) image (frame) only; (2) video only; (3) text only; (4) image and text; and (5) video and text. In image only setting (1), we use the model  $M_1$  CNN to predict one face from all detected faces from the frame as the speaker. If there are no faces, we randomly choose a character from the candidate speaker list. The video only (2) is almost the same as (1) but uses  $M_1$  TalkNet to predict one face from the video. In the text only setting (3), we compare  $M_2$  (DeBERTa-v3) with 3-shot GPT 3.5 using its in-context learning ability to perform this task. For  $M_2$ , although it is good at judging whether two utterances are said by the same speaker, it is not trained to identify the speaker for a single utterance. Therefore, it can only make guesses according to the relations between sentences. GPT 3.5, however, possesses some ability to understand the candidate speaker list and identifying names from utterances, so it can make full use of the textual context to achieve more accurate reasoning. In image-text setting (4), the jointed  $M_1$  (CNN) +  $M_2$  model is used together with a quadratic binary optimization solver, and we also try to replace the output of  $M_2$  with ground truth labels (denoted by  $M_2^\dagger$ ) to explore bottlenecks and possible improvement directions. We also test GPT4-o (OpenAI 2024) under this setting. Video-text setting (5) is almost the same as (4), but again CNN is replaced by TalkNet for  $M_1$ .

We also fine-tune three popular and powerful multi-modal pre-trained models: Violet (Fu et al. 2021), LLaVA (Liu et al. 2023) and Emu (Sun et al. 2023). As LLaVA only accepts one image at a time during pre-training, we also use one single turn (one utterance and one frame) during

		5 turns noisy		8 turns noisy	
0	random (std.dev.)	31.82 (0.25)	32.61 (0.47)	28.54 (0.49)	29.03 (0.27)
<i>Frame Only</i>					
1	$M_1$ (CNN)	72.88	63.72	72.90	62.51
<i>Video Only</i>					
2	$M_1$ (TalkNet)	80.89	70.91	81.00	70.50
<i>Text Only</i>					
3	$M_2$	33.24	33.85	29.09	29.33
4	GPT 3.5 (3-shot)	37.21	37.24	33.35	32.81
<i>Use image and text modality</i>					
5	Violet	32.66	33.16	27.73	28.86
6	LLaVA v1.5-13B	46.30	42.39	45.73	41.41
7	Emu-14B	61.76	58.23	60.96	56.46
8	$M_1$ (CNN) + $M_2$	75.81	68.61	74.53	67.21
9	$M_1$ (CNN) + $M_2^\dagger$	84.90	78.01	90.80	83.93
10	GPT-4o (0-shot)	66.36	65.60	63.64	61.02
11	Human	82.25	-	84.49	-
<i>Use video and text modality</i>					
12	$M_1$ (TalkNet) + $M_2^\dagger$	83.21	74.12	83.60	75.00
13	$M_1$ (TalkNet) + $M_2^\dagger$	90.88	83.09	95.10	89.69

Table 2: Accuracy on the test and test-noisy set of Friends-MMC.  $M_1$  and  $M_2$  denote the visual and textual model of our baseline method, respectively. For  $M_1$ , we use CNN or TalkNet to take image or video as input.  $\dagger$  indicates that we use ground truths instead of textual model outputs ( $M_2$ ) to serve as upper bounds.

fine-tuning and testing, neglecting information from adjacent turns. Violet and Emu, on the other hand, are pre-trained with videos or image-text interleaved data, so we can directly use all 5 or 8 turns as input.

To check how much useful information can be provided by using frame as the visual context, we also report the human performance of this task. we randomly sample 80 dialogue sessions from each (5 turns / 8 turns) test set, provide dialogue contents, frames, face bounding boxes & names to participants, and ask them to select a speaker for each turn from the candidate speaker set (*i.e.*, the characters that appear in all frames). This process requires intensive efforts from humans, according to their post-interview, as the task of selecting speakers requires careful observation and a thorough understanding of the dialogue contents. We thus only perform the human studies on the test set, since we believe the human performance on the test-noisy set should be apparently worse. See appendix for details of all the baselines mentioned above.

**Main Results** Results can be found in Table 2: (1) visual context acquired by the vision model  $M_1$ , including which face appears in the frame and looks like a speaking face, serves as the most critical clues, shown by the performance of  $M_1$  (line 1, 2). We can conclude that this speaker identification task is still vision dominant, especially when videos are provided. (2) Speaker relations acquired by the text model  $M_2$  also play a vital supporting role to bring a significant improvement of 2% ~ 5% from  $M_1$  to  $M_1 + M_2$  (line 8 vs. line 1, line 12 vs. line 2). The textual context benefits this task not only by providing dialogue contents, but also aims for more real scenarios where the speaker does not

appear in the scene. What’s more, the benefit is greater when the paired visual model is more accurate. It makes sense that one has to accurately identify speakers of some turns before it is able to identify other turns using this speaker relation information. (3) Directly fine-tuning a multi-modal pre-trained model (line 5, 6, 7) or using proprietary multi-modal models (line 10) is not a good option despite we have tried different models with various parameter sizes. We believe the key aspects that are essential to solve this problem remain difficult to be understood by these candidate models, as this task differs a lot from the original training objectives. It also may be due to the reason that this speaker identification task is difficult to be formatted as the proper and shorten input of the model and thus to be easily learned. (4) Comparing the line 8-9 and 12-13, our current method of fine-tuning  $M_2$  model to predict whether two utterances are spoken by the same speaker still has a lot of room for improvement. This strong upper bound result also evidences that our designed objective for the textual module is meaningful.

The influence of the hyper-parameter  $\alpha$  to the results is listed in the appendix.

## Conversation Response Prediction

### Task Introduction

Conversation response prediction aims to predict the textual content of the last turn given the visual context of all turns and the textual content of all previous turns. In this study, we demonstrate that speaker information is critical to the response prediction task of multi-modal multi-party conversations. Though predicting the next utterance in multi-modal dialogues has been a widely studied topic, none of the previous work focuses on the specific nature of multi-party dialogues, *i.e.*, taking the speaker information into account.

We hypothesize that the potential advantages come from speaker information in two aspects: (1) The model may learn the speaking style of each speaker. When the speaker information like name is available, it can be used as a **global speaking style indicator** for generating better responses; (2) Speaker information can be a **local context indicator** to infer the response is from which person’s point of view, *i.e.*, should be consistent with which utterance in the local conversation.

To validate these two aspects, we constructed 3 test sets: (a) a test set with random speaker names: speakers of all utterances are randomly assigned. Here the potential advantages of (1) and (2) are both broken. (b) a test set with random history speaker names: similar to (a) but we keep the speaker name of the utterance we need to predict unchanged. Here only the potential advantage of (2) is broken. (c) a test set with shuffled names: we replace all names in both speaker information and utterance content with another name according to a predefined random shuffle mapping, *e.g.*, replace “Ross” with “Joey”, and “Joey” with “Chandler”. Here only the potential advantage of (1) is broken. Note that in all cases the speaker information of the train set is not modified.

Our experiments are in the following settings: train and test without speaker name; train with ground truth speaker

Model	Speaker	5 turns	8 turns
Llama2-7B	No	30.69	36.98
	Random	31.23	43.32
	Random History	31.63	43.40
	Shuffled	35.20	48.60
	Ground truth	36.89	49.36
	$M_1(\text{CNN}) + M_2$	34.16	45.81
	$M_1(\text{TalkNet}) + M_2$	34.56	46.64
Emu-14B	No	30.49	31.09
	Random	29.35	31.55
	Random History	29.45	31.25
	Shuffled	33.02	35.17
	Ground truth	34.06	36.30
	$M_1(\text{CNN}) + M_2$	31.98	33.89
	$M_1(\text{TalkNet}) + M_2$	32.97	34.64

Table 3: Accuracy of conversation response prediction by selecting one from a set of ten utterances as candidates.

information and test with ground truth or random or shuffled or our identified speaker name ( $M_1 + M_2$ ). Note that when using our  $M_2$  model, the last utterance to predict is removed from the input, *i.e.*, the speaker of top  $m - 1$  turns is predicted using both visual and textual context, and the speaker of the last turn is predicted using visual context. In addition, we finetune Emu-14B as a dialogue model with visual input, and Llama2-7B as a dialogue model without visual input. As mentioned in (Wang, Wang, and Zhao 2023), it is often hard to predict the next utterance solely by the dialogue history, hence the text generation metrics (*e.g.*, rouge) may not be very valid when evaluating on our movie dialogue dataset. Therefore, following (Wang, Wang, and Zhao 2023), we build a response selection task to evaluate models instead of using generative metrics. We randomly select 9 other utterances from the test set as negative responses, and ask the model to select the one with the lowest perplexity from totally 10 (plus ground truth) candidates as output.

## Experiment Results

As shown in Table 3, regardless of which pre-trained model is used, or the length of dialogue context, in all cases adding speaker information always improves the performance. Comparing to using the random speakers and random history speakers, the performance with shuffled speakers is much closer to the one with ground truth speakers. Therefore, in terms of our hypothesis on the two potential advantages that come from speaker information, this result indicates that speaker information mainly works by indicating the local context instead of serving as a global speaking style indicator. Overall, our baseline method of speaker identification also benefits the response prediction task as it achieves a significant better performance than the ones with no or random speakers.

## Related Works

### Multi-party Conversations

Multi-party conversations (MPC), as opposed to two-party conversations, is a more practical and challenging scenario which studies conversations that involve more than two interlocutors. Research on MPC often includes three sub-

topics: speaker prediction, utterance prediction, and addressee prediction. (Ouchi and Tsuboi 2016; Hu et al. 2019) built MPC datasets from Ubuntu technical dialogues and proposed baseline models for MPC tasks. Recent studies on MPC usually train and evaluate models jointly on those three objectives. (Gu et al. 2021) propose MPC-BERT, which finetunes BERT (Devlin et al. 2019) on several self-supervised tasks, and achieve state-of-the-art results on the above MPC tasks. (Su and Zhou 2022) identify speakers of utterances by clustering them with pairwise relations encoded by a dialogue content encoder. GIFT (Gu et al. 2023) revises the model structure of transformer encoders to make the self-attention layer be aware of the information flow of MPC. Details regarding MPC can be found in these recent surveys (Gu, Tao, and Ling 2022; Ganesh, Palmer, and Kann 2023).

### Multi-Modal Dialogue Datasets

There have been a number of works on building multi-modal dialogue datasets (Das et al. 2016; Mostafazadeh et al. 2017; Shuster et al. 2020; AlAmri et al. 2019; Zheng et al. 2022; Zang et al. 2021; Feng et al. 2023). Despite the diversity in modality (image or video) and the position of the visual information in the dialogue history, all above datasets are limited as the interlocutors are outside the visual contexts rather than “situated” inside them, and only includes two interlocutors instead of being “multi-party”.

Dialogue in movie/TV series is a typical data source of multi-party conversation with “situated” visual context. Recent large-scale movie dialogue datasets include OpenViDial (Meng et al. 2020; Wang et al. 2021) and VSTAR (Wang et al. 2023). However, these datasets do not have any kind of speaker information, which hinders a deeper-level understanding of the conversation. Perhaps the dataset most similar to ours is MELD (Porcia et al. 2018), which is also a speaker-aware multi-modal multi-party dialogue dataset collected from *Friends* but focuses only on emotion recognition, and does not annotate faces in the visual context.

## Conclusion

We work on a new field of research, multi-modal multi-party conversation (MMC), to bridge the gap between existing researches on multi-modal dialogue and real-world applications such as face-to-face conversations or meetings. We build Friends-MMC, an MMC dataset in which each utterance is paired with video context, speaker, face bounding box and face name annotation. We propose two new tasks of MMC, namely conversation speaker identification and conversation response prediction, and design a baseline to solve both tasks. Our future directions include collecting more diverse data rather than movies, and exploring ways to make better use of speaker information for response generation.

## Acknowledgments

The authors thank the kind suggestions and support from AI Data Technology Laboratory of Huawei Noah’s Ark Lab.

## References

- AlAmri, H.; Cartillier, V.; Das, A.; Wang, J.; Lee, S.; Anderson, P.; Essa, I.; Parikh, D.; Batra, D.; Cherian, A.; Marks, T. K.; and Hori, C. 2019. Audio Visual Scene-Aware Dialog. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7550–7559.
- Cao, Q.; Shen, L.; Xie, W.; Parkhi, O. M.; and Zisserman, A. 2017. VGGFace2: A Dataset for Recognising Faces across Pose and Age. *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, 67–74.
- Das, A.; Kottur, S.; Gupta, K.; Singh, A.; Yadav, D.; Moura, J. M. F.; Parikh, D.; and Batra, D. 2016. Visual Dialog. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1080–1089.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *North American Chapter of the Association for Computational Linguistics*.
- Feng, J.; Sun, Q.; Xu, C.; Zhao, P.; Yang, Y.; Tao, C.; Zhao, D.; and Lin, Q. 2023. MMDialog: A Large-scale Multi-turn Dialogue Dataset Towards Multi-modal Open-domain Conversation. In *Annual Meeting of the Association for Computational Linguistics*.
- Fu, T.-J.; Li, L.; Gan, Z.; Lin, K.; Wang, W. Y.; Wang, L.; and Liu, Z. 2021. VIOLET : End-to-End Video-Language Transformers with Masked Visual-token Modeling. *ArXiv*, abs/2111.12681.
- Ganesh, A.; Palmer, M.; and Kann, K. 2023. A Survey of Challenges and Methods in the Computational Modeling of Multi-Party Dialog. *Proceedings of the 5th Workshop on NLP for Conversational AI (NLP4ConvAI 2023)*.
- Gu, J.-C.; Ling, Z.; LIU, Q.; Liu, C.; and Hu, G. 2023. GIFT: Graph-Induced Fine-Tuning for Multi-Party Conversation Understanding. In *Annual Meeting of the Association for Computational Linguistics*.
- Gu, J.-C.; Tao, C.; and Ling, Z. 2022. Who Says What to Whom: A Survey of Multi-Party Conversations. In *International Joint Conference on Artificial Intelligence*.
- Gu, J.-C.; Tao, C.; Ling, Z.; Xu, C.; Geng, X.; and Jiang, D. 2021. MPC-BERT: A Pre-Trained Language Model for Multi-Party Conversation Understanding. In *Annual Meeting of the Association for Computational Linguistics*.
- Gurobi Optimization, LLC. 2023. Gurobi Optimizer Reference Manual.
- He, P.; Gao, J.; and Chen, W. 2021. DeBERTaV3: Improving DeBERTa using ELECTRA-Style Pre-Training with Gradient-Disentangled Embedding Sharing. *arXiv:2111.09543*.
- Hu, W.; Chan, Z.; Liu, B.; Zhao, D.; Ma, J.; and Yan, R. 2019. GSN: A Graph-Structured Network for Multi-Party Dialogues. In *International Joint Conference on Artificial Intelligence*.
- Kalogeiton, V. S.; and Zisserman, A. 2020. Constrained Video Face Clustering using INN Relations. In *British Machine Vision Conference*.
- Li, K.; He, Y.; Wang, Y.; Li, Y.; Wang, W.; Luo, P.; Wang, Y.; Wang, L.; and Qiao, Y. 2023a. VideoChat: Chat-Centric Video Understanding. *ArXiv*, abs/2305.06355.
- Li, K.; Wang, Y.; He, Y.; Li, Y.; Wang, Y.; Liu, Y.; Wang, Z.; Xu, J.; Chen, G.; Luo, P.; Wang, L.; and Qiao, Y. 2023b. MVBench: A Comprehensive Multi-modal Video Understanding Benchmark. *ArXiv*, abs/2311.17005.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual Instruction Tuning. *ArXiv*, abs/2304.08485.
- Maaz, M.; Rasheed, H. A.; Khan, S. H.; and Khan, F. S. 2023. Video-ChatGPT: Towards Detailed Video Understanding via Large Vision and Language Models. *ArXiv*, abs/2306.05424.
- Meng, Y.; Wang, S.; Han, Q.; Sun, X.; Wu, F.; Yan, R.; and Li, J. 2020. OpenViDial: A Large-Scale, Open-Domain Dialogue Dataset with Visual Contexts. *ArXiv*, abs/2012.15015.
- Meng, Z.; Mou, L.; and Jin, Z. 2018. Towards Neural Speaker Modeling in Multi-Party Conversation: The Task, Dataset, and Models. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Mostafazadeh, N.; Brockett, C.; Dolan, W. B.; Galley, M.; Gao, J.; Spithourakis, G. P.; and Vanderwende, L. 2017. Image-Grounded Conversations: Multimodal Context for Natural Question and Response Generation. *ArXiv*, abs/1701.08251.
- OpenAI. 2024. Hello GPT-4o. <https://openai.com/index/hello-gpt-4o/>. Accessed: 2024-12-10.
- Ouchi, H.; and Tsuboi, Y. 2016. Addressee and Response Selection for Multi-Party Conversation. In *Conference on Empirical Methods in Natural Language Processing*.
- 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*, abs/1810.02508.
- Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. FaceNet: A unified embedding for face recognition and clustering. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 815–823.
- Shuster, K.; Humeau, S.; Bordes, A.; and Weston, J. 2020. Image-Chat: Engaging Grounded Conversations. In *Annual Meeting of the Association for Computational Linguistics*.
- Su, Z.; and Zhou, Q. 2022. Speaker Clustering in Textual Dialogue with Pairwise Utterance Relation and Cross-corpus Dialogue Act Supervision. In *Proceedings of the 29th International Conference on Computational Linguistics, COLING 2022, Gyeongju, Republic of Korea, October 12-17, 2022*, 734–744. International Committee on Computational Linguistics.
- Sun, Q.; Yu, Q.; Cui, Y.; Zhang, F.; Zhang, X.; Wang, Y.; Gao, H.; Liu, J.; Huang, T.; and Wang, X. 2023. Generative Pretraining in Multimodality. *CoRR*, abs/2307.05222.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S. E.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A.

2014. Going deeper with convolutions. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1–9.
- Tao, R.; Pan, Z.; Das, R. K.; Qian, X.; Shou, M. Z.; and Li, H. 2021. Is Someone Speaking?: Exploring Long-term Temporal Features for Audio-visual Active Speaker Detection. *Proceedings of the 29th ACM International Conference on Multimedia*.
- Wang, S.; Meng, Y.; Li, X.; Sun, X.; Ouyang, R.; and Li, J. 2021. OpenViDial 2.0: A Larger-Scale, Open-Domain Dialogue Generation Dataset with Visual Contexts. *ArXiv*, abs/2109.12761.
- Wang, Y.; Wang, Y.; and Zhao, D. 2023. Overview of the NLPCC 2023 Shared Task 10: Learn to Watch TV: Multi-modal Dialogue Understanding and Response Generation. In *12th National CCF Conference, NLPCC 2023, Foshan, China, October 12-15, 2023, Proceedings, Part III*, volume 14304, 412–419.
- Wang, Y.; Zheng, Z.; Zhao, X.; Li, J.; Wang, Y.; and Zhao, D. 2023. VSTAR: A Video-grounded Dialogue Dataset for Situated Semantic Understanding with Scene and Topic Transitions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, 5036–5048. Association for Computational Linguistics.
- Xu, Y.; Wang, S.; Li, P.; Luo, F.; Wang, X.; Liu, W.; and Liu, Y. 2023. Exploring Large Language Models for Communication Games: An Empirical Study on Werewolf. *ArXiv*, abs/2309.04658.
- Zang, X.; Liu, L.; Wang, M.; Song, Y.; Zhang, H.; and Chen, J. 2021. PhotoChat: A Human-Human Dialogue Dataset With Photo Sharing Behavior For Joint Image-Text Modeling. *ArXiv*, abs/2108.01453.
- Zhang, H.; Du, W.; Shan, J.; Zhou, Q.; Du, Y.; Tenenbaum, J. B.; Shu, T.; and Gan, C. 2023. Building Cooperative Embodied Agents Modularly with Large Language Models. *CoRR*, abs/2307.02485.
- Zhang, S.; Zhu, X.; Lei, Z.; Shi, H.; Wang, X.; and Li, S. 2017. S3FD: Single Shot Scale-Invariant Face Detector. *2017 IEEE International Conference on Computer Vision (ICCV)*, 192–201.
- Zheng, Y.; Chen, G.; Liu, X.; and Lin, K. W. 2022. MM-Chat: Multi-Modal Chat Dataset on Social Media. *ArXiv*, abs/2108.07154.
- Zhou, P.; Madaan, A.; Potharaju, S. P.; Gupta, A.; McKee, K. R.; Holtzman, A.; Pujara, J.; Ren, X.; Mishra, S.; Nematzadeh, A.; Upadhyay, S.; and Faruqui, M. 2023. How FaR Are Large Language Models From Agents with Theory-of-Mind? *ArXiv*, abs/2310.03051.
- Zhu, D.; Chen, J.; Shen, X.; Li, X.; and Elhoseiny, M. 2023. MiniGPT-4: Enhancing Vision-Language Understanding with Advanced Large Language Models. *ArXiv*, abs/2304.10592.