

VSTAR: A Video-grounded Dialogue Dataset for Situated Semantic Understanding with Scene and Topic Transitions

Yuxuan Wang^{1,2,3*}, Zilong Zheng^{2†}, Xueliang Zhao¹

Jinpeng Li¹, Yueqian Wang¹, Dongyan Zhao^{1,2,3,4†}

¹ Wangxuan Institute of Computer Technology, Peking University, Beijing, China

² Beijing Institute for General Artificial Intelligence (BIGAI), Beijing, China

³ Center for Data Science, AAIS, Peking University, Beijing, China

⁴ National Key Laboratory of General Artificial Intelligence, Beijing, China

{wyx, lijinpeng}@stu.pku.edu.cn, zlzheng@bigai.ai

{xl.zhao, wangyueqian, zhaody}@pku.edu.cn

<https://vstar-benchmark.github.io/>

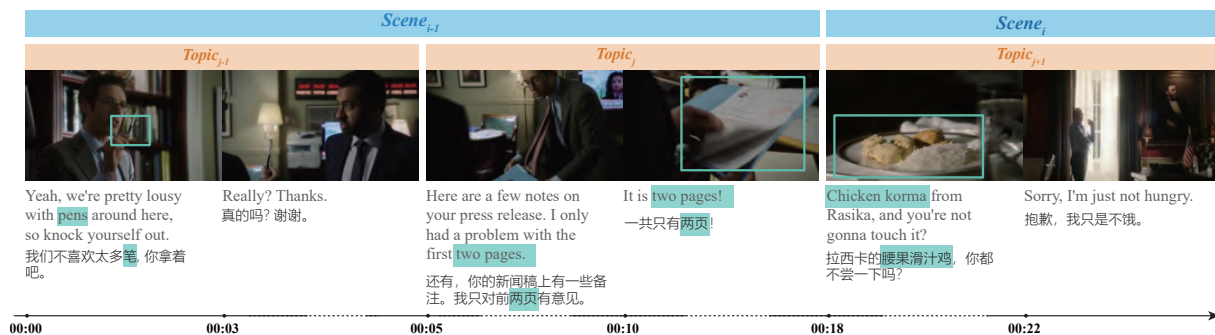


Figure 1: An overview of our VSTAR dataset. A 30 seconds video clip involves two dialogue scenes in which the environments and interlocutors are totally different.

Abstract

Video-grounded dialogue understanding is a challenging problem that requires machine to perceive, parse and reason over situated semantics extracted from weakly aligned video and dialogues. Most existing benchmarks treat both modalities the same as a frame-independent visual understanding task, while neglecting the intrinsic attributes in multimodal dialogues, such as scene and topic transitions. In this paper, we present **Video-grounded Scene&Topic AwaRe dialogue (VSTAR)** dataset, a large scale video-grounded dialogue understanding dataset based on 395 TV series. Based on VSTAR, we propose two benchmarks for video-grounded dialogue understanding: scene segmentation and topic segmentation, and one benchmark for video-grounded dialogue generation. Comprehensive experiments are performed on these benchmarks to demonstrate the importance of multimodal information and segments in video-grounded dialogue understanding and generation.

*This work was partially conducted when Yuxuan Wang was a research intern at BIGAI.

† Correspondence to Zilong Zheng and Dongyan Zhao.

1 Introduction

“Every film should have its own world, a logic and feel to it that expands beyond the exact image that the audience is seeing.”
— Christopher Nolan

Have you seen movie “Memento”? If so, you may feel astonished for the mysterious movie plots but also clap for the excellent directing and editing skills of director Christopher Nolan. In this movie, the main character plays a role of a murder victim with short-term memory. Truth, memories, fake stories are cut into pieces and re-ordered in the movie. Such nonlinear narrative structure is not unique to Nolan’s movie but is frequently used in story-based videos to increase the mysteriousness. The audience has to identify the type and chronological order of each scene in order to understand the entire movie. We call such ability as **scene transition identification**. In our daily dialogue, similarly, topic shifting commonly occurs during chit-chatting. One typical example is a speaking trick “callback” that is used by comedians, where a previously mentioned joke verbatim is repeated later. The listener has to comprehend relations between different topics in order to understand the

entire discourse. We call such ability as **topic transition identification**.

However, these abilities have barely been investigated in modern literature of multimodal conversational understanding. Over the recent years, we have witnessed a growing trend of modeling longer and more diverse video-grounded dialogues due to growing computing capacity and hardware acceleration. However, most tasks are based on multi-turn Visual Question Answering (VQA) for its simplicity in answer evaluation (Das et al., 2017; AlAmri et al., 2019), which yields huge differences compared with realistic conversations.

To fill the gap between the current multimodal dialogue systems and realistic video-grounded dialogue understanding, in this work, we introduce a new challenge, **Video-grounded Scene&Topic AwaRe dialogue (VSTAR)** dataset, aiming to address the frequent scene and topic transition within open-domain video-grounded dialogues. We start by constructing a VSTAR dataset by collecting a total of 395 TV series and 8,159 TV episodes together with their corresponding story-line and keywords; Figure 1 shows a typical example. To test the machine’s abilities on scene and topic transition identifications, we annotate video scene boundary and dialogue topic boundary through determining the semantic transition among short video-dialogue turn pairs. Table 1 shows the statistics and main differences of our new dataset compared with previous multimodal dialogue benchmark. Below, we point out the distinguished challenges existing in VSTAR:

Complicated video understanding. We carefully selected story-based TV series as the main data source. Compared with movies or homemade short videos, TV series containing more complicated plots: there are many remarkable scene transitions and topic shifts within each video clip, which brings in extra challenge in video understanding.

Multimodal scene&topic transition identification. It is worth noting that identifying scene and topic boundary in a video clip is non-trivial. Besides the complexity of long videos, both identifications require reasoning over multi-modalities. For example in Figure 1, it is hard to separate topics or scenes solely based on visual cues or dialogue text.

High-level contextual information. In VSTAR, the object-level links between video and language are much weaker than high-level semantic connections. Different from showing captions that are

directly related with videos as VQA-based tasks do, reasoning over contents in VSTAR requires the capability of high-level multimodal contextual understanding, *i.e.*, making connections between scenes, topics, and multimodal contexts.

We benchmark our dataset via three challenging tasks: scene segmentation, topic segmentation, and response generation. Detailed task formulation and evaluation are introduced in Section 4. Moreover, we propose a sliding-window based discriminative model (SWST) for segmentation tasks and an autoregressive generative model (AVDT) for dialogue generation. Extensive experiments are performed to evaluate our model and baseline methods in video-grounded dialogue understanding and generation; refer to Section 5 for detailed results and analysis.

In summary, our contributions are three-fold: (i) we collect and annotate VSTAR, a large-scale video-grounded dialogue dataset with scene and topic boundary annotations; (ii) we formalize three challenging tasks regarding video-grounded scene and topic segmentation and dialogue generation; (iii) we benchmark three tasks and analyze experimental results with baseline methods and two new transformer-based models: SWST and AVDT.

2 Related work

Recently, multimodal dialogue systems have attracted interest of many researchers with a number of benchmarks and datasets proposed. VisualDialog (Das et al., 2017; Seo et al., 2017; de Vries et al., 2017; Chattopadhyay et al., 2017; Zheng et al., 2019) treated the problem as a multi-turn VQA that aims to develop a dialogue agent that can answer questions given dialogue histories and corresponding images. Empowered by today’s high computational capacity, similar dialogue systems have been extended to the video domain (AlAmri et al., 2019; Le et al., 2021): questions about a given video are positioned in a multi-turn QA-pairs. In each QA-pair turn, a question usually exhibits different types of cross-turn relations to other questions in prior turns. To mimic the situations where common dialogues happen. Twitch-FIFA (Pasunuru and Bansal, 2018) introduces a video-context dialogue dataset based on live-broadcast soccer game videos and chats from *Twitch.tv*. Meng et al. (2020); Wang et al. (2021a) construct OpenViDial dataset based on movies and TV series. However, the subtitles extracted by OCR in this dataset contain much noise

Table 1: Comparisons of different multimodal dialogue datasets.

Dataset	Vision	Language	Scene	Topic	# Dialogues	# Turns	Turns/Clip	Words/Turn
VisualDialog	Image	QA	✗	✗	120K	2.4M	20.0	4.0
Twitch-FIFA	Live Video	Dialogue	✗	✗	15K	161K	10.4	6.0
AVSD	Recorded Video	Dialogue	✗	✗	11K	118K	20.0	9.5
MoiveNet [†]	Movies	Dialogue	✓	✗	-	421K	-	7.2
OpenViDial 2.0	Movies & TV Series	Dialogue	✗	✗	116K	5.6M	48.0	8.3
VSTAR (Ours)	TV Series	Dialogue	✓	✓	185K	4.6M	25.1	6.7

[†]: We compute statistics of the sub-dataset with scene boundary annotations.

and monologue. One of the closest work to us is the recently proposed MovieNet (Huang et al., 2020), a large scale dataset for movie understanding. In this work, shot is viewed as the minimal visual unit of a movie. However, shot boundary is annotated automatically by existing tool (Sidiropoulos et al., 2011), which leads to inaccurate predictions. Besides, few methods take the multi-modal information as input which make current models learn little high-level semantic information.

3 The VSTAR Dataset

The VSTAR is collected from 395 TV series (8,159 episodes, 185K 90-second clips) with carefully cleaned dialogues and metadata information. The collection and data cleaning details are as follow:

Data Source We carefully selected and purchased video sources from the copies of Blu-ray disks and online Reddit Open Directories. Specifically, we filtered out animation and documentary series because the former differ a lot in terms of visual input styles and are unrealistic, while the latter mostly contain monologues. To ensure the quality of selected series, we selected the TV shows rated by over 1,000 IMDb users. The ultimate genre distribution is shown in Figure 2a. As seen, there are 19 genres covering almost all common genres of TV series. For research purposes, we segmented each TV episode into 90-second video clips. In the end, we got 185K multi-modal dialogue clips.

Metadata We crawled metadata for each episode from IMDb¹ as complementary information. Each TV episode is paired with genres, keywords, and storylines. Compared to daily chat, the conversations in TV series are much longer and contain richer background knowledge. Therefore the metadata will be helpful for further dialogue understanding. We show the wordcloud of the story-lines in


¹<https://www.imdb.com>

Figure 2b. It is interesting to observe that most salient words are relevant with “work and life”.


3.1 Annotation in VSTAR

Video Scene Boundary A scene, according to the previous definition (Rasheed and Shah, 2003; Huang et al., 2020; Rao et al., 2020), is a plot-based semantic unit in which a certain activity occurs among a specific group of individuals. Recently popular methods use off-the-shelf shot segmentation tools first and then determine whether the shot boundaries are scene boundaries. Considering the detection error of the shot segmentation tools, we did not take previous shot-based methods to annotate scene boundaries. Instead, we segment the TV episode into short videos by subtitle timeline. Specifically, each short video is paired with a dialogue turn. Annotators are then asked to look through these short videos with subtitles, and find if the short video is the start of a dialogue scene. With the help of multi-modal information, the dialogue scene boundaries are clearer. Thus the annotation procedure is more efficient. In order to keep consistency with previous work, we modified the annotated boundaries as the end of a dialogue scene. Finally we got 265k dialogue scene segments with 1.4 scene boundaries in each dialogue clip on average. Comparisons between VSTAR and other datasets for video scene segmentation are shown in Table 2. We demonstrate that VSTAR is significantly larger than existing datasets. In addition, we show the distributions of the number of dialogue scenes in a TV episode in Figure 3a. The number of scene segments in a TV episode is mostly between 10-60. To the best of our knowledge, VSTAR is the first dataset whose scene boundaries are labeled with the help of multimodal information.

Dialogue Topic Boundary We perform dialogue topic boundary annotation and dialogue scene boundary annotation at the same time. Specifically, we take the video as auxiliary information to




(a) Distribution of VSTAR Genres (y-axis in log scale).




(b) Story-lines wordcloud

Figure 2: Statistics of metadata in VSTAR



(a) Distribution of # scene segments



(b) Distribution of # topic segments

Figure 3: Statistics of annotations in VSTAR

#	Scene	Video	Source
# OVSD	300	21	MiniFilm
# BBC	670	11	Documentary
# MovieScenes	21k	150	Movie
# MovieNet	42k	318	Movie
# VSTAR (Ours)	265k	8159	TV Episode

Table 2: Comparisons of dialogue scene annotation in VSTAR.

	sentence	sent/seg	Language
DiaSeg_711	19K	5.6	English
Doc2Dial	19K	3.5	English
ZYS	12K	6.4	Chinese
VSTAR	4.6M	9.3	English

Table 3: Comparisons of dialogue topic annotations.

determine whether a dialogue turn is the end of a dialogue topic. In total, we obtain 499k dialogue topic segments with 2.7 topic boundaries in each dialogue clip on average. Each scene segment contains 1.88 topic segments. Comparisons between VSTAR and other datasets for dialogue topic segmentation are shown in Table 3. VSTAR is 200× larger than previous datasets in scale. And the dia-

logue topic length is longer than current datasets, which makes the dialogue topic segmentation task more challenging. As shown in Figure 3b, the number of dialogue topics in a TV episode varies from less than 10 to more than 160, which demonstrates the diversity of VSTAR.

Annotation Process We recruited 30 highly-educated students (undergraduates and above) with high English proficiency for the annotation. Each student is assigned with 4 groups of dialogues, each of which includes 40K continuous dialogue turns. For each dialogue group, we randomly sampled 5% data and checked them manually. If the error rate is more than 4%, we asked the annotator to re-annotate the whole dialogue turn sequence. We repeated this validation procedure three times. In the end, 4% data did not meet our requirements and were all dropped off. During the annotation procedure, the salary for annotating each utterance is determined by the average time of annotation and local labor compensation standard.

4 Benchmarks and Models

We set three benchmarks based on VSTAR for video-grounded dialogue understanding and gener-

ation. In Sec.4.1, we first introduce the task formulation along with their evaluation metrics. Then in Sec.4.2 we present our proposed transformer-based video dialogue frameworks to benchmark our tasks.

4.1 Task formulation

VSTAR consists of a set of video-grounded dialogue clips $(U, V) \in \mathcal{D}$, where $U = \{u_1, \dots, u_N\}$ serves as a dialogue clip with u_i denoting the i -th dialogue utterance, $V = \{v_1, \dots, v_N\}$ signifies the corresponding video clip with v_i paired with dialogue turn u_i , N refers to the number of dialogue turns. More precisely, v_i can be separated into a sequence of RGB image frames $\{z_{i,1}, \dots, z_{i,K}\}$, where $z_{i,k}$ is a RGB tensor of frame k , K is the number of frames in v_i .

Video-grounded dialogue scene segmentation

A dialogue scene segment is a series of video-grounded dialogue pieces that sharing the same visual scene context. We thereby formulate the dialogue scene segmentation task as a binary classification problem: given a clip $(U, V) = \{(u_i, v_i), i = 1 \dots K\}$, the model is asked to predict $s_i \in \{0, 1\}$ indicating the dialogue scene boundary. We take three commonly used metrics for evaluation:

- **AP**. We compute Average Precision (AP) of $s_i = 1$ for each video piece v_i .
- **mIoU**. Following Huang et al. (2020), we use mIoU to measure the averaged intersection-over-union (IoU) between predicted dialogue scene segments and their closest ground truth dialogue scene segments.
- **micro-F1**. Inspired by Mun et al. (2022), we use micro-F1 as an additional evaluation metric to compare algorithms.

Video-grounded dialogue topic segmentation

Similar to scene segmentation, we formulate dialogue topic segmentation as a turn-level classification problem. Concretely, given a video-grounded dialogue clip (U, V) , we need to predict if the i -th dialogue utterance is the end of a dialogue topic. Following Xing and Carenini (2021), we apply three standard metrics to evaluate the performance of the proposed segmentation model in this benchmark:

- P_k **error score** (Beeferman et al., 2004). P_k is a penalty computed via a sliding window of length k .
- **WinDiff** (Pevzner and Hearst, 2002). WinDiff is calculated based on the intersection between ref-

erence segments and predicted segments within a moving window.

- **macro-F1**. We utilize the F1 score to make a balanced comparison of precision and recall.

Video-grounded dialogue response generation

For each dialogue clip (U, V) , we set the first $N - 1$ dialogue turns $\{u_1, \dots, u_{N-1}\}$ as the dialogue context C and the last dialogue turn u_N as the gold reply r . We choose four commonly used reference-based metrics: BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), METEOR (Lavie and Agarwal, 2007) and CIDEr (Vedantam et al., 2015).

4.2 Transformer-based Video Dialogue Model

In this section, we propose a transformer-based discriminative model, namely SWST, to benchmark two segmentation tasks. For the response generation task, we develop a transformer-based generative model following encoder-decoder framework (AVDT); Figure 4 depicts our overall architecture.

Sliding-window-based Segmentation Transformer (SWST)

Inspired by currently popular works (Rao et al., 2020; Chen et al., 2021; Mun et al., 2022) in video scene segmentation, we adopt the sliding window scheme to learn the contextual representation of a dialogue scene. Specifically, the window can be denoted as a pair of short video sequence $W_i^v = \{v_{i-L}, \dots, v_i, \dots, v_{i+L}\}$, and the corresponding dialogue turn sequence $W_i^u = \{u_{i-L}, \dots, u_i, \dots, u_{i+L}\}$, with (v_i, u_i) as the center of the window, K as the number of neighbor pieces before and after the center. Our goal is to train a model by maximizing the expected log-likelihood:

$$\theta^* = \arg \max_{\theta} \mathbb{E} [\log p_{\theta}(s_i | W_i^v, W_i^u)] \quad (1)$$

Figure 4A depicts the architecture of our Transformer-based video-grounded dialogue scene segmentation model. For visual feature W_i^v , we follow Lei et al. (2021) to randomly sample M_l frames for i -th video piece v_i instead of using the full-length short video. Then, we utilize ResNet-50 (He et al., 2016) pretrained on ImageNet dataset (Deng et al., 2009) to extract 1,000 dim visual features for each frame. For dialogue feature W_i^u , we utilize the same tokenizer and embedding matrix as in BERT to obtain its initialization. We concatenate the sparse-sampled short video sequence W_i^v and the dialogue turn sequence W_i^u as the model

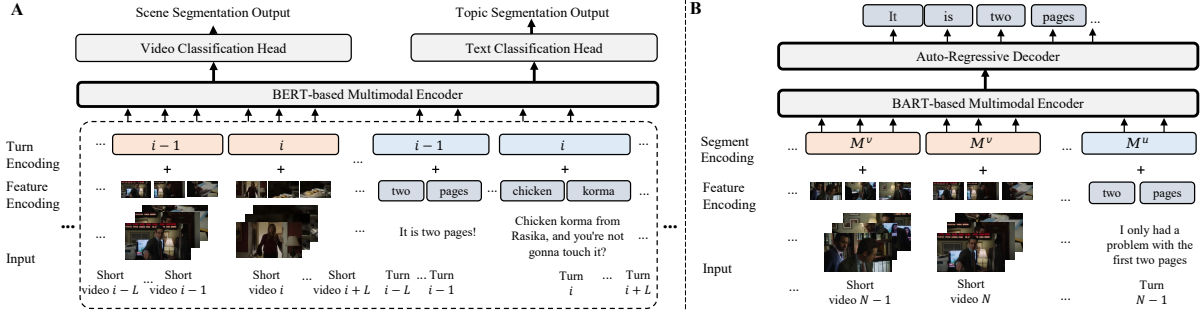


Figure 4: **Transformer-based model architecture.** (A) Sliding-window based segmentation transformer (SWST) for scene and topic boundaries identification. The dashed rectangle indicates the current sliding window for turn i . (B) Autoregressive video-grounded dialogue transformer (AVDT) for dialogue response generation.

input. The scene segmentation model leverages the same architecture and parameters in BERT² (Devlin et al., 2019) as its initialization:

$$e_i^v = f_{scene}([W_i^v; W_i^u]), \quad (2)$$

where $f_{scene} : \mathbb{R}^{(4L+2) \times D_s} \rightarrow \mathbb{R}^{(4L+2) \times D_e}$ represents the BERT-based contextual relation network, D_s and D_e denote dimensions of the input and output features, $e_i^v = \{e_{i-L}^v, \dots, e_i^v, \dots, e_{i+L}^v\}$ represents the output feature sequence. After that, we apply a dialogue scene boundary detection head h_S to predict the result for the contextualized representation. We then use cross-entropy loss to optimize the contextual relation network f_{scene} and the dialogue scene boundary detection head h_S . In the test procedure, we binarize the prediction score with a threshold $\tau = 0.5$ to get the result.

Similarly, we adopt the contextual relation network that has the same structure with f_{scene} to encode the multi-modal inputs. And we use a linear layer as the dialogue scene boundary detection head, which is optimized by the ground truth dialogue topic labels.

Autoregressive Video-grounded Dialogue Transformer (AVDT) Given a new dialogue context C associated with a video clip V , our goal is to learn a generative model $p(r|V, C; \theta)$ from dialogue \mathcal{D} . Figure 4(B) illustrates the architecture of our autoregressive generative model. The model is composed of a BART-based multi-layer Transformer (Lewis et al., 2020) and an autoregressive decoder. We concatenate dialogue context $C = \{u_i\}_{i=1}^{N-1}$ and video clip $V = \{v_i\}_{i=1}^N$ as the encoder inputs. Considering the computing complexity, we sample one frame from each short video. Same as in the segmentation model, we

use ResNet-50 to extract features of the sampled frames. We signify the dialogue scene segment sequence as $\{1, \dots, M^v\}$ and the dialogue topic segment sequence as $\{1, \dots, M^u\}$, where M^v and M^u are the number of scene segments in V and topic segments in C , respectively. Then, we add these segment tokens to inputs to learn a scene & topic aware context representation.

We adopt two other image embedding backbones to investigate the impact of the frame representation in our model: OD-based region feature extractor and ViT-based patch feature extractor. Specifically, we use a Faster R-CNN (Ren et al., 2015) trained on Visual Genome (Krishna et al., 2016) to extract the OD-based Region feature embedding. Each region feature is a 2048-d Region-of-Interest (RoI). Following ViT (Dosovitskiy et al., 2021), we reshape the frame $Z_j \in \mathbb{R}^{(C \times H \times W)}$ into a sequence of flattened 2D patches $Z_j^p \in \mathbb{R}^{N \times (P^2 \times C)}$, where P is the size of each image patch and $N = H \times W / P^2$ is the number of patches.

5 Experiments

We split the whole dataset into Train, Val, and Test with the ratio 17:1:1 on utterance level. For each task, we demonstrate the main results compared with baseline methods along with ablation studies on each module. Detailed information about baseline methods and implementation details are introduced in ?? and ??.

5.1 Video Scene Segmentation

We choose two popular methods as our baselines, **ShotCol** (Chen et al., 2021) and **Bassal** (Mun et al., 2022). In practice, we did not extract the shot boundaries with external tools but used short video pieces as video units. We additionally implement the **Random** method by setting the ratio of

²<https://huggingface.co/bert-base-uncased>

Model	mIoU \uparrow	AP \uparrow	F1 \uparrow
Random	0.251	0.060	0.075
ShotCol	0.427	0.412	0.365
Bassal	0.466	0.442	0.401
SWST <i>text-only</i>	0.453	0.351	0.380
SWST <i>video-only (1 frame)</i>	0.448	0.419	0.385
SWST <i>video-only (3 frames)</i>	0.481	0.474	0.430
SWST (3 frames)	0.536	0.543	0.503

Table 4: Results of dialogue scene segmentation task.

Model	WinDif \downarrow	P_k \downarrow	F1 \uparrow
Random	0.765	0.603	0.370
TextTiling	0.636	0.581	0.480
BERT+Greedy	0.615	0.565	0.486
BERT+CS	0.541	0.512	0.527
BERT+CS SP	0.531	0.422	0.610
SWST <i>text-only</i>	0.374	0.326	0.644
SWST <i>full</i>	0.336	0.281	0.690

Table 5: Results of the dialogue topic segmentation task.

scene segmentation label/non-segmentation label the same as in the test set. The overall results are shown in Table 4. Compared with **ShotCol** and **Bassal** which focused on learning better shot-level representations based on ResNet-50, our method using offline-extracted video features can achieve similar performance with a difference within 2 points on F1 score. Concretely, our method performs better than ShotCol (Chen et al., 2021) but worse than Bassal (Mun et al., 2022). This phenomenon shows that a better video encoder does help the model to distinguish the scene boundary. However, training a video encoder is really a time-consuming procedure.

Ablation Studies We ablated our method by adopting different lengths of sampled short videos. We find our method of taking 3 frames to represent the short video outperforms the 1-frame version. The frequency of change prediction between 1 sampled frame and 3 sampled frames is 0.137. We believe the length of frames is an essential part not to be ignored for our transformer-based model. We also compare models with uni-modal inputs against ones with multi-modal inputs. The results demonstrate that with the help of text input, our method improves from 0.481 on mIoU to 0.536 (11.4% relatively), from 0.474 on AP to 0.543 (14.6% relatively), from 0.430 on F1 to 0.503 (17.0% relatively). We are delighted to find the text information is very helpful to video scene segmentation.

5.2 Dialogue topic segmentation

We compare several currently popular baselines: **TextTiling** (Hearst, 1997), **GreedySeg** (Xu et al., 2021) and **BERT+CS** (Xing and Carenini, 2021). In addition, we train another **BERT+CS SP** model under the supervision of the ground truth label for comparison. We implement **Random** algorithm following (Xing and Carenini, 2021). The overall results are represented in Table 5. The **BERT+CS SP** model with supervision signals improves 0.527 on F1 to 0.610 (15.7% relatively). This result shows the importance of our dialogue topic boundary annotations. Compared with **BERT+CS SP** model, our sliding window-based approach SWST improves 0.610 on F1 to 0.644(5.6% relatively). To further investigate the validity of the visual information in dialogue topic segmentation, we add the video clip to the inputs. The result turns out that the visual information is important for the dialogue topic boundary detection where it leads to a performance gain of 7.1% in F1.

5.3 Video-grounded Response Generation

We choose two commonly used transformer-based models as our baselines: **CoarseVisual** (Wang et al., 2021b) and **RLM** (Li et al., 2020). Visual features in these models are all extracted through ResNet-50 pre-trained on ImageNet dataset. The overall results are shown in Table 6. Results show that our model outperforms baseline methods across all metrics.

Analysis & Ablation studies We further examined the performance of our model on different input settings (rows 4-7) and made three fundamental observations: (i) When the input is only a query (row 4), our model is slightly worse than the input is the full dialogue clip history (row 5), which demonstrates that the dialogue history containing much noise is of limited help for response generation. (ii) Though our model using the text-only setting (row 5) can reach comparable results with the standard setting (row 6), the results improve across all metrics when we increase the number of input frames (rows 6-7). Such observation is similar to Huang et al. (2021) — the utility of multi-modality is positively correlated with the data scale. (iii) Compared with the setting without segment embedding (row 6), our method performs much better. Specifically, the Rouge-L score increases from 0.077 to 0.082 with an increment of 6.5%; the CIDEr increases from 0.126 to 0.145 with an

Model	BLEU1	BLEU2	BLEU3	BLEU4	METEOR	ROUGE-L	CIDEr
OpenViDial coarse	0.075	0.026	0.013	0.006	0.035	0.063	0.066
RLM	0.072	0.032	0.017	0.010	0.032	0.061	0.079
AVDT	0.089	0.040	0.022	0.014	0.041	0.082	0.145
AVDT w/o seg query-only	0.082	0.035	0.019	0.011	0.037	0.073	0.119
AVDT w/o seg text-only	0.085	0.037	0.021	0.013	0.037	0.075	0.126
AVDT w/o seg (1 frame)	0.087	0.039	0.021	0.013	0.039	0.077	0.126
AVDT w/o seg (3 frames)	0.090	0.040	0.022	0.014	0.041	0.081	0.139
AVDT Faster-RCNN	0.089	0.040	0.022	0.014	0.040	0.080	0.137
AVDT ViT	0.092	0.041	0.023	0.014	0.041	0.082	0.144

Table 6: Results of the dialogue response generation task. AVDT denotes our Autoregressive Video-grounded Transformer. All evaluation metrics denote better generation performance with higher scores.

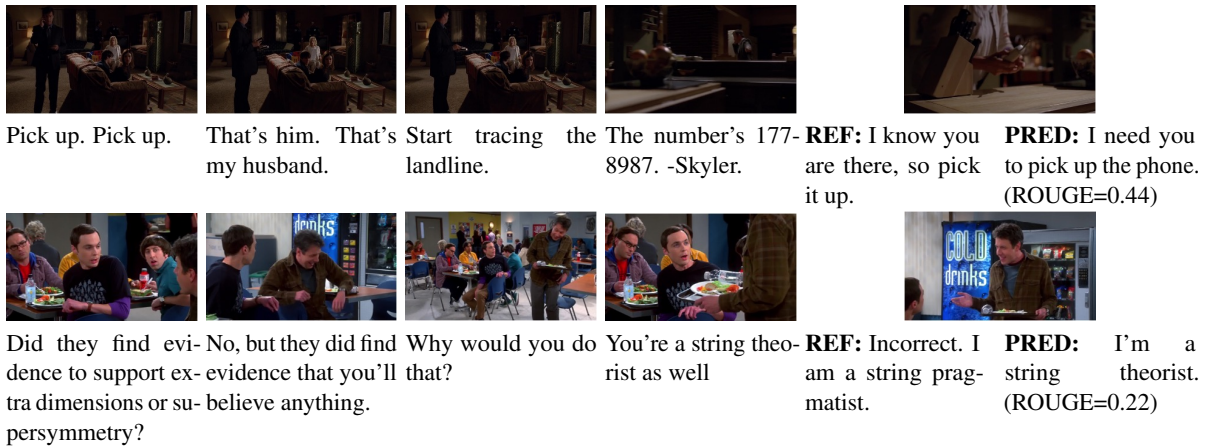


Figure 5: Examples of human evaluation. Each image indicates a short video clip, **PRED** denotes our predicted response, **REF** denotes the reference human response.

increment of 15.1%. These improvements show that segment information is important for dialogue generation in our dataset.

Additionally, we investigate the contribution of different visual backbones. Specifically, we use different frame representations while keeping other parts of our model unaltered. The results (rows 8-9) show similar performance between these feature representations and the ViT-based patch features even perform slightly better than other offline-extracted features with high computational overload. This phenomenon validates our hypothesis that the current encoder-decoder model can not make full use of the visual information for the video-grounded dialogue generation task, which yields future investigation on video-grounded dialogue modeling.

Human Evaluation We follow Sun et al. (2022) to run human evaluation by comparing our generated responses with baseline methods. Specifically, we select 20 highly educated students with proficient English skills as evaluators and randomly

	Win	Lose	Tie	Kappa
AVDT vs. OpenViDial	0.20	0.16	0.64	0.71
AVDT vs. Human	0.08	0.71	0.21	0.73

Table 7: Human evaluation result

sample 300 video-dialogue pairs with corresponding responses as test cases. To each evaluator, two responses from different models are presented, which are shuffled to hide the sources. The evaluators then judge which response is more coherent to the current dialogue scene and topic. The agreement among the annotators is measured by Fleiss' Kappa (Fleiss, 1971). Table 7 shows the human evaluation results and Figure 5 shows some qualitative comparison. We additionally add a comparison with human annotation as an upper bound.

6 Limitations and Ethics Concerns

We point out potential limitations and ethical concerns of this work.

Limitation: Data and Modeling The dialogues

in our dataset are made by playwright, which are slightly different from daily chat. Second, the automatic evaluation metrics for the response generation task can not perfectly reflect the interactivity of dialogue system. Lastly, Our autoregressive generative model simply add the segment embedding to the inputs. Similar to the position encoding in transformer, our coarse method does not make good use of the segmentation, and lacks interpretability.

Ethics: Copyright and Licensing The data source are publicly available TV series. Its collection and annotation procedure is designed for video-grounded dialogue understanding and generation purpose, and does not involve privacy issues. Following LSMDC (Rohrbach et al., 2016) and MovieNet (Huang et al., 2020), we will polish an agreement and release TV shows content under very strict conditions but will open-source all the scrawling code, pretrained features and sampled images.

7 Conclusion

In this paper, we introduce VSTAR, a scene and topic-aware video-grounded dialogue understanding and generation dataset. The main purpose of our dataset is to improve the situated multimodal semantic perception capability of the dialogue system, so that the dialogue system can generate the response that are both semantically and logically consistent with the dialogue scene in common situation. We introduce three challenging benchmarks on different aspects of video-grounded dialogue understanding and generation, *i.e.*, discovering scene transition and topic transition on video-grounded dialogue, and getting proper response. Furthermore, we propose three new baselines for corresponding benchmarks. Experiment results shows the multi-modal information can benefit dialogue understanding. Besides, we find scene boundary and topic scene boundary contribute to generate more relevant and coherent responses. By introducing VSTAR, we hope to shed light on future research towards building conversational agents that can comprehend complicated realistic multimodal signals.

Acknowledgements

This work is supported by the National Key Research and Development Program of China (No.2020AAA0106600).

References

- Huda AlAmri, Vincent Cartillier, Abhishek Das, Jue Wang, Stefan Lee, Peter Anderson, Irfan Essa, Devi Parikh, Dhruv Batra, Anoop Cherian, Tim K. Marks, and Chiori Hori. 2019. Audio visual scene-aware dialog. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7550–7559.
- Doug Beeferman, Adam L. Berger, and John D. Lafferty. 2004. Statistical models for text segmentation. *Machine Learning*, 34:177–210.
- Prithvijit Chattopadhyay, Deshraj Yadav, Viraj Prabhu, Arjun Chandrasekaran, Abhishek Das, Stefan Lee, Dhruv Batra, and Devi Parikh. 2017. Evaluating visual conversational agents via cooperative human-ai games. In *HCOMP*.
- Shixing Chen, Xiaohan Nie, David D. Fan, Dongqing Zhang, Vimal Bhat, and Raffay Hamid. 2021. Shot contrastive self-supervised learning for scene boundary detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9791–9800.
- Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José M. F. Moura, Devi Parikh, and Dhruv Batra. 2017. Visual dialog. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1080–1089.
- Harm de Vries, Florian Strub, A. P. Sarath Chandar, Olivier Pietquin, H. Larochelle, and Aaron C. Courville. 2017. Guesswhat?! visual object discovery through multi-modal dialogue. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4466–4475.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, K. Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. *International Conference on Learning Representations (ICLR)*.
- Joseph L. Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76:378–382.

- Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.
- Marti A. Hearst. 1997. Text tiling: Segmenting text into multi-paragraph subtopic passages. *Comput. Linguistics*, 23:33–64.
- Qingqiu Huang, Yu Xiong, Anyi Rao, Jiase Wang, and Dahua Lin. 2020. Movienet: A holistic dataset for movie understanding. *European Conference on Computer Vision (ECCV)*.
- Yu Huang, Chenzhuang Du, Zihui Xue, Xuanyao Chen, Hang Zhao, and Longbo Huang. 2021. What makes multimodal learning better than single (provably). In *Advances in Neural Information Processing Systems (NIPS)*.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2016. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision (IJCV)*, 123:32–73.
- Alon Lavie and Abhaya Agarwal. 2007. Meteor: An automatic metric for mt evaluation with high levels of correlation with human judgments. In *Proceedings of the second workshop on statistical machine translation*, pages 228–231.
- Hung Le, Chinnadhurai Sankar, Seungwhan Moon, Ahmad Beirami, Alborz Geramifard, and Satwik Kottur. 2021. Dvd: A diagnostic dataset for multi-step reasoning in video grounded dialogue. *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Jie Lei, Linjie Li, Luowei Zhou, Zhe Gan, Tamara L. Berg, Mohit Bansal, and Jingjing Liu. 2021. Less is more: Clipbert for video-and-language learning via sparse sampling. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7327–7337.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Zekang Li, Zongjia Li, Jinchao Zhang, Yang Feng, Cheng Niu, and Jie Zhou. 2020. Bridging text and video: A universal multimodal transformer for video-audio scene-aware dialog. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Proceedings of the ACL-02 Workshop on Automatic Summarization*, pages 74–81.
- Yuxian Meng, Shuhe Wang, Qinghong Han, Xiaofei Sun, Fei Wu, Rui Yan, and Jiwei Li. 2020. Openvidial: A large-scale, open-domain dialogue dataset with visual contexts. *ArXiv*, abs/2012.15015.
- Jonghwan Mun, Minchul Shin, Gunsoo Han, Sangho Lee, Seong Jong Ha, Joonseok Lee, and Eun-Sol Kim. 2022. Boundary-aware self-supervised learning for video scene segmentation. *ArXiv*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 311–318. Association for Computational Linguistics.
- Ramakanth Pasunuru and Mohit Bansal. 2018. Game-based video-context dialogue. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Lev Pevzner and Marti A. Hearst. 2002. A critique and improvement of an evaluation metric for text segmentation. *Computational Linguistics*, 28:19–36.
- Anyi Rao, Linning Xu, Yu Xiong, Guodong Xu, Qingqiu Huang, Bolei Zhou, and Dahua Lin. 2020. A local-to-global approach to multi-modal movie scene segmentation. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10143–10152.
- Zeeshan Rasheed and Mubarak Shah. 2003. Scene detection in hollywood movies and tv shows. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2:II–343.
- Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39:1137–1149.
- Anna Rohrbach, Atousa Torabi, Marcus Rohrbach, Niket Tandon, Christopher Joseph Pal, H. Larochelle, Aaron C. Courville, and Bernt Schiele. 2016. Movie description. *International Journal of Computer Vision (IJCV)*, 123:94–120.
- Paul Hongsuck Seo, Andreas M. Lehrmann, Bohyung Han, and Leonid Sigal. 2017. Visual reference resolution using attention memory for visual dialog. In *Advances in Neural Information Processing Systems (NIPS)*.
- Panagiotis Sidiropoulos, Vasileios Mezaris, Yiannis Kompatsiaris, Hugo Meinedo, Miguel M. F. Bugalho, and Isabel Trancoso. 2011. Temporal video segmentation to scenes using high-level audiovisual features. *IEEE Transactions on Circuits and Systems for Video Technology*, 21:1163–1177.

- 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 Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 2854–2866, Dublin, Ireland. Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL).
- 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 (CVPR)*, pages 4566–4575.
- Shuhe Wang, Yuxian Meng, Xiaoya Li, Xiaofei Sun, Rongbin Ouyang, and Jiwei Li. 2021a. Openvidial 2.0: A larger-scale, open-domain dialogue generation dataset with visual contexts. *ArXiv*, abs/2109.12761.
- Shuhe Wang, Yuxian Meng, Xiaofei Sun, Fei Wu, Rongbin Ouyang, Rui Yan, Tianwei Zhang, and Jiwei Li. 2021b. Modeling text-visual mutual dependency for multi-modal dialog generation. *ArXiv*, abs/2105.14445.
- Linzi Xing and Giuseppe Carenini. 2021. Improving unsupervised dialogue topic segmentation with utterance-pair coherence scoring. In *Proceedings of the Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*.
- Yi Xu, Hai Zhao, and Zhuosheng Zhang. 2021. Topic-aware multi-turn dialogue modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Zilong Zheng, Wenguan Wang, Siyuan Qi, and Song-Chun Zhu. 2019. Reasoning visual dialogs with structural and partial observations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Left blank.
- A2. Did you discuss any potential risks of your work?
Left blank.
- A3. Do the abstract and introduction summarize the paper’s main claims?
Left blank.
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Left blank.

- B1. Did you cite the creators of artifacts you used?
Left blank.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Left blank.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Left blank.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Left blank.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Left blank.

C Did you run computational experiments?

Left blank.

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Left blank.

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Left blank.

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Left blank.

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Left blank.

D **Did you use human annotators (e.g., crowdworkers) or research with human participants?**

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

Left blank.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

Left blank.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

Left blank.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

Left blank.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Left blank.