

TikTalk: A Video-Based Dialogue Dataset for Multi-Modal Chitchat in Real World

Hongpeng Lin*
hopelin@ruc.edu.cn
Renmin University of China

Ludan Ruan*
ruanld@ruc.edu.cn
Renmin University of China

Wenke Xia*
xwk529086017@gmail.com
Renmin University of China

Peiyu Liu
liupeiyustu@ruc.edu.cn
Renmin University of China

Jingyuan Wen
wenjingyuan@ruc.edu.cn
Renmin University of China

Yixin Xu
xu_yixin@ruc.edu.cn
Renmin University of China

Di Hu
dihu@ruc.edu.cn
Renmin University of China

Ruihua Song[†]
songruihua_bloon@outlook.com
Renmin University of China

Wayne Xin Zhao
batmanfly@gmail.com
Renmin University of China

Qin Jin
qjin@ruc.edu.cn
Renmin University of China

Zhiwu Lu
luzhiwu@ruc.edu.cn
Renmin University of China

ABSTRACT

To facilitate the research on intelligent and human-like chatbots with multi-modal context, we introduce a new video-based multi-modal dialogue dataset, called TikTalk. We collect 38K videos from a popular video-sharing platform, along with 367K conversations posted by users beneath them. Users engage in spontaneous conversations based on their multi-modal experiences from watching videos, which helps recreate real-world chitchat context. Compared to previous multi-modal dialogue datasets, the richer context types in TikTalk lead to more diverse conversations, but also increase the difficulty in capturing human interests from intricate multi-modal information to generate personalized responses. Moreover, external knowledge is more frequently evoked in our dataset. These facts reveal new challenges for multi-modal dialogue models. We quantitatively demonstrate the characteristics of TikTalk, propose a video-based multi-modal chitchat task, and evaluate several dialogue baselines. Experimental results indicate that the models incorporating large language models (LLM) can generate more diverse responses, while the model utilizing knowledge graphs to introduce external knowledge performs the best overall. Furthermore, no existing model can solve all the above challenges well. There is still a large room for future improvements, even for LLM with visual extensions. Our dataset is available at <https://ruc-aimind.github.io/projects/TikTalk/>.

KEYWORDS

Multi-modal Dialogue; Chitchat; Real world; Dataset

1 INTRODUCTION

The ability to perceive and interact with multi-modal information, such as vision and audio, is a significant manifestation of conversational robots' progress toward artificial general intelligence. Recently, with the excellent performance of conversation models [1,

*Equal contribution.

[†]Corresponding author.

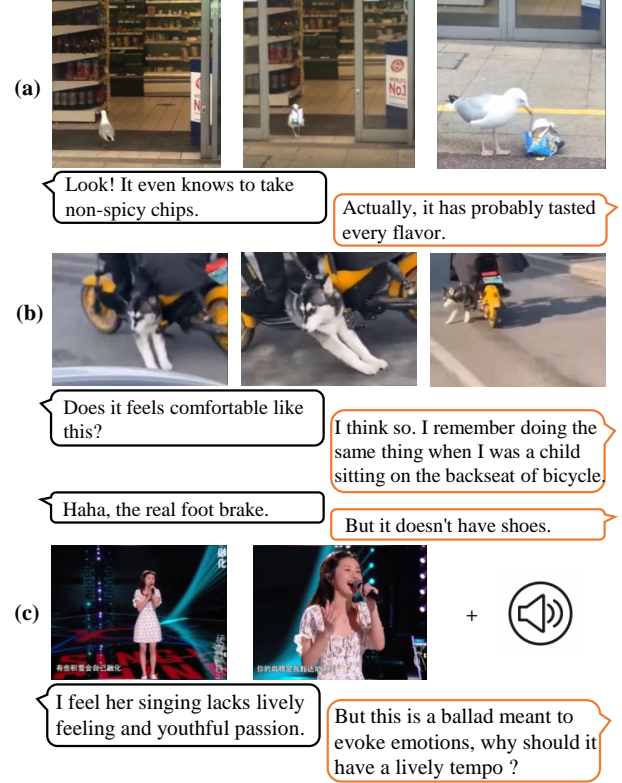


Figure 1: Dialogue examples from our TikTalk dataset. (a) A seagull carrying food from a store. (b) A dog intentionally stretches out its paw to rub against the ground while its owner rides an electric bike. (c) A singer is singing a ballad.

4, 36, 57, 60] and large language models [5, 17, 30, 42, 47, 56], there have been widespread interests in introducing multi-modal information into conversations. GPT-4 [31] and some other works [3, 22] have extended text-based dialogue agents to support both textual

and visual inputs. They provide surprisingly excellent results when prompted with questions related to visual content. However, current ways to ask questions about visual content [2, 6] are just designed to test the model’s abilities in objective perception and simple comprehension of visual information, and they still differ largely from human daily dialogues with multi-modal context.

To facilitate the research on multi-modal dialogue, various forms of datasets have been proposed. IGC [29], Image-Chat [38] and PhotoChat [53] construct dialogues around given images via crowdsourcing. MMDD [18] and DialogCC [19] match textual dialogue datasets with image caption datasets. MMChat [59] and MMDialog [13] are derived from image posts and comment-reply pairs on social media. OpenViDial [28, 45] and M³ED [58] extract dialogues of characters and corresponding scenes from movies and TV series.

However, there still exist certain gaps between these previous datasets and real-world multi-modal dialogues. Firstly, people engage in spontaneous dialogues under their shared experiences of multi-modal information from the real world, whereas most multi-modal dialogue datasets do not simulate this scenario. For example, the dialogues collected via crowdsourcing follow explicit guidelines and lack spontaneity and authenticity, leading to limited quality and coverage. Textual dialogue and image caption datasets for multi-modal dialogues are not naturally aligned. Additionally, character dialogues in movies and TV series are designed to meet the demands of storyline, and may be manipulated for dramatic effects. Secondly, real-world multi-modal dialogue scenarios involve various context, enabling individuals to engage in personalized and diverse dialogues. Most of the existing datasets provide limited visual information from static images. Furthermore, the specific environmental context in movies and TV series, as well as the predominant visual content of the dialogue scenes being the talking characters, may also restrict the diversity of conversations.

In this paper, we aim to introduce a novel video-based dialogue dataset that aligns with multi-modal chitchat scenarios in real world. Firstly, we construct a Chinese multi-modal dialogue dataset called TikTalk by collecting videos and their corresponding comments and replies from a video-sharing platform. TikTalk consists of 38K videos and 367K dialogues. Figure 1 shows several examples from TikTalk, where users engage in spontaneous dialogues after watching videos, offering a better simulation of real-world chitchat context. These dialogues involve a diverse range of multi-modal context. For example, the responses in Figure 1 (a) and (c) are derived from visual and audio content, respectively. In Figure 1 (b), the dialogue requires external knowledge, including personal childhood experiences and the association between the concept of “brake” and the action of rubbing against the ground. Moreover, users may capture different interests from the same context, as depicted in Figure 1 (b), where one user cares for whether the dog is comfortable, while the other notices that it does not wear shoes while rubbing against the ground. Secondly, we quantitatively compare our dataset with previous ones from the new perspective of which type of context is necessary when generating responses. Results show that the responses in TikTalk require a more diverse range of context. Additionally, we define a multi-modal chitchat task based on TikTalk and summarize three new challenges: (1) Perceiving more diverse multi-modal context and comprehending their complex interaction; (2) Capturing human interests from rich information to generate the

desired responses; (3) Introducing external knowledge related to the video as supplementary context. Thirdly, we evaluate some dialogue baselines on TikTalk, including two mainstream architectures: (1) Using adapter to bridge the semantic gap between vision encoder and language language model; (2) Directly fusing multi-modal features. We find that introducing visual information and external knowledge can improve the quality of responses. Regarding the use of knowledge to generate responses, large language models performs worse than knowledge graphs. However, existing baselines still have some issues, such as failing to utilize audio information effectively and a lack of exploration of the second challenge.

There are three main contributions of our paper:

- (1) We construct a new video-based dialogue dataset called TikTalk, simulating real-world multi-modal chitchat scenarios;
- (2) We quantitatively demonstrate the characteristics of TikTalk, propose a video-based multi-modal chitchat task and summarize three new challenges brought by TikTalk;
- (3) We evaluate several baselines in two mainstream frameworks. Experimental results demonstrate that there is still a lot of room for future improvements on TikTalk.

2 RELATED WORK

2.1 Multi-modal Dialogue Datasets

Dialogue has been a subject of research interest for NLP community for many years. Many textual dialogue datasets are available [11, 25, 40, 49, 50, 55]. There have been attempts to introduce multi-modal information into dialogue datasets. We summarize main multi-modal dialogue datasets and their characteristics in Table 1.

Some datasets [2, 6, 7] aim to address multi-turn Q&A tasks. They are created with the purpose of measuring models performance on objective perception for given multi-modal information. Some studies use crowdsourcing, textual dialogue datasets, or collect comment-reply pairs from social media to construct image-based dialogue corpora. IGC [29] and Image-Chat [38] ask crowd-workers to talk about given images and related topics with explicit guidelines. In most cases, these utterances are highly image-dependent. PhotoChat [53] simulates a conversation scenario of photo sharing. MMDD [18] and DialogCC [19] attempt to align textual dialogue datasets and image caption datasets. MMChat [59] and MMDialog [13] collect images and dialogues from social media. MMDialog may contain images in both context and responses. Compared to these works, our dialogues are based on videos containing rich temporal multi-modal information, rather than static images.

Researchers have also made attempts at constructing video-based dialogue datasets from different perspectives. OpenViDial datasets [28, 45] extract dialogue between characters from movies and TV series and use corresponding screenshots as vision context. The characters’ lines are designed to drive plots and have artistic qualities. YTD-18M [14] crawls online videos and constructs dialogues through transcripts. However, the dialogue between characters in video scenes still accounts for the majority of the data, and converting transcripts into dialogues may lead to inaccuracies. MLED [33] and M³ED [58] are designed to predict the emotions of dialogues in TV series. Live comment generation [27, 48] aims to generate comments about a specific moment in a video or to

Table 1: A summary of main multi-modal dialogue datasets and their characteristics. The dialogue types can be classified into three categories, including question-answer dialogues(Q&A), dialogues that take place in multi-modal environments(in-scene), and dialogues based on multi-modal information(multi-modal based). Modalities include vision(v), text(t), and audio(a).

Dataset	Dialogue Type	Modalities	Dialogue Source	Turns	Language	Vision Source
VisDial [6]	Q&A	v, t	crowdsourcing	2.47M	English	MS COCO
AVSD [2]	Q&A	a, v, t	crowdsourcing	236K	English	Charades
IGC [29]	multi-modal based	v, t	crowdsourcing	25.3K	English	VQG
Image-Chat [38]	multi-modal based	v, t	crowdsourcing	401K	English	YFCC100M
PhotoChat [53]	multi-modal based	v, t	crowdsourcing	156K	English	OpenImage V4
MMDD [18]	multi-modal based	v, t	text datasets	346K	English	MS COCO, Flickr30K
DialogCC [19]	multi-modal based	v, t	text datasets	929K	English	CC3M
MMChat [59]	multi-modal based	v, t	social media	314K	Chinese	Weibo
MMDial [13]	multi-modal based	v, t	social media	4.92M	English	Social Media
OpenViDial [28]	in-scene	v, t	movies&TVs	1.1M	English	Movies&TVs
OpenViDial 2.0 [45]	in-scene	v, t	movies&TVs	5.6M	English	Movies&TVs
YTD-18M [14]	in-scene	a, v, t	transcripts	54M	English	YouTube
MLED [33]	in-scene	a, v, t	TVs	13.7K	English	TVs
M ³ ED [58]	in-scene	a, v, t	TVs	24.4K	Chinese	TVs
TikTalk(Ours)	multi-modal based	a, v, t	social media	827K	Chinese	Douyin

respond to other comments. However, the comments often lack clear context and contain many redundant corpora.

Our TikTalk dataset comprises of comment-reply dialogues collected from a video-sharing platform. Users engage in spontaneous chitchat based on their shared multi-modal information experiences from the videos, which helps recreate real-world dialogue context. This provides valuable information for understanding the interplay of multi-modal information in human communication.

2.2 Multi-modal Dialogue Models

With the outstanding performance of large language models [5, 17, 30, 42, 47, 56] on text-based dialogue, many studies have explored how to incorporate multi-modal information into dialogue models. Due to the diverse characteristics of multi-modal dialogue datasets, there are significant differences in task definitions and model settings. For example, PhotoChat [53] propose the task of retrieving relevant images of dialogues. Divter [41] takes dialogue context as input and generates textual response and image as visual response.

Most of existing multi-modal dialogue models take image and dialogue context as raw inputs and generate textual responses. Their frameworks can be broadly classified into two categories. The first architecture, such as Flamingo [3] and BLIP-2 [22], aims to incorporate multi-modal information into large language models. They design adapter that bridge the semantic gap between vision encoder and large language model, and use the powerful generative ability of large language model to generate responses. The second approach directly integrates information from different modalities to achieve multi-modal interaction and generate responses. Some works [14, 39, 46, 52, 59] concatenate vision and text features and input them into multi-modal fusion networks. Some studies also utilize images or context to introduce intermediate information that aids in response generation. VisAD [37] fuses image and text to obtain a multi-modal representation, and then predicts the keywords and retrieves images related to the response. VICTOR [26] leverages a caption model to obtain image descriptions and extends

concepts in the dialogue using a knowledge base. Maria [23] extracts region features and corresponding concepts from images. We evaluate some baselines on TikTalk and modify Maria by adding audio and external knowledge as inputs, respectively.

3 TIKTALK DATASET

We collect data from the Chinese version of TikTok, also known as Douyin, to construct a multi-modal chitchat dataset, which we call **TikTalk**. Douyin covers a wide range of video content across more than 25 broad categories, including science, food, gaming, travel, and more. Users have spontaneous chitchat based on the videos, which is similar to the multi-modal chitchat scenes in real world.

3.1 Data Construction

We collect videos posted on Douyin in 2021, including their titles, comments, and reply relationships. In order to protect privacy, we don't crawl any user information. Not all videos and comments constitute high-quality multi-modal dialogues, so we use some rules to select potentially useful data. First, we remove videos with only one level of comments, as these comments without replies cannot constitute dialogue corpora. Secondly, since users have the option to like videos and comment on the Douyin platform, videos and comments with high numbers of likes are more likely to meet human expectations and have higher quality. We leverage this as a screening criterion and retain videos and comments that meet a certain threshold of likes. Through such filtering, we can obtain high-quality data that is easy to scale. Furthermore, we clean the comments that we keep. We use regular expressions to filter out useless contents, such as "@[someone]" (a way of mentioning another user), merge repeated words, such as "hahaha" (laughing with different lengths), and remove some unethical comments. Comments often come with emoticons that can express users' emotions. Therefore, we move emoticons to metadata to keep only the plain text in the dialogues. Finally, we split the hierarchical comments into conversations, forming the final dialogue corpora.

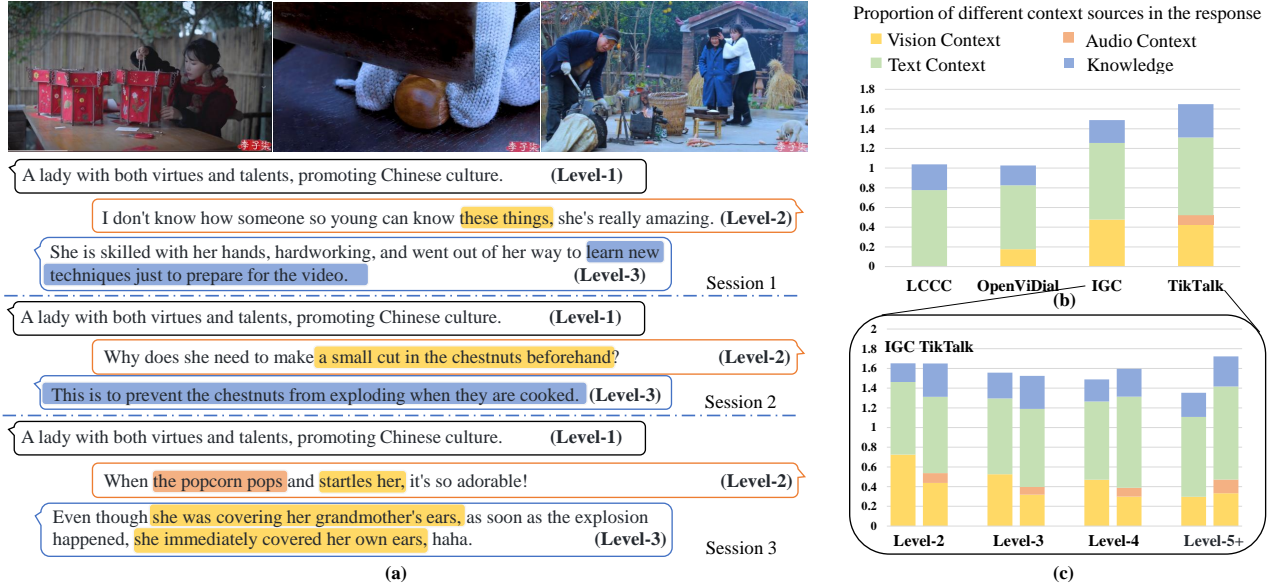


Figure 2: Examples and comparisons of responses involving different context sources. (a) Examples of response involving various context sources from TikTok. We mark in the response the parts related to vision, audio, and external knowledge context with the corresponding colors in (b). (b) Proportions of Responses involving different context sources among four dialogue datasets. (c) Proportions of responses involving different context sources across different levels in IGC and TikTok.

Table 2: Data division of TikTok. “Avg-v”, “Avg-t”, and “Avg-l” denote the average duration of videos, the average turns of dialogues, and the average length of utterances.

Split	Videos	Dialogues	Turns	Avg-v	Avg-t	Avg-l
Train	35,703	363,589	816,170	34.01s	2.24	16.42
Val	1,000	1,353	3,495	36.93s	2.58	16.06
Test	2,000	2,728	7,087	36.42s	2.60	16.25
Total	38,703	367,670	826,752	34.03s	2.25	16.41

3.2 Data Statistics

Finally, we select 38,703 videos and 367,670 dialogues from a pool of 153,340 videos to construct the TikTok dataset. The dataset is split into training set, validation set, and test set, consisting of 35,703, 1,000, and 2,000 videos, respectively. The detailed statistics can be found in Table 2. These videos have an average duration of 34.03 seconds and provide a variety of multi-modal information. Each dialogue contains an average of 2.25 turns, and due to the transient nature of the chitchat data, responses are more likely to involve video information and external knowledge. To construct high-quality validation and test sets and facilitate evaluation, we select dialogues with at least five genuine responses and chose five of them as the ground-truth responses for that dialogue.

3.3 Data Analysis

In order to analyze the characteristics of TikTok, we quantitatively compare the proportions of different context sources involved in the responses in different datasets. Moreover, we define a video-based multi-modal chitchat task and summarize three new challenges.

3.3.1 Comparison with Other Dialogue Datasets. We observe that in TikTok a response cannot be generated or even understood without certain context or knowledge. As shown in the examples in

Figure 2 (a), a user would not ask the question “why does she need to make a small cut in the chestnuts beforehand” (Level-2 utterance of Session 2) without the vision context of the middle frame. To answer this question, we need to know some external knowledge (Level-3 utterance of Session 2). Sometimes, audio context may be useful, such as in the Level-2 utterance of Session 3, where we need to know that there is a loud sound when the popcorn pops. Therefore, we are wondering whether other datasets have the same phenomena and what the fractions are in different datasets.

We choose three representative dialogue datasets for comparison and randomly pick 300 samples from each dataset: (1) LCCC [49]: textual dialogues collected from social media. There are totally 550 responses in the 300 dialogues; (2) OpenViDial [28]: dialogues and aligned key frames extracted from movies and TV series. We obtain 1,019 responses from the 300 sampled key frames; (3) IGC [29]: image-based chitchat dialogues created by crowdsourcing. There are 1,182 responses associated with the 300 sampled images. For our TikTok, there are totally 2,472 responses in the 300 sampled videos. Then, we develop an annotation website and ask annotators to view a video and text context and then judge whether they cannot understand a response if without (1) vision context, (2) audio context, (3) text context, (4) external knowledge. Note that a response can involve context information from multiple modalities. When a response is hard to judge, they can skip labeling it.

The comparison results of the four datasets are shown in Figure 2 (b). Firstly, TikTok has the highest overall proportion, indicating that it has the richest multi-modal context, including 42% vision, 10% audio and 34% knowledge. Secondly, TikTok has the highest proportion of responses derived from external knowledge (about 34%), which suggests that people tend to introduce more new information in chitchat when there is more rich multi-modal context.

Thirdly, only 18% of the responses in the OpenViDial dataset are related to vision context, suggesting that OpenViDial is less suitable for investigating how vision context affects response generation. In contrast, this proportion is 48% and 42% in IGC and TikTalk respectively, suggesting that they are more visually oriented. On the other hand, IGC has a higher visual relevance ratio, indicating that its responses rely more on images compared to TikTalk. Furthermore, except for OpenViDial, the proportion of text context is close to 80% for all three dialogue datasets. This indicates that the dialogue in OpenViDial is less articulated than daily dialogue and requires a greater understanding of the overall storyline.

To further explore how it changes when the level of utterances goes deeper in TikTalk, we compare it with IGC in Figure 2(c). The results indicate that when the level goes deeper, there are still 35% utterances are related to vision context and the proportion of audio context also stays similar. This is different from our observation in IGC, where the proportion of vision context has a significant decrease trend when the level goes deeper. This may be caused by the richer information provided by a video (in TikTalk) rather than an image (in IGC). Video users can focus on different points of interest and details in the video, and use them to continue discussions.

3.3.2 Task Definition and New Challenges. Based on the above comparisons, We formulate the task of video-based multi-modal chitchat as follows. Suppose that we have a multi-modal dialogue dataset \mathcal{D} . Each dialogue is represented as (C^v, C^a, C^t, R) , where C^v and C^a represent the vision context and audio context in the video, C^t is the textual context formed by the preceding utterances $\{u_i\}_{i=1}^n$, and R is a textual response. The goal of multi-modal chitchat is to learn a model $P(R|C^v, C^a, C^t)$ from \mathcal{D} to generate an ideal response R to the given context C^v , C^a , and C^t .

There are three new challenges in addressing this multi-modal dialogue problem on TikTalk:

1) **Multi-modal Diversity:** Due to the diversity of multi-modal information in TikTalk, users generate diverse responses involving different modalities for a given utterance, e.g., the three Level-2 utterances of different session in Figure 2 (a), which are replied by users to the same Level-1 comment after viewing the video. Without the vision context, the user cannot post the question of why to make a small cut in the chestnuts. Without the audio context, it is difficult for the users to realize the moment when the popcorn pops. Therefore, it is essential to perceive more context information and comprehend the interaction among different information.

2) **Capturing Interests:** Compared with image-based conversations, it is more challenging to automatically capture points of human interest in video to generate relevant responses. For example, users focus on different points of interest in a video and issue the three Level-2 utterances shown in Figure 2 (a). Despite having all the visual and audio information, we are still uncertain which part should be taken as the point of interest to evoke the desired response. When models use a single image as input, they can easily take it as the point of interest. However, when the image is extended to a video, it is non-trivial to decide which video or audio clips should be used as a multi-modal prompt along with a text prompt. It is also difficult to understand some human interests in video clips. For example, the Level-3 utterance in Session3 involve understanding the change of the lady covering whose ears.

3) **Complementing Knowledge:** Despite users being unaware of it, they use a lot of knowledge beyond the video when composing a response. For instance, the Level-3 utterance of Session 1 uses the background information of the video creator, Li Ziqi, who is diligent in learning new skills for preparing her videos. For another example, to reply the Level-2 utterance in Session 2, we need the knowledge that making a cut before cooking is to prevent the chestnuts from exploding. Such external knowledge mostly likely can be neither learnt from training data nor generalized for a new video. It is necessary to introduce external knowledge in generating responses somehow. Using large language models may be one way while complementing knowledge graphs may be another. More detailed statistics and analyses of TikTalk are shown in Appendix A.

4 EXPERIMENTS

4.1 Evaluate Metrics

4.1.1 Automatic Metrics. We adopt two types of automatic metrics to evaluate the performance of models from different perspectives:

- **Similarity:** For each dialogue, we calculate BLEU-2, BLEU-4 [32], Meteor [8], Rouge-L [24] and CIDEr [44]. Each conversation in the validation and test sets has five ground-truth responses. We calculate these metrics for each ground-truth response separately and take the averages as the results for each conversation. We adopt the public NLG evaluation code¹ to calculate these metrics.

- **Diversity:** We adopt Dist-2 and Dist-3 [21] to measure the diversity of generated responses.

4.1.2 Human Evaluation. To further evaluate the quality of generated responses, we conduct human evaluation on four metrics:

- **Sensibleness:** Whether the response is meaningful and coherent, and conforms to logic and common sense.
- **Specificity:** Whether the response is relevant to specific context, provides sufficient details, and avoids being too vague or general.
- **Multi-modal Relevance:** The degree of correlation between the response and the multi-modal information in the video.
- **Overall Quality:** Overall assessment of the quality of the generated response based on the video and dialogue context.

Each metric has five levels, ranging from 1 point (worst) to 5 points (best). We randomly select the generation results of 100 dialogues from the test set and ask three annotators to independently evaluate them. The final score is calculated as the mean of their ratings.

4.2 Baselines and Experimental Setup

We evaluate some state-of-the-art baselines from different tasks and settings. Existing architectures of multi-modal dialogue models can mainly be classified into two categories, as shown in Figure 3. Inspired by the findings in Section 3.3, we also have trials to modify Maria by adding audio context and knowledge.

- **DialoGPT:** [57] is an extension of GPT-2 [34] that has been pre-trained on a massive corpus of conversational data. We fine-tune a Chinese version of the DialoGPT model² on TikTalk.

- **ChatGLM-6B**³: [54] is a conversational large language model based on the GLM architecture [10]. It employs similar techniques

¹<https://github.com/Maluuba/nlg-eval>

²<https://github.com/yangjianxin1/GPT2-chitchat>

³<https://github.com/THUDM/ChatGLM-6B>

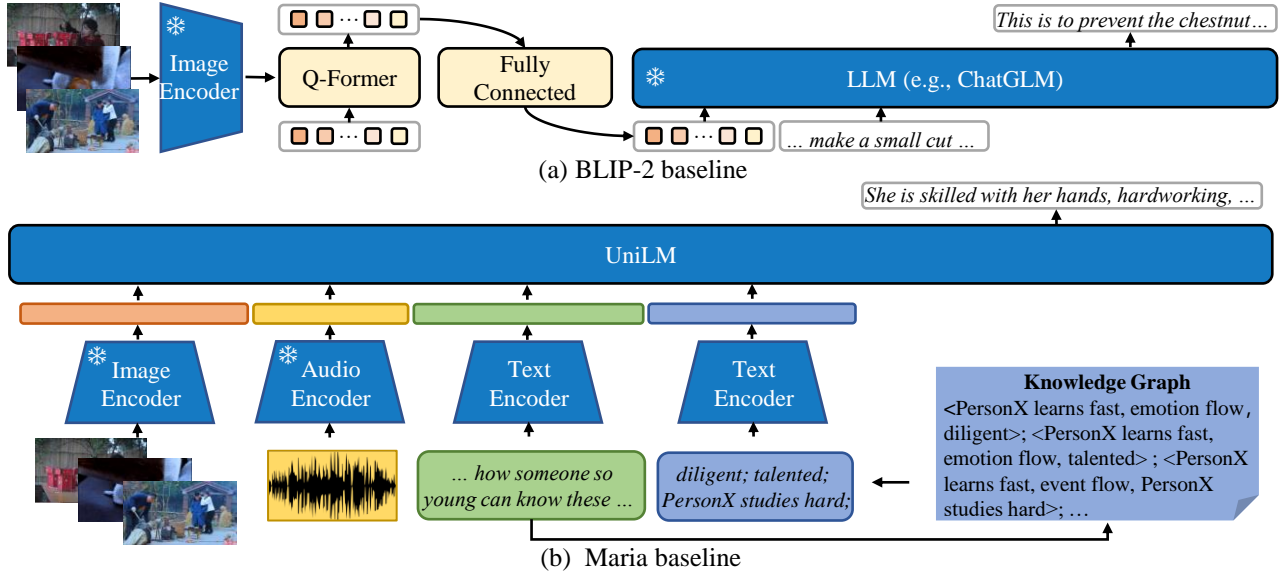


Figure 3: Two mainstream frameworks for multi-modal dialogue tasks, as exemplified by BLIP-2 and the modified model based on Maria. (a) These models design adapters (such as Q-Former) to bridge the semantic gap between frozen large language models and pre-trained vision encoders and use large language models to generate responses. (b) The extracted information from various modalities is concatenated directly to perform multi-modal interaction and fusion to generate responses. Based on Maria, we introduce audio information and external knowledge as additional inputs, respectively.

to ChatGPT. As it is biased towards receiving input prompts in the form of questions, we fine-tune it on TikTalk, which enables it to generate appropriate responses for non-Q&A type dialogues.

- **Livebot**: is proposed for live comment generation [27]. It is based on the Transformer architecture [43] and consists of three modules: a video encoder, a context encoder, and a response decoder.
- **BLIP-2**: [22] is a vision-language framework, as shown in Figure 3 (a). It extracts vision features using a frozen pre-trained image encoder, trains a Q-Former to align them with the textual semantic space, and finally uses a frozen large language model to decode the aligned vision features and textual prompts to generate outputs. We employ ViT-G/14 from EVA-CLIP [12] as visual encoder and the fine-tuned ChatGLM-6B as large language model. They are both frozen during the training of Q-Former. We explore two settings: (1) **BLIP-2_Img** encodes one randomly sampled frame of the video, which is the same as the original BLIP-2; (2) **BLIP-2_Video** extracts multiple frame features and averages them.
- **Maria**: is a visual conversation model [23] that employs the UniLM architecture [9] for multi-modal modeling. We extract consecutive video frames as visual inputs and use MCP (Masked Context Prediction) and MRP (Masked Response Prediction) as training objectives. In addition, we make two modifications to Maria, which involve introducing audio and external knowledge as inputs to Maria, respectively, as shown in Figure 3 (b). **Maria+Audio** uses the VGGish network [16] to extract audio features from the video. Then we take consecutive audio frames as another segment as video frames do. **Maria+C3KG** retrieves dialogue flows from a commonsense dialogue knowledge base as implicit context. As we observe in Figure 2, about 34% of the responses in TikTalk are related to external knowledge. We use C³KG [20], a Chinese commonsense conversational knowledge graph. Based on people’s emotions and

reactions to daily events, C³KG constructs three types of dialogue flows, including event, concept, and emotion-based empathy flow. For each conversation, we use Sentence-BERT [35] to encode textual context. By measuring the semantic similarity, we retrieve the heads of dialogue flow in C³KG which are most relevant to the dialogue context. Next, we obtain the corresponding commonsense conversational knowledge and employ it as knowledge input. In addition, we add MKP (Masked Knowledge Prediction) as a training objective. We use a probability of 15% to mask textual context and knowledge, while in response prediction the probability is 70%. The implementation details are shown in Appendix B.

4.3 Main Results

Experimental results for the automatic metrics are presented in the upper part of Table 3. The modified Maria+C³KG model outperforms the other baselines in terms of similarity metrics, and improves over Maria in all metrics. This validates the usefulness of introducing knowledge as implicit context in multi-modal dialogue. While the performance of Maria+Audio was inferior to that of Maria. There are three possible reasons for this phenomenon. Firstly, as depicted in Figure 2 (b), the proportion of responses involving audio context is less than 7%, in contrast to over 40% for visual context and over 30% for knowledge context. The available data may not be sufficient to adequately train the model to effectively utilize audio context. Secondly, unlike the text modality, the audio modality is low-level and challenging to extract the high-level information relevant to our problem. Consequently, the model may not effectively learn how to utilize audio context. In our experiment, we employ VGGish as a feature extractor, which is inclined to extract ambient sound information rather than speech information. Thirdly, given

Table 3: The evaluation results of automatic metrics for different baseline models on TikTalk (above) and the ablation study on Maria (below). We use two types of automatic metrics to evaluate the models: similarity and diversity metrics.

Type	Methods	BLEU-2	BLEU-4	Meteor	Rouge-L	CIDEr	Dist-2	Dist-3
T baselines	DialoGPT	3.76	0.73	4.87	10.43	13.50	26.82	55.36
	ChatGLM-6B	3.14	0.95	3.84	9.34	16.73	36.28	58.79
V+T baselines	Livebot	0.33	0.06	1.20	4.40	5.94	0.10	0.12
	BLIP-2_Img	3.42	0.87	3.77	8.34	12.05	38.40	64.63
	BLIP-2_Video	3.41	0.75	3.72	8.11	11.61	37.32	62.44
	Maria	4.31	1.48	5.07	11.60	26.19	29.31	48.04
V+T+A baseline	Maria+Audio	3.01	0.96	4.15	10.19	20.67	24.60	41.11
V+T+K baseline	Maria+C ³ KG	4.74	1.64	5.32	11.84	26.95	32.82	54.62
Maria ablations	text	3.35	1.18	4.56	11.16	24.90	26.83	43.80
	Audio+text	2.92	0.95	4.13	10.19	20.52	23.95	39.78
	C ³ KG+text	4.34	1.56	5.09	11.65	26.78	31.24	51.38

Table 4: Results of human evaluation. We evaluate the responses on sensibility(Sensible.), specificity(Specific.), multi-modal relevance(MM Relevant.), and overall quality(Quality).

Methods	Sensible.	Specific.	MM Relevant.	Quality
DialoGPT	2.37	1.91	1.73	2.00
ChatGLM-6B	2.87	2.12	2.07	2.36
Livebot	2.29	1.03	1.68	1.68
BLIP-2_Img	2.66	2.25	2.03	2.33
BLIP-2_Video	2.79	2.43	2.14	2.48
Maria	3.06	2.27	2.15	2.53
Maria+Audio	2.98	2.17	2.15	2.45
Maria+C ³ KG	3.29	2.41	2.31	2.71

that the majority of cases heavily depends on visual and textual information, the inclusion of low-level and noisy audio tokens has the potential to disrupt the original generation models. Due to the rich knowledge contained in large language models, the models that utilize ChatGLM-6B perform exceptionally well on diversity metrics. Especially, BLIP-2 build upon this by incorporating visual information from videos, further enhancing the diversity of generated responses. The performance of Livebot is poor on all metrics, possibly due to its design for generating live comments, leading to the generation of numerous repetitive and similar sentences.

The human evaluation results in Table 4 show that Maria+C³KG performs the best in terms of sensibleness, multi-modal relevance, and overall quality, while BLIP-2_Video achieves the highest score in specificity. Maria+C³KG outperforms Maria on all four metrics, while the performance of Maria+Audio decreases, which is consistent with the results of automatic metrics. ChatGLM-6B outperforms DialoGPT thanks to its larger pre-training corpus and the training process similar to ChatGPT, which enable such large language models to possess inherent knowledge. BLIP-2_Video performs better than BLIP-2_Img, and we attribute this to the fact that it takes into account the entire video information. Their responses are more specific compared to ChatGLM-6B, possibly because they streamline the visual information and have clearer focuses.

4.4 Ablation Study

The results of our ablation study are shown in the last three rows of Table 3. We evaluate Maria, Maria+Audio, and Maria+C³KG by removing visual information. We observe decrease in the results

of the three models, which shows that the visual context enhance the quality of generated responses. In contrast, the model that combines audio and text performs worse than the one that only includes text, suggesting that simply fusing audio features is not a viable solution. On the other hand, the Maria+C³KG model without visual information, i.e. C³KG+text, shows the smallest degradation among the three ablations, and even slightly outperforms Maria. This suggests that external knowledge is equally important as visual context, which is consistent with the analysis results in Section 3.3.

4.5 Case Study

To further investigate the performance of the baselines, we present some response cases in Figure 4. In Figure 4(a), Maria+C³KG is able to infer the reason for the lyric allocation. Figure 4(b) shows that the baselines using the large language model can provide responses based on common sense. Impressively, the response from Maria+Audio contains audio information that is not mentioned in the textual context in Figure 4(c). These three cases illustrate that the external knowledge and the audio information all play a role in multi-modal dialogue. However, as shown in Figure 4(d), for more difficult logical reasoning problems, these models may be powerless and produce unreasonable responses. And no model can produce satisfactory responses in all these cases. In summary, due to the diversity of video scenarios, TikTalk requires complex perception, understanding, and reasoning capabilities. While existing baselines can sometimes produce reasonable responses, they fall far short of meeting the expectation of multi-modal dialogue in the real world.

5 CONCLUSION

In this paper, we introduce TikTalk, a video-based multi-modal chitchat dataset. We collect videos from a video-sharing platform and users’ spontaneous dialogues based on them. We analyze the characteristics of TikTalk through quantitative comparisons with other dialogue datasets. Additionally, we propose a multi-modal chitchat task and summarize three new challenges posed by TikTalk, including perceiving and understanding more diverse multi-modal context and information interaction, capturing the human interests from rich information, and supplementing with external knowledge. We evaluate some baselines and experimental results show that existing methods exhibit some advantages, but they cannot address the above challenges well and generate ideal responses. There still



Figure 4: Examples of generated response from models. (a) Two singers are singing on the stage. (b) Some people are making zongzi. (c) A video blogger is recording a new song. (d) An actress is learning how to use a prop gun from the crew.

remains a considerable gap between the performance of multi-modal dialogue models and humans, indicating great potential for further improvements on TikTalk. We hope this work can contribute to the research on multi-modal dialogue.

6 LIMITATIONS AND FUTURE WORK

For the new challenges presented by TikTalk, existing baselines have demonstrated some abilities but still have limitations. Firstly, simple fusion of audio information does not lead to improved model performance. Experimental results indicate that we need better methods to leverage audio in videos, including sound and transcripts in text. Secondly, current models rely on attention mechanisms to capture the points of human interest in videos, while such implicit way may not be optimal. In the future, we may try ground utterances with frames or regions in frames to explicitly supervise our algorithm to automatically capture human interests. Third, knowledge is surprisingly helpful. Although the BLIP-2 based on LLM performs worse than Maria integrated by a knowledge graph, it is still open whether the BLIP-2 model based on LLM can leverage the knowledge graph too. Actually they may excel in different knowledge domains. LLM knows why the chestnuts shall be made a cut before cooking while the knowledge graph that we use knows better that a person may possess many skills due to one's diligence. The flows in C³KG is like a chain-of-thoughts (CoT) in the emotion, reaction, and event domains. If we can enable LLM the ability of

CoT in these domains somehow, perhaps the BLIP-2 based on LLM can think step by step and infer the appropriate response.

ACKNOWLEDGMENTS

This work is supported by the Fundamental Research Funds for the Central Universities, and the Research Funds of Renmin University of China (21XNLG28) and National Natural Science Foundation of China (No. 62276268).

REFERENCES

- [1] Daniel Adiwardana, Minh-Thang Luong, David R So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, et al. 2020. Towards a human-like open-domain chatbot. *arXiv preprint arXiv:2001.09977* (2020).
- [2] Huda Alamri, Vincent Cartillier, Abhishek Das, Jue Wang, Anoop Cherian, Irfan Essa, Dhruv Batra, Tim K Marks, Chiori Hori, Peter Anderson, et al. 2019. Audio visual scene-aware dialog. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7558–7567.
- [3] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems* 35 (2022), 23716–23736.
- [4] Siqi Bao, Huang He, Fan Wang, Hua Wu, Haifeng Wang, Wenquan Wu, Zhihua Wu, Zhen Guo, Hua Lu, Xinxian Huang, et al. 2022. PLATO-XL: Exploring the Large-scale Pre-training of Dialogue Generation. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2022*. 107–118.
- [5] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416* (2022).

- [6] Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José MF Moura, Devi Parikh, and Dhruv Batra. 2017. Visual dialog. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 326–335.
- [7] Harm De Vries, Florian Strub, Sarah Chandar, Olivier Pietquin, Hugo Larochelle, and Aaron Courville. 2017. Guesswhat?! visual object discovery through multi-modal dialogue. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5503–5512.
- [8] Michael Denkowski and Alon Lavie. 2014. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the ninth workshop on statistical machine translation*. 376–380.
- [9] Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. *Advances in neural information processing systems* 32 (2019).
- [10] Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. 2022. GLM: General Language Model Pretraining with Autoregressive Blank Infilling. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 320–335.
- [11] Nouha Dziri, Ehsan Kamalloo, Kory Mathewson, and Osmar R Zaiane. 2019. Augmenting Neural Response Generation with Context-Aware Topical Attention. In *Proceedings of the First Workshop on NLP for Conversational AI*. 18–31.
- [12] Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. 2022. Eva: Exploring the limits of masked visual representation learning at scale. *arXiv preprint arXiv:2211.07636* (2022).
- [13] Jiazhan Feng, Qingfeng Sun, Can Xu, Pu Zhao, Yaming Yang, Chongyang Tao, Dongyan Zhao, and Qingwei Lin. 2022. MMDialoG: A Large-scale Multi-turn Dialogue Dataset Towards Multi-modal Open-domain Conversation. *arXiv preprint arXiv:2211.05719* (2022).
- [14] Seungju Han, Jack Hessel, Nouha Dziri, Yejin Choi, and Youngjae Yu. 2023. CHAMPAGNE: Learning Real-world Conversation from Large-Scale Web Videos. *arXiv preprint arXiv:2303.09713* (2023).
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [16] Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, et al. 2017. CNN architectures for large-scale audio classification. In *2017 IEEE international conference on acoustics, speech and signal processing (icassp)*. IEEE, 131–135.
- [17] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. 2022. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556* (2022).
- [18] Nyoungwoo Lee, Suwon Shin, Jaegul Choo, Ho-Jin Choi, and Sung-Hyon Myaeng. 2021. Constructing Multi-Modal Dialogue Dataset by Replacing Text with Semantically Relevant Images. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*. 897–906.
- [19] Young-Jun Lee, Byungsoo Ko, Han-Gyu Kim, and Ho-Jin Choi. 2022. DialogCC: Large-Scale Multi-Modal Dialogue Dataset. *arXiv preprint arXiv:2212.04119* (2022).
- [20] Dawei Li, Yanran Li, Jiayi Zhang, Ke Li, Chen Wei, Jianwei Cui, and Bin Wang. 2022. C3KG: A Chinese Commonsense Conversation Knowledge Graph. In *Findings of the Association for Computational Linguistics: ACL 2022*. 1369–1383.
- [21] Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and William B Dolan. 2016. A Diversity-Promoting Objective Function for Neural Conversation Models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 110–119.
- [22] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. BliP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597* (2023).
- [23] Zujie Liang, Huang Hu, Can Xu, Chongyang Tao, Xiubo Geng, Yining Chen, Fan Liang, and Daxin Jiang. 2021. Maria: A Visual Experience Powered Conversational Agent. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. 5596–5611.
- [24] Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*. 74–81.
- [25] Pierre Lison and Jörg Tiedemann. 2016. OpenSubtitles2016: Extracting Large Parallel Corpora from Movie and TV Subtitles. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*. 923–929.
- [26] Yiting Liu, Liang Li, Beichen Zhang, and Qingming Huang. 2022. Think Beyond Words: Exploring Context-Relevant Visual Commonsense for Diverse Dialogue Generation. In *Findings of the Association for Computational Linguistics: EMNLP 2022*. 3106–3117.
- [27] Shuming Ma, Lei Cui, Damai Dai, Furu Wei, and Xu Sun. 2019. Livebot: Generating live video comments based on visual and textual contexts. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 6810–6817.
- [28] Yuxian Meng, Shuhe Wang, Qinghong Han, Xiaofei Sun, Fei Wu, Rui Yan, and Jiwei Li. 2020. Opnvidial: A large-scale, open-domain dialogue dataset with visual contexts. *arXiv preprint arXiv:2012.15015* (2020).
- [29] Nasrin Mostafazadeh, Chris Brockett, William B Dolan, Michel Galley, Jianfeng Gao, Georgios Spithourakis, and Lucy Vanderwende. 2017. Image-Grounded Conversations: Multimodal Context for Natural Question and Response Generation. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. 462–472.
- [30] OpenAI. 2022. Introducing chatgpt. <https://openai.com/blog/chatgpt>.
- [31] OpenAI. 2023. GPT-4 Technical Report. [arXiv:2303.08774](https://arxiv.org/abs/2303.08774) [cs.CL].
- [32] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*. 311–318.
- [33] Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2019. MELD: A Multimodal Multi-Party Dataset for Emotion Recognition in Conversations. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 527–536.
- [34] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1, 8 (2019), 9.
- [35] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 3982–3992.
- [36] Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, et al. 2021. Recipes for Building an Open-Domain Chatbot. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*. 300–325.
- [37] Lei Shen, Haolan Zhan, Xin Shen, Yonghao Song, and Xiaofang Zhao. 2021. Text is not enough: Integrating visual impressions into open-domain dialogue generation. In *Proceedings of the 29th ACM International Conference on Multimedia*. 4287–4296.
- [38] Kurt Shuster, Samuel Humeau, Antoine Bordes, and Jason Weston. 2020. Image-Chat: Engaging Grounded Conversations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 2414–2429.
- [39] Kurt Shuster, Eric Michael Smith, Da Ju, and Jason Weston. 2021. Multi-Modal Open-Domain Dialogue. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 4863–4883.
- [40] Alessandro Sordani, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and William B Dolan. 2015. A Neural Network Approach to Context-Sensitive Generation of Conversational Responses. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 196–205.
- [41] Qingfeng Sun, Yujing Wang, Can Xu, Kai Zheng, Yaming Yang, Huang Hu, Fei Xu, Jessica Zhang, Xiubo Geng, and Daxin Jiang. 2022. Multimodal Dialogue Response Generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 2854–2866.
- [42] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. LLaMA: Open and Efficient Foundation Language Models. *arXiv preprint arXiv:2302.13971* (2023).
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [44] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4566–4575.
- [45] Shuhe Wang, Yuxian Meng, Xiaoya Li, Xiaofei Sun, Rongbin Ouyang, and Jiwei Li. 2021. OpenViDial 2.0: A Larger-Scale, Open-Domain Dialogue Generation Dataset with Visual Contexts. *arXiv preprint arXiv:2109.12761* (2021).
- [46] Shuhe Wang, Yuxian Meng, Xiaofei Sun, Fei Wu, Rongbin Ouyang, Rui Yan, Tianwei Zhang, and Jiwei Li. 2021. Modeling text-visual mutual dependency for multi-modal dialog generation. *arXiv preprint arXiv:2105.14445* (2021).
- [47] Shuohuan Wang, Yu Sun, Yang Xiang, Zhihua Wu, Siyu Ding, Weibao Gong, Shikun Feng, Junyuan Shang, Yanbin Zhao, Chao Pang, et al. 2021. Ernie 3.0 titan: Exploring larger-scale knowledge enhanced pre-training for language understanding and generation. *arXiv preprint arXiv:2112.12731* (2021).
- [48] Weiyang Wang, Jieting Chen, and Qin Jin. 2020. Videoic: A video interactive comments dataset and multimodal multitask learning for comments generation. In *Proceedings of the 28th ACM International Conference on Multimedia*. 2599–2607.
- [49] Yida Wang, Pei Ke, Yinhe Zheng, Kaili Huang, Yong Jiang, Xiaoyan Zhu, and Minlie Huang. 2020. A large-scale chinese short-text conversation dataset. In

- Natural Language Processing and Chinese Computing: 9th CCF International Conference, NLPCC 2020, Zhengzhou, China, October 14–18, 2020, Proceedings, Part I* 9. Springer, 91–103.
- [50] Yu Wu, Wei Wu, Chen Xing, Ming Zhou, and Zhoujun Li. 2017. Sequential Matching Network: A New Architecture for Multi-turn Response Selection in Retrieval-Based Chatbots. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 496–505.
 - [51] An Yang, Junshu Pan, Junyang Lin, Rui Men, Yichang Zhang, Jingren Zhou, and Chang Zhou. 2022. Chinese clip: Contrastive vision-language pretraining in chinese. *arXiv preprint arXiv:2211.01335* (2022).
 - [52] Ze Yang, Wei Wu, Huang Hu, Can Xu, Wei Wang, and Zhoujun Li. 2021. Open domain dialogue generation with latent images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 14239–14247.
 - [53] Xiaoxue Zang, Lijuan Liu, Maria Wang, Yang Song, Hao Zhang, and Jindong Chen. 2021. PhotoChat: A Human-Human Dialogue Dataset With Photo Sharing Behavior For Joint Image-Text Modeling. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. 6142–6152.
 - [54] Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, Weng Lam Tam, Zixuan Ma, Yufei Xue, Jidong Zhai, Wenguang Chen, Zhiyuan Liu, Peng Zhang, Yuxiao Dong, and Jie Tang. 2023. GLM-130B: An Open Bilingual Pre-trained Model. In *The Eleventh International Conference on Learning Representations (ICLR)*. <https://openreview.net/forum?id=-Aw0rrrPUF>
 - [55] Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing Dialogue Agents: I have a dog, do you have pets too?. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 2204–2213.
 - [56] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068* (2022).
 - [57] Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. DialoGPT: Large-Scale Generative Pre-training for Conversational Response Generation. In *ACL, system demonstration*.
 - [58] Jinming Zhao, Tenggao Zhang, Jingwen Hu, Yuchen Liu, Qin Jin, Xinchao Wang, and Haizhou Li. 2022. M3ED: Multi-modal Multi-scene Multi-label Emotional Dialogue Database. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 5699–5710.
 - [59] Yinhe Zheng, Guanyi Chen, Xin Liu, and Jian Sun. 2022. MMChat: Multi-Modal Chat Dataset on Social Media. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*. 5778–5786.
 - [60] Hao Zhou, Pei Ke, Zheng Zhang, Yuxian Gu, Yinhe Zheng, Chujie Zheng, Yida Wang, Chen Henry Wu, Hao Sun, Xiaocong Yang, et al. 2021. Eva: An open-domain chinese dialogue system with large-scale generative pre-training. *arXiv preprint arXiv:2108.01547* (2021).

A STATISTICAL DETAILS OF TIKTALK

A.1 Statistics of Topics

We gather topic category names from Douyin and assign videos to these categories. Actually, we employ Sentence-BERT to calculate the similarity between their descriptions and the topic category names. We show the number of videos in each category in table 5.

Table 5: Statistics of videos in different Topics on TikTalk.

Topic Videos	Science	Short play	Gaming	Lifestyle	Recording
	8448	4684	4265	3592	2487
	Food	Creativity	Parenting	Campus	Others
	1973	1660	1596	822	9176

A.2 Distribution of Human Interests

Human interests encompass various aspects of a video that capture users’ attention, prompting them to respond based on their specific points of interest. For instance, as depicted in Figure 2, the three Level-1 comments illustrate distinct user focuses: one user directs attention to the video clip demonstrating the process of cutting chestnuts, while another expresses interest in the auditory and visual experience of popcorn exploding. Consequently, human interests involve multiple modalities and necessitate a nuanced comprehension that evolves over time.

We analyze the human interests in visual modality. Specifically, we conduct preliminary statistical analysis to examine the correlation between visual information at various time segments and human responses. Utilizing the Chinese-CLIP [51] model, we measure the similarity between responses annotated as relevant to the visual context and all video frames. Subsequently, we identify the time segment in the video that contains the most similar frame. We then perform statistical analysis across different time segment buckets and present the results in table 6. It indicates that users display relatively uniform interests over time in responses related to the visual context.

Table 6: The distribution of human interests in the visual modality across different time segments in videos.

Time Segment (%)	0~20	20~40	40~60	60~80	80~100
Relevant	224	172	229	194	227
Responses	(21.4%)	(16.4%)	(21.9%)	(18.6%)	(21.7%)

A.3 Classification of External Knowledge

We classify external knowledge into two primary categories: general knowledge and personalized knowledge. For instance, in Session 2 of Figure 2, the Level-3 commentator provides an explanation

regarding the necessity of making a small cut in chestnuts before cooking, which falls under general knowledge. On the other hand, personalized knowledge encompasses information pertaining to the video blogger, as well as the personal experiences and preferences of users. For instance, in Session 1 of Figure 2, the Level-3 commentator demonstrates awareness of the video blogger’s background and highlights the extensive preparations made in advance for the video, explaining this to the Level-2 commentator. Users may watch other videos by the same blogger to acquire additional knowledge, thereby combining personalized information to formulate their responses. To evaluate these categories, we randomly assign 100 responses related to external knowledge annotated in Figure 2 (b). The results indicate that 72 out of 100 examples are classified as general knowledge, while the remaining 28 examples are associated with personalized knowledge.

B IMPLEMENTATION DETAILS

In our implementation of Maria, we adopt ResNet101 [15] to extract the video frame features. The maximum length of the input sequence is 180 textual tokens and 50 frames for both vision and audio. We initialize the parameters with a pre-trained model Bert-base-Chinese and train 20 epochs with a linear decay of learning rate from 3×10^{-5} .

We conduct statistical analysis of the training efficiency of the Maria models using various modal combinations, employing 4 Nvidia V100 (32G) GPUs. The results are presented in table 7. When incorporating an additional modal input for Maria, the training time increases by 37% to 75%. Despite this added cost, it is deemed acceptable due to the substantial improvement in generation quality achieved by combining certain modalities. Inference efficiency is also acceptable as online inference can benefit from acceleration strategies such as distributed processing, caching, and compression. In our future work, we will explore methods to expedite and optimize model training and inference.

Regarding the human evaluation, we calculate the Fleiss’ Kappa, which yields a value of 0.52 (-1~1), indicating a substantial agreement among the three annotators. We also try a small-scale test with more annotators, while the results are very similar to our current annotations.

Table 7: Traing efficiency of different combination of modalities. Text (T), Image (I), Audio (A), Knowledge (K).

Modalities	Training Efficiency(min/epoch)
M1: T	10
M2: T+I	17.5 (vs. M1, +75%)
M3: T+I+A	24 (vs. M2, +37%)
M4: T+I+K	30 (vs. M2, +71%)
M5: T+I+A+K	37 (vs. M3, +54%)