

Scene-Text Grounding for Text-Based Video Question Answering

Sheng Zhou^{1b}, Junbin Xiao^{1b}, Xun Yang^{1b}, Peipei Song^{1b}, Dan Guo^{1b}, *Member, IEEE*,
 Angela Yao^{1b}, *Member, IEEE*, Meng Wang^{1b}, *IEEE Fellow*, Tat-Seng Chua^{1b}

Abstract— Existing efforts in text-based video question answering (TextVideoQA) are criticized for their opaque decision-making and heavy reliance on scene-text recognition. In this paper, we propose to study *Grounded TextVideoQA* by forcing models to answer questions and spatio-temporally localize the relevant scene-text regions, thus decoupling QA from scene-text recognition and promoting research towards interpretable QA. The task has three-fold significance. First, it encourages scene-text evidence versus other short-cuts for answer predictions. Second, it directly accepts scene-text regions as visual answers, thus circumventing the problem of ineffective answer evaluation by stringent string matching. Third, it isolates the challenges inherited in VideoQA and scene-text recognition. This enables the diagnosis of the root causes for failure predictions, e.g., wrong QA or wrong scene-text recognition? To achieve Grounded TextVideoQA, we propose the T2S-QA model that highlights a disentangled temporal-to-spatial contrastive learning strategy for weakly-supervised scene-text grounding and grounded TextVideoQA. To facilitate evaluation, we construct a new dataset *ViTXT-GQA* which features 52K scene-text bounding boxes within 2.2K temporal segments related to 2K questions and 729 videos. With ViTXT-GQA, we perform extensive experiments and demonstrate the severe limitations of existing techniques in Grounded TextVideoQA. While T2S-QA achieves superior results, the large performance gap with human leaves ample space for improvement. Our further analysis of oracle scene-text inputs posits that the major challenge is scene-text recognition. To advance the research of Grounded TextVideoQA, our dataset and code are at <https://github.com/zhousheng97/ViTXT-GQA.git>

Index Terms—Text-Based Video Question Answering, Scene Text Grounding, Interpretability.

I. INTRODUCTION

TEXT-based Video Question Answering (TextVideoQA) [55], [41], [21] is an emerging task that requires models to answer questions pertaining to scene texts in dynamic visual

This work was supported in part by the National Natural Science Foundation of China (U22A2094, 62272144, 72188101, 62020106007, and 62272435), the Major Project of Anhui Province (2408085J040, 202203a05020011), and the Fundamental Research Funds for the Central Universities (JZ2024HGTTG0309, JZ2024AHST0337, and JZ2023YQTD0072). (Corresponding authors: Junbin Xiao; Xun Yang; Dan Guo.)

Sheng Zhou, Dan Guo, and Meng Wang are with the School of Computer Science and Information Engineering, Hefei University of Technology, Hefei, 230601, China (email: hzgn97@gmail.com; guodan@hfut.edu.cn; eric.mengwang@gmail.com).

Junbin Xiao, Angela Yao, and Tat-Seng Chua are with the School of Computing, National University of Singapore, 117418, Singapore. (email: {junbin, ayao, chuats}@comp.nus.edu.sg).

Xun Yang and Peipei Song are with the School of Information Science and Technology, University of Science and Technology of China, Hefei, 230026, China. (email: xyang21@ustc.edu.cn; beta.songpp@gmail.com).

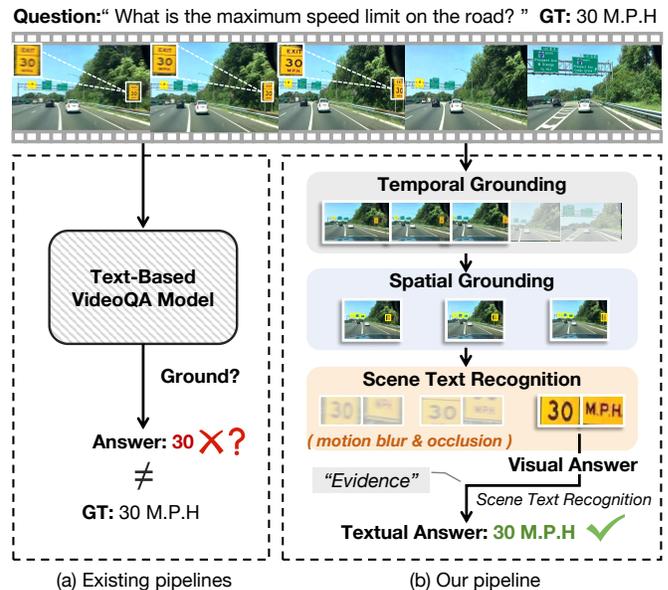


Fig. 1: Comparison between existing research and our work for TextVideoQA. (a) Existing research has two major problems: 1) Opaque decision-making; they hardly tell if their answers (e.g., “30”) are originated from the relevant scene texts in the videos, or attributed to other short-cuts. 2) Heavy reliance on scene-text recognition; their low QA accuracy could be due to a failure in decoding the textual answer (e.g., “30 M.P.H.”) from the corresponding scene text region. (b) We establish a novel pipeline by temporal-spatially localizing the scene text region and then decoding them into textual answers. We also enable direct evaluation on the grounded scene-text region.

contents. It challenges the models in identifying a very limited portion of crucial scene texts from a large amount of visual content to answer a particular question, in which motion blur and occlusion often result in wrong predictions (refer to the example in Fig. 1 (a)). Existing methods [22], [60], [11] that perform well on image-based TextVQA [39], [2] struggle to achieve good performances in the video domain [55]. However, the key factors causing performance loss remain unclear due to the opaque decision-making process. For example, is it because of poor QA or poor scene-text recognition in the video? Additionally, even for the correct predictions, these methods rarely tell if their answers are originated from relevant scene texts in the videos, or attributed to other short-cut paths. This severely impedes further improvements.

To tackle this problem, we propose to study a novel *Grounded TextVideoQA* task by requiring models to answer the question and localize the relevant scene-text regions as evidence. Such a setting enjoys three-fold advantages: **First**, the grounded scene-text regions serve as visual evidence to support textual answers, thus enabling a reliable TextVideoQA system. **Second**, the grounded scene-text regions can serve as visual answers in scenarios where the accuracy metric fails to align with human intuition to effectively measure the predictions' correctness. For instance, the answer "30" in Fig. 1 (a) should be an acceptable (understandable) answer to human. **Third**, by allowing both visual answers and text answers, we decouple scene-text recognition from VideoQA, thus enabling better diagnose the source of the problem if the models arrive at a wrong prediction. This significantly benefits model debugging towards improvements.

To solve *Grounded TextVideoQA*, precisely grounding (localizing) the question-relevant scene texts is a key challenge, especially in both temporal and spatial dimensions under weak supervision. To tackle this, we propose a *Temporal-to-Spatial (T2S)* grounding and then *Question-Answering (QA)* model T2S-QA. T2S-QA highlights a temporally- and spatially- disentangled contrastive learning strategy for weakly-supervised scene text grounding in space and time. Specifically, at the first stage, T2S-QA employs temporal grounding to distinguish positive frames (*i.e.*, frames with question-relevant scene texts) from negative frames (*i.e.*, frames without question-relevant scene texts) in a video. Given that motion blur and occlusion often obscure scene text in video frames, it subsequently refines the selection by identifying a few key positive frames that are most relevant to the question as the grounded frames. At the second stage, T2S-QA applies spatial grounding in each grounded frame to differentiate positive scene texts from negative ones. Since answers typically pertain to a small subset of the rich scene text present in a video frame, T2S-QA further selects partial positive scene texts covering the answer as the final grounding results. Finally, to decode the textual answers from the grounded scene texts, we feed the grounded scene texts, the grounded frames (serve as contexts), and the question into a transformer-based answer decoder for answer prediction.

As there is a lack of a benchmark for evaluating Grounded TextVideoQA, we construct the ViTXT-GQA dataset that considers both answer grounding and QA. ViTXT-GQA is built over the existing largest TextVideoQA dataset M4-ViteVQA [55] by extending its QAs in the validation and test sets with spatio-temporal labels (*i.e.*, timestamps and spatial bounding boxes). ViTXT-GQA consists of 2,055 QAs, 729 videos, 2,227 annotated temporal segments, and 52,494 bounding boxes. The labels are manually annotated and checked to be key for deriving accurate textual and visual answers.

Based on our ViTXT-GQA dataset, we source the baseline methods from three groups: TextVQA [20], TextVideoQA [55], and weakly-grounded VideoQA [14], [29]. For methods from image-based TextVQA and TextVideoQA, we conduct a post-hoc analysis of the attention distribution corresponding to the predictions. For methods designed for weakly-grounded VideoQA, we analyze their grounding outputs derived from the respective grounding modules. Our findings reveal that

existing models struggle to predict visually grounded answers, despite their relatively decent QA behavior. The results show our T2S-QA improves over these methods for both grounding and QA. Additionally, we evaluate two advanced multimodal large language models (MLLMs) on ViTXT-GQA: the open-source Qwen2-VL [44] and the closed-source GPT-4o-mini [34]. The experimental results show that T2S-QA has better scene text grounding capabilities than them, indicating the huge room for improvement of current MLLMs in scene-text grounding. Moreover, we find the performance gap between our T2S-QA and human is still large. For example, humans can correctly ground 77% of the answers to the questions, whereas T2S-QA only achieves 28%. In addition, we find that there is a clear discrepancy between grounding and QA performance for both human (77% vs. 64%) and T2S-QA (28% vs. 10%). On the one hand, the higher grounding and much lower QA accuracy of T2S-QA indicate that incorrect scene text recognition still seriously jeopardizes the final QA performance. On the other hand, the lower QA accuracy of human confirms the ineffectiveness of exact word matching for answer evaluation, and suggests the advantage and importance of our grounding evaluation to directly accept scene texts as visual answers. Finally, the large discrepancy between T2S-QA and human reveals that the poor TextVideoQA results of existing methods are largely attributed to the challenge posed by scene-text detection and recognition in the video domain.

Our main contributions are summarized as follows:

- We propose to study a new Grounded TextVideoQA task. It promotes research towards interpretable QA systems by not only providing visual evidence to answers, but also decoupling VideoQA from scene text recognition for better model analysis and development.
- We propose a cascaded temporal-to-spatial scene-text grounding and question-answering model, T2S-QA, with a disentangled temporal-spatial contrastive learning strategy for weakly-supervised scene-text grounding.
- We manually construct the ViTXT-GQA dataset to assess the models' capabilities of scene-text answer grounding.
- We thoroughly analyze the model behaviors, revealing significant challenges in performing Grounded TextVideoQA.

II. RELATED WORK

A. Text-Based Video Question Answering

TextVideoQA aims to answer questions by jointly reasoning about scene texts and visual contents in videos. Unlike TextVQA [20], [56], [18], [57], TextVideoQA [55], [41], [21], [58] focuses on developing the ability of machines to answer questions based on scene texts in dynamic environments. Zhao et al. [55] first identify this challenge and build a benchmark M4-ViteVQA, whose videos factor nine scenarios (*i.e.*, shopping, traveling, driving, vlog, sport, advertisement, movie, game, and talking), and develop a baseline model T5-ViteVQA. Tom et al. [41] propose the RoadTextVQA dataset on driving videos, focusing on questions that require understanding scene texts on road signs. Jahagirdar *et al.* [21] introduce NewsVideoQA to study TextVideoQA in news videos.

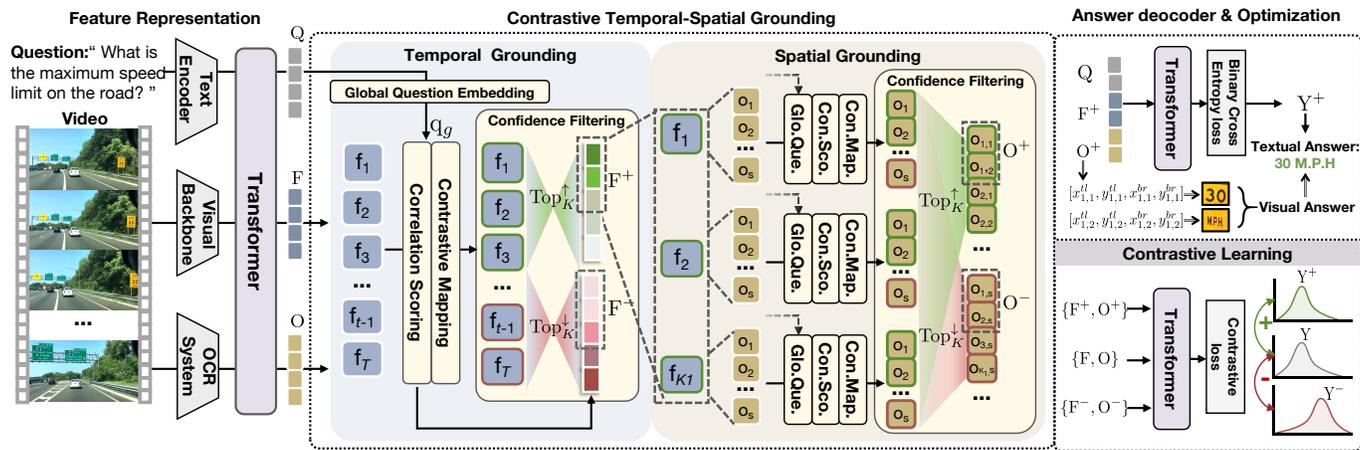


Fig. 2: Overview of our T2S-QA model. It mainly consists of three components: (1) the Feature Representation prepares the features of the question Q , video frames F , and OCR tokens O ; (2) the Contrastive Temporal-Spatial Grounding adopts a two-stage fine-grained grounding approach; (3) the Answer Decoder integrates the grounded frames F^+ , the grounded OCR tokens O^+ , and the question Q to achieve answer generation. In the Optimization process, we introduce a contrastive learning mechanism to help improve the question answering and answer grounding capabilities of our model.

The above studies have evaluated the QA ability of existing vision-and-language models on TextVideoQA, and shown that there is still large room for improvements. However, these works primarily emphasize the accuracy of the predicted textual answers, overlooking the research on answer interpretability. In this study, we propose Grounded TextVideoQA to explore both answer grounding and question answering. Moreover, we propose a grounding-then-answering method to benchmark the task and help analyze the interplay between scene-text recognition and VQA in videos.

B. Answer Grounding in VQA.

A key challenge faced by the VQA [38], [9], [35], [29] community is the issue of language bias [23], [40], [17], [32]. Specifically, models can answer questions by exploiting language priors in datasets without considering the visual content. To some extent, this problem is difficult to detect as commonly used evaluation metrics solely take account into textual answers to assess model performance. In image-based VQA tasks, many works [4], [5], [37], [13] propose new benchmarks for answer grounding to support models to localize visual evidence that humans rely on to answer visual questions. This is due to the observation that answer grounding can serve as a valuable foundation for debugging VQA models. However, related research is relatively scarce in VideoQA. TVQA [24] and TVQA+ [25] pioneer in study grounded VideoQA, yet in full supervision. NEX-T-GQA [49] firstly studies weakly-supervised grounding for VideoQA, yet focusing on temporal dimension. Moreover, all of them concentrate on grounding the normal visual objects or video moments. In contrast, we focus on spatio-temporally grounding the scene text regions, with our techniques also centering on weakly-supervised grounding.

C. Weakly-Supervised Video Grounding for VideoQA.

In VideoQA task [51], [48], [43], [7], several methods [36], [47], [49], [52], [14], [29] introduce weakly-supervised

video grounding techniques to improve the QA performance and model interpretability [59]. These methods locate critical frames [36], [27], [28], [52] and visual objects [14], [29], [50] to aid question answering. For example, [36] and [52] select the temporal proposal with the highest similarity to the language query. [49] design a Gaussian mask learning module to ground temporal segments for answer prediction.

Compared with VideoQA, TextVideoQA focuses on understanding local and even tiny scene texts in the video, where frame-level grounding struggles to capture subtle scene text details. In this work, we study spatio-temporal grounding [54] for the proposed Grounded TextVideoQA task. For video spatio-temporal grounding, MIST [14] designs an iterative attention module to select video segment and region features and fuse them for question answering. TranSTR [29] locates key video frames and objects and uses object-enhanced frame features for QA. However, these methods rely on the similarity between frame (object) features and questions to select critical elements, lacking effective self-supervision to constrain the selection during training. To address this, inspired by disentangled learning approaches [31] and contrastive learning strategy [30], [16], [15], we design a disentangled temporal and spatial contrastive learning strategy that distinguishes positive and negative video frames and scene text regions in the video and enlarges the semantic difference between positive and negative features. This approach injects strong self-supervision signals into the model to distinguish question-related video frames and scene texts, thereby effectively achieving scene text grounding. Our method distinguishes positive and negative video frames and scene text regions in the video and enlarges the semantic difference between positive and negative features. This approach injects strong self-supervision signals into the model to distinguish question-related video text frames and scene texts, thereby effectively achieving scene text grounding.

III. METHOD

Given a video and a question, Grounded TextVideoQA aims to encapsulate visual content, scene text, question semantics, and generate textual answers. Meanwhile, the scene text should be explicitly localized to substantiate the predicted textual answer or directly act as visual answer. To solve the task, localizing the question-related scene text is of crucial importance. As such, we propose a cascaded Temporal-to-Spatial Grounding and QA (T2S-QA) architecture as depicted in Fig. 2. T2S-QA comprises three core modules: 1) Feature Representation module extracts features for the question, video frames, and scene texts; 2) the Contrastive Temporal-Spatial Grounding (TSG) module executes temporal grounding to identify question-relevant video frames, followed by spatial grounding to pinpoint critical scene text regions within the identified frames; 3) the Answer Decoder module integrates the grounded video frames, scene text, and question to decode the final answer. For optimization, we curate positive and negative pairs to optimize the alignment between scene texts and the question spatio-temporally.

A. Feature Representation

Question features. Given a question, we extract its feature $Q = \{q_l\}_{l=1}^L \in \mathbb{R}^{L \times d}$ by a language backbone, where L is the question length, and d is the dimension of word embedding.

Video features. Given a video, we uniformly sample T frames $V = \{v_t\}_{t=1}^T$ and extract frame feature $F = \{f_t\}_{t=1}^T \in \mathbb{R}^{T \times d}$ to represent a video. For the t -th video frame, its features include two parts: 1) visual feature $f_t^v \in \mathbb{R}^{1024}$ extracted by a pre-trained visual backbone which is frozen during training; 2) temporal id feature $f_t^{te} \in \mathbb{R}^{50}$ where the temporal id t is assigned to denote the t -th video frame and encoded by an embedding layer as a temporal information. We merge two features to obtain the final video frame feature.

$$f_t = \text{LN}(W_1 f_t^v + W_2 f_t^{te}), \quad (1)$$

where $\text{LN}(\cdot)$ denotes layer normalization, W_1 and W_2 are learnable parameters.

OCR features. Given the sampled video frames $V = \{v_t\}_{t=1}^T$, following [55], we employ a scene text detection system [46] to obtain bounding boxes and track ids of the scene texts in each video frame, and then use a scene text recognition system [12] to read the scene texts according to their bounding boxes. For simplicity, the scene text detection system and recognition system are collectively referred to as optical character recognition (OCR) systems. Finally, the OCR token set for a video is denoted as:

$$\{\mathbf{B}_t, \text{Ta}_t\}_{t=1}^T = \text{Detector}(v_t), \quad (2)$$

$$\{r_{t,s}\}_{s=1}^S = \text{Recognizer}(v_t, B_t), \quad (3)$$

$$o_{t,s} = \{r_{t,s}, o_{t,s}^b, o_{t,s}^{ta}\}_{t=1,s=1}^{T,S}, \quad (4)$$

where $\text{Detector}(\cdot)$ is the scene text detection system applied to frame v_t . $\text{Recognizer}(\cdot)$ is the scene text recognition system applied to frame v_t and its bounding boxes \mathbf{B}_t , where $B_t = \{o_{t,1}^b, o_{t,2}^b, \dots, o_{t,S}^b\}$ is the set of bounding boxes in frame t and S is the OCR token number in each frame, where each

box contains the coordinates $[x_{t,s}^{tl}, y_{t,s}^{tl}, x_{t,s}^{br}, y_{t,s}^{br}]$ for the top-left and bottom-right corners. $\text{Ta}_t = \{o_{t,1}^{ta}, o_{t,2}^{ta}, \dots, o_{t,S}^{ta}\}$ is set of track ids for the detected text regions in frame t . $r_{t,s}$ is the recognized scene text content.

After that, we obtain OCR features $O = \{o_{t,s}\}_{t=1,s=1}^{T,S} \in \mathbb{R}^{T \times S \times d}$ in the video. For the s -th OCR token in the t -th frame, its features consist of five aspects: 1) FastText feature $o_{t,s}^f \in \mathbb{R}^{300}$ [3] providing essential word information based on the recognized scene text content $r_{t,s}$; 2) pyramidal histogram of characters (PHOC) feature $o_{t,s}^p \in \mathbb{R}^{604}$ [1] representing character features based on the recognized scene text content $r_{t,s}$; 3) normalized bounding box feature $o_{t,s}^b \in \mathbb{R}^4$ as the visual geometric features; 4) temporal id feature $o_{t,s}^{te} \in \mathbb{R}^{50}$ to provides temporal features consistent with the video frames; 5) track id feature $o_{t,s}^{ta} \in \mathbb{R}^{50}$ as the identification features. To project the representations into a common d -dimensional space, we apply a linear mapping on the above features respectively and sum them up to get the final representation of the OCR token features.

$$o_{t,s} = \text{LN}(W_3 o_{t,s}^f + W_4 o_{t,s}^p + W_5 o_{t,s}^{te} + W_6 o_{t,s}^{ta}) + \text{LN}(W_7 o_{t,s}^b), \quad (5)$$

where $\text{LN}(\cdot)$ denotes layer normalization, $W_3 \sim W_7$ are learnable parameters.

To reduce the semantic gap between different modalities, we apply a two-layer transformer Ψ [42] to establish associations between visual and textual content, aiding in the update of features. Here, we obtain the new question feature $Q = \{q_l\}_{l=1}^L$, frame feature $F = \{f_t\}_{t=1}^T$, and OCR token feature $O = \{o_{t,s}\}_{t=1,s=1}^{T,S}$.

$$[Q; F; O] = \Psi([Q; F; O]), \quad (6)$$

B. Contrastive Temporal-Spatial Grounding

Accurately localizing question-relevant scene text in a video is challenging, especially under weak supervision. To address this problem, we design an adaptive two-stage approach to localize key video frames and OCR tokens for answering questions: a Temporal Grounding (TG) stage that identifies the question-relevant video frames and a Spatial Grounding (SG) stage that pinpoints the critical OCR tokens.

1) Temporal Grounding: In most cases, scene text relevant to questions often appears in specific video frames as shown in Fig. 3. To align video content with question semantics and ground key video frames as context for effective question answering, we design a multi-step approach as follows.

Global Question Embedding. A global sentence representation, q_g , is derived from word embeddings q_l using self-attention, capturing overall question semantics by emphasizing keywords.

$$q_g = \sum_{l=1}^L \text{Softmax}(W_8 q_l) \cdot q_l, \quad (7)$$

Correlation Scoring. The correlation-based softmax operation computes the relevance distribution between the global question embedding and the video frame features. Guided by the contrastive learning objective in Eqn. (14), we calculate

the cross-attention scores for the positive frame distribution $f^+ \in \mathbb{R}^{T \times 1}$ and the negative frame distribution $f^- \in \mathbb{R}^{T \times 1}$, respectively. This approach explicitly distinguishes between positive (relevant) and negative (irrelevant) frames, ensuring that the positive frames align with the question semantics. The contrastive learning objective further amplifies the semantic differences between positive and negative frames, enhancing the capture of question-relevant video information.

$$f^+ = \text{Softmax}(W_9 F \cdot W_{10} Q_g^\top), \quad (8)$$

$$f^- = \text{Softmax}(W_{11} F \cdot W_{12} Q_g^\top), \quad (9)$$

Contrastive Mapping. To identify video frames that are relevant or irrelevant to the question, we apply the Gumbel-Softmax function to obtain the distribution of positive (f^+) and negative (f^-) frames. This operation enables a discrete distinction, ensuring that the positive frames align with the question semantics.

$$M_f = \text{Gumbel-Softmax}([f^+; f^-]), \quad (10)$$

where $[\cdot]$ denotes concatenation, where $W_8 \sim W_{12}$ are learnable parameters. The first and second columns of M_f index the positive frames (*i.e.*, M_{f_0}) and negative frames (*i.e.*, M_{f_1}).

Frame Filtering. Not all positive video frames carry clear and unobstructed scene text to answer the question. We need to further filter the positive and negative video frames based on their relevance to the given question. Specifically, we utilize Eqns (11) and (12) to enhance the semantic distinction between positive and negative frame features. Leveraging the cross-correlation scores computed by Eqns (8) and (9), we identify the top K_1 video frames from the positive (negative) frame set that exhibit the highest (lowest) relevance to the question. The obtained frame sets $\{F^+, F^-\}$ demonstrate significant semantic divergence, quantified by their relevance to the question. Guided by the contrastive learning objective, the model is optimized to further amplify this divergence, enabling grounding of question-relevant video frames F^+ .

$$F^+ = \text{Top}_K^\uparrow(f^+, M_{f_0} F) |_{K=K_1}, \quad (11)$$

$$F^- = \text{Top}_K^\downarrow(f^-, M_{f_1} F) |_{K=K_1}, \quad (12)$$

Up to now, we have introduced how to ground the frames F^+ that carry the key scene text for answering questions. These frames will be inputted into the answer decoder (Sec. III-C1) to serve as context for answer prediction.

2) **Spatial Grounding:** Our preliminary analysis reveals that most answers only occupy a very small area in the video frame, as shown in Fig. 3. Therefore, we first need to distinguish the scene texts related to the question from numerous irrelevant visual elements, and then locate a few critical strongly associated scene texts as visual answers. To achieve that, SG facilitates adaptive OCR token selection similar to Eqn. (7) \sim (12). Concretely, for the i -th critical frame, SG is fed with S object features detected on the frame, enabling it to differentiate positive OCR tokens from negative ones based on their similarity to the given question. For positive OCR tokens, we select K_2 OCR tokens in each frame with the *highest* attention weight from O^+ and gather their corresponding OCR features as $O^+ \in \mathbb{R}^{K_1 \times K_2 \times d}$. For

negative OCR tokens, we select K_2 OCR tokens in each frame with the *lowest* attention weight from O^- and gather their corresponding OCR features as $O^- \in \mathbb{R}^{K_1 \times K_2 \times d}$. It is worth noting that the detected OCR results vary for each frame, and SG is applied independently to each frame. Hence, even if we use hyperparameter K_2 to select OCR tokens in each frame, the numbers of critical OCR tokens in each frame are different.

Up to now, we have obtained the pair of the K_1 positive video frames and $K_1 \times K_2$ positive OCR tokens, namely $\{F^+, O^+\}$, and the pair of the K_1 negative video frames and $K_1 \times K_2$ negative OCR tokens, namely $\{F^-, O^-\}$. Notably, we take the bounding boxes of O^+ as grounded OCR tokens for answer grounding evaluation, the evaluation details are elaborated in Sec. V-A.

C. Answer Decoder and Optimization

To utilize the grounded video frames and OCR tokens to generate answers, we adopt a contrastive learning mechanism as an auxiliary objective to enhance the grounding accuracy. Specifically, we obtain the positive frame-OCR token pair $\{F^+, O^+\}$ and the negative frame-OCR token pair $\{F^-, O^-\}$ in TSG, and take raw video frames and OCR tokens $\{F, O\}$ as anchor. We next feed these pair instances into the answer decoder to obtain three answer predictions, *i.e.*, anchor Y and its contrastive counterparts Y^+ and Y^- . We elaborate on the answer generation and optimization process below.

1) **Answer Decoder:** Following [20], the answer decoder is composed of a three-layer transformer $\hat{\Psi}$ and two classifiers — a vocabulary classifier ψ_{voc} and an OCR token classifier ψ_{ocr} . Take the example of the answer Y^+ generated by the positive frame-OCR token pair $\{F^+, O^+\}$, we concatenate the question Q , the positive frames F^+ and the positive OCR tokens O^+ , and a hidden state D and input them into the transformer block as follows:

$$[Q_{t+1}; F_{t+1}^+; O_{t+1}^+; D_{t+1}] = \hat{\Psi}([W_{13} Q_t; W_{14} F_t^+; W_{15} O_t^+; W_{16} D_t]), \quad (13)$$

where $W_{13} \sim W_{16}$ are learnable parameters, the hidden state D_0 is initialized by the positional embedding [20], and where the initialization of $F_0^+ = F^+$, $O_0^+ = O^+$, and $Q_0 = Q$. By iteratively performing the transformer L_a times, we obtain $D = [d_1, \dots, d_{L_a}] \in \mathbb{R}^{L_a \times d}$. At the t -th decoding step, the classifier ψ_{voc} is optimized to predict the probability score y_t^{voc} over a preset vocabulary. The classifier ψ_{ocr} calculates the relevance score y_t^{ocr} of D_t and the OCR token O^+ . We implement the Argmax function on y_t^{voc} and y_t^{ocr} to predict the word y_t . Thus, the generated answer is denoted as $Y^+ = \{y_1, \dots, y_{L_a}\}$. The generation of the negative answer Y^- and the anchor answer Y is the same as above.

2) **Model Optimization:** By far, we obtain the three answers Y , Y^+ , and Y^- based on the anchor, positive frame-OCR token pair, and negative frame-OCR token pair. Next, we optimize the model from two aspects: 1) we design a contrastive loss to improve the grounding accuracy of our model; 2) we adopt a cross-entropy loss to help the model utilize the localized features to answer questions.

Contrastive loss. To enhance the grounding accuracy of the Grounded TextVideoQA model, we exploit the causal

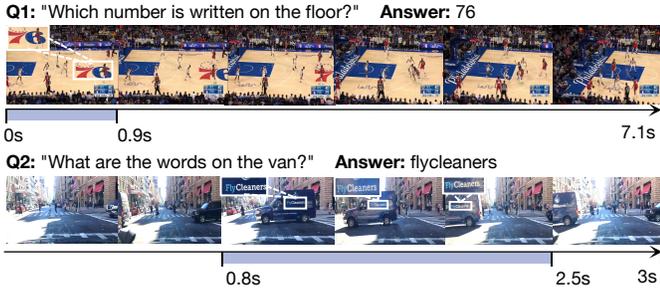


Fig. 3: Examples of spatial-temporal labels in ViTXT-GQA.

TABLE I: Statistics of ViTXT-GQA dataset.

Split	Vid.	Que.	Tem. Span	Box Num.	Ave. Box
Train	5,444	18,220	-	-	-
Val	343	981	1,087	25,445	25.88
Test	386	1,074	1,140	27,049	25.19

invariance between the original video and the positive video frames. The contrastive learning loss is designed to amplify the semantic differences between the positive frames (OCR tokens) and the negative frames (OCR tokens). We construct a contrastive objective referring to InfoNCE [33] as follows:

$$\mathcal{L}_{cons} = -\log\left(\frac{\exp(\text{sim}(Y^+, Y^T)/\tau)}{\exp(\text{sim}(Y^+, Y^T)/\tau) + \exp(\text{sim}(Y^-, Y^T)/\tau)}\right), \quad (14)$$

where $\text{sim}(\cdot)$ denotes cosine similarity function, and we empirically set temperature coefficient $\tau = 0.1$ following [6], [26]. **Binary cross-entropy loss.** For TextVideoQA, the answer prediction can be regarded as a multi-label classification problem with the classifier ψ_{voc} and the classifier ψ_{ocr} . Binary cross-entropy loss function is suitable for multi-label classification [20], [55]. Considering that the grounded positive video frames and OCR tokens are beneficial for question answering, we utilize this loss function to minimize the semantic distance between the positive answer and the ground truth answer.

$$\mathcal{L}_{bce} = -Y \log(\sigma(Y^+)) - (1 - Y) \log(1 - \sigma(Y^+)), \quad (15)$$

where $\sigma(\cdot)$ is sigmoid function, Y is the ground truth answer. The overall training objective of T2S-QA is as follows:

$$\mathcal{L} = \mathcal{L}_{bce} + \lambda \mathcal{L}_{cons}, \quad (16)$$

and λ is a trade-off hyperparameter.

IV. DATASET

A. ViTXT-GQA Construction

Data Source. We choose M4-ViteVQA [55] dataset as our data source to augment with spatio-temporal labels. Compared with others (Table II), M4-ViteVQA has a richer set of video scenarios, such as shopping, traveling, driving, vlog, sport, advertisement, movie, game, and talking. Specifically, we choose its hard data split where video clips do not overlap among the train/val/test set. The data split consists of 5,444/525/680 videos and 18,220/1,620/2,103 questions on train/val/test sets. **Dataset Filtering.** We only label the validation and test sets for a weakly-supervised setting. To ensure the clarity of answer

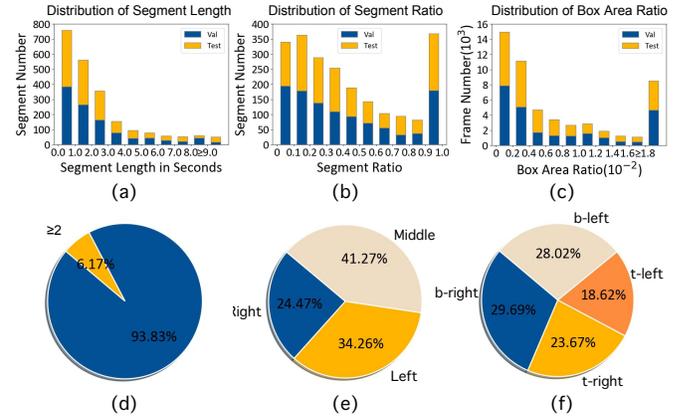


Fig. 4: Analysis of annotated spatio-temporal labels.

grounding, the questions whose answers do not involve scene texts or involve multiple text regions on a single frame are removed before the annotation. Moreover, the questions whose answer texts are blurry or occluded in the video frame are also removed. Consequently, 343/386 videos and 981/1,074 questions on val/test sets are to be annotated.

Label Collection. We use Vidat [53] as our video annotation tool, and invite 15 undergraduate students for annotation. The annotators are instructed to first comprehend the questions and video contents, and then annotate scene texts pertaining to the answers in the video with bounding boxes. Noteworthy, all frames that carry the answer should be annotated for a given question. Moreover, if an answer comprises multiple words, we ask for a single bounding box that can encompass all answer words. Yet, our preliminary statistics show that the majority of answers include only a single word. Note that there is at most one annotated box on a video frame. Fig. 3 illustrates two examples from the ViTXT-GQA dataset, where we need to annotate the temporal segment timestamps for the appearance of the scene text corresponding to the answer for each question. Besides, we annotate a bounding box on each frame, even if the answer consists of multiple words. To ensure quality and reduce subjectivity, each QA is reviewed by at least two people. The annotation process took about two months.

B. Dataset Analysis

The statistics of ViTXT-GQA are shown in Table I and Fig. 4. In Table I, we annotate 1,087 and 1,140 temporal spans on the validation and the test sets, including 25,445 and 27,049 box annotations. The average box annotation per video is about 25. Fig. 4 (a) and (b) show the temporal ratio distributions of the annotated temporal segments. In 722/734 questions of the validation/test sets, the annotated temporal segment ratio in the video is less than 50%; in 123/118 questions of the validation/test sets, the annotated segment ratio is greater than 98%. In the validation set and test set, the average ratio of annotated segments is 41.75% and 44.33% respectively. Fig. 4 (d) shows that a significant 93.83% of the questions have one temporal segment. Fig. 4 (e) reflects the spatial distribution of the annotated bounding boxes within the temporal segments. The annotated segments are evenly distributed across the left,

TABLE II: Benchmark comparison. ST: Scene Text. TG: Temporal Grounding. SG: Spatial Grounding. Acc: Accuracy. IoU: Intersection over Union. GQA: Accuracy & IoU.

Datasets	ST	TG	SG	Evaluation Metrics
TVQA [24]	×	✓	×	Acc, IoU
TVQA+ [25]	×	✓	✓	Acc, IoU
NEXt-GQA [49]	×	✓	×	Acc, IoP, GQA
NewsVideoQA [21]	✓	×	×	Acc, ANLS
RoadTextVQA [41]	✓	×	×	Acc, ANLS
M4-ViteVQA [55]	✓	×	×	Acc, ANLS
ViTXT-GQA	✓	✓	✓	Acc, ANLS, IoU, GQA

middle, and right sections of the video, with 41.27% of the bounding boxes distributed in the middle portion of the video. We also examine the spatial distribution characteristics of the bounding boxes. Fig. 4 (f) shows the spatial distribution of the bounding boxes within the image. It can be observed that the boxes are uniformly distributed in the top-left, bottom-left, top-right, and bottom-right corners of the video frames, with a slightly higher occurrence of boxes in the bottom-right corner, reaching 29.69%. Fig. 4 (c) presents the statistics of the ratio between the annotated bounding box area and the area of the corresponding video frame. It is found that the majority of the boxes have an area ratio smaller than 0.2%, which reflects the challenge of locating scene text in videos for QA.

C. Dataset Comparison

To highlight the uniqueness of ViTXT-GQA, we compare it with TextVideoQA benchmarks and visually grounded VideoQA benchmarks in Tab. II.

Comparison with TextVideoQA. NewsVideoQA [21], RoadTextVQA [41], and M4-ViteVQA [55] aim to predict textual answers. ViTXT-GQA differs in two main aspects: (1) it provides visual evidence to support the answers, and (2) it extends the VQA setting by allowing visual answers. This satisfies more practical applications and facilitates better diagnostics of model performance. For example, are the incorrect answer predictions due to scene text recognition errors or overly strict answer evaluation? ViTXT-GQA is also more challenging because: (1) models need to achieve multiple objectives (*i.e.*, grounding and QA) while maintaining consistency, and (2) models must accomplish not only temporal grounding but also spatial grounding, which is challenging to locate scene text that answers the question in visually rich videos.

Comparison with VideoQA. The fully supervised benchmark TVQA [24] and TVQA+ [25] respectively provides temporal annotations and spatio-temporal annotations in the training data, where labels can resolve referential ambiguities in questions or improve QA performance through well-localized visual content. ViTXT-GQA differs from them by focusing on recognizing and grounding *scene text* as visual evidence for QA. Additionally, we place emphasis on weakly supervised spatio-temporal grounding to better investigate model generalization. Compared with NEXt-GQA [49], which focuses on weakly supervised temporal grounding, we consider fine-grained spatio-temporal grounding and additionally incorporate scene text understanding, making it more challenging and offering broader applications.

You are an expert in analyzing video data, and you should answer a question based on the given video. Please briefly answer the question with a single word, a phrase or a short sentence. Question: <question>

(a) Question Answering Prompt

You are an expert in video data analysis tasked with answering questions based on a given video. Please briefly answer the question with a single word, a phrase or a short sentence.

To guide your reasoning, please follow these steps:

1. Identify the video frame most relevant to the question.
2. Within this frame, focus specifically on the scene text, as the answer is derived from this text. Locate the scene text regions most relevant to the question.
3. Using the identified frame and relevant scene text regions, answer the question clearly and concisely in a single word, a phrase, or a short sentence.

For example, when answering the question 'What number is written on the yellow wall?', you need to follow these steps:

1. Find the video frame where the yellow wall appears;
2. Find the scene text information in the yellow wall area on the video frame;
3. Answer the question based on the scene text.

Now, you need to answer the question: Question: <question>

(b) Question Answering with CoT Prompt

Step-1: Temporal Grounding

Input: You are an expert in video analysis. Analyze the following eight video frames to determine which frame contains the most relevant information to answer the given question: '{question}'. Respond with a single integer indicating the sequence number of the relevant frame, e.g., 3.

Output: frame id

Step-2: Spatial Grounding

Input: Given the following image and a question: '{question}'. The answers consists of the scene texts in the image. Please find the bounding box containing the text directly relevant to answering this question (e.g., a price, a name, or a date). The bounding box should precisely surround the relevant text area. Please provide the coordinates of the bounding box in the following format: <[box_start]>(top-left x, top-left y, bottom-right x, bottom-right y)<[box_end]>. Output only the coordinates in this exact format.

Output: <bbox>

Step-3: Question Answering

Input: You are analyzing an image and a bounding box on the image. The question is: '{question}'. The bounding box coordinates are: '{bbox}'. Based on the content within the bounding box on the image, answer the question concisely using a single word, a phrase, or a short sentence.

Output: <answer>

(c) Question Answering with Grounded CoT Prompt

Fig. 5: Prompts for MLLMs to perform ViTXT-GQA.

V. EXPERIMENTS

In this section, we show the experimental results to answer three research questions.

- **Q1:** How effective is our T2S-QA model?
- **Q2:** What is the contribution of each component in T2S-QA?
- **Q3:** How does scene-text recognition affect TextVideoQA?

A. Settings

Implementation Details. We adopt a pre-trained BERT-Base [8] model to encode the question. The encoder is also fine-tuned. For video embedding, each frame is encoded by a ViT-L [10] model pre-trained on ImageNet-21k. We standardize the video frame rate to 10 frames per second and set frame number $T = 64$, OCR token number $S = 15$ in each frame, the maximal question length $L = 20$, and the common hidden dimension $d = 768$. For answer decoder and model optimization, we set the step number of answer decoding to $L_a = 12$ and the trade-off hyperparameter $\lambda = 100$. We collect the top 5,000 frequent words from the answers in the training set as a fixed vocabulary. During training, we set batch size to 48 and train

TABLE III: Performance comparison on ViTXT-GQA test set. Models use ABINet [12] by default for scene text recognition. We additionally upgrade to CLIPOCR [45]. The best and second-best results are **bolded** and underlined respectively.

Method	Top 1 × 1						Top 5 × 5					
	IoU@0.5	IoU@0.3	Acc	ANLS	GQA@0.5	GQA@0.3	IoU@0.5	IoU@0.3	Acc	ANLS	GQA@0.5	GQA@0.3
M4C [20]	2.45	4.41	9.66	0.1362	0.64	0.70	9.13	14.35	9.87	0.1344	2.29	2.56
T5-ViteVQA [55]	1.30	2.15	9.45	0.1381	0.36	0.45	16.82	24.81	9.45	0.1381	1.81	2.03
MIST [14]	1.80	2.94	8.47	0.1320	0.38	0.42	10.69	16.78	8.53	0.1316	1.72	1.18
TranSTR [29]	2.72	4.26	9.45	0.1302	0.45	0.45	14.94	23.01	9.27	0.1430	2.63	3.05
T2S-QA (ABINet)	<u>2.99</u>	<u>4.89</u>	<u>9.72</u>	<u>0.1445</u>	<u>0.72</u>	<u>0.72</u>	<u>28.14</u>	<u>41.62</u>	<u>10.20</u>	<u>0.1472</u>	<u>2.69</u>	<u>3.62</u>
T2S-QA (CLIPOCR)	3.21	4.95	9.87	0.1477	0.88	0.88	28.60	41.94	10.45	0.1520	3.50	4.23

TABLE IV: Performance of MLLMs (zero-shot) on ViTXT-GQA. The evaluation is conducted under the Top 1 × 1 setting.

Method	IoU@0.5	Acc	ANLS	GQA@0.5
<i>Question Answering</i>				
GPT-4o-mini w/ Standard	-	34.19	0.5070	-
Qwen2-VL w/ Standard	-	<u>41.34</u>	<u>0.5649</u>	-
<i>Question Answering with CoT Prompt</i>				
GPT-4o-mini w/ CoT	-	37.99	0.5574	-
Qwen2-VL w/ CoT	-	41.90	0.5663	-
<i>Question Answering with Grounded CoT Prompt</i>				
GPT-4o-mini w/ Grounded CoT	0.19	34.46	0.5058	0.00
Qwen2-VL w/ Grounded CoT	2.14	26.07	0.3628	<u>1.02</u>
Qwen2-VL w/ One-shot Grounded CoT	<u>2.94</u>	26.98	0.3664	1.32
T2S-QA (CLIPOCR)	3.21	9.87	0.1477	0.88

for a maximum of 24,000 iterations. Our model is trained using the Adam optimizer, with a learning rate of 1e-4 and a staircase learning rate schedule, where we multiply the learning rate by 0.1 at 10,000 and at 20,000 iterations. The best checkpoint is selected using the validation set accuracy.

Evaluation. We consider QA and answer grounding evaluation. For QA, we adopt Accuracy (**Acc**) in percentage $\in [0, 100]$ and Average Normalized Levenshtein Similarity (**ANLS**) [2] $\in [0, 1]$. Acc considers a predicted answer to be correct if and only if it matches the ground-truth answer word by word. ANLS measures the character-level compositional similarity between the predicted answer and the ground truth. For answer grounding, we adopt Intersection over Union (**IoU**) between the grounded and the annotated OCR token boxes. We consider IoU at two thresholds of 0.3 and 0.5 for better analysis. To evaluate joint QA and answer grounding, we consider grounded QA accuracy (**GQA**) [49]. Specifically, we report GQA@0.5 and GQA@0.3 metrics which stand for the percentages of correctly answered questions with grounding IoU ≥ 0.5 and 0.3 respectively. Finally, we set two evaluations considering the severe challenge of video scene-text grounding. The first is a normal evaluation of the Top 1 × 1, which measures the Top-1 grounded scene text region within the Top-1 temporally localized video frame. The second is Top 5 × 5, which relaxes the requirements from Top-1 to Top-5 for both spatial and temporal grounding, *i.e.*, measuring if the identified Top-5 frames and the Top-5 regions within each frame contains the ground-truth scene-text region.

B. Main Comparison (RQ1)

As there is no existing method for Grounded TextVideoQA. We adapt several approaches from TextVQA, TextVideoQA, and VideoQA as baselines in Tab. III. For a fair comparison, all methods in Tab. III share feature encoding and answer decod-

ing backbones. Besides, we evaluate the zero-shot performance of the current advanced MLLMs Qwen2-VL [44] and GPT-4o-mini [34] in Tab. IV. We elaborate on these model details as follows.

M4C [20] is an image-based TextVQA model. We adapt it to the video domain by randomly sampling a single frame. In practice, we sample the first, middle, and end frames, and average their results for reference. To obtain spatial grounding results, we employ post-hoc attention analysis as described in Eqn. 7 ~ 8. We select Top K_2 OCR tokens according to the attention distribution of Eqn. 8. M4C utilizes the sampled frame features and OCR token features for the QA prediction.

T5-ViteVQA [55] is a TextVideoQA model. To obtain its spatio-temporal grounding results, we refer to the same method as above for post-hoc attention analysis. We select Top $K_1 \times K_2$ OCR tokens from all OCR tokens of the video according to the attention distribution. Similarly, the result of post-hoc attention does not affect QA prediction; T5-ViteVQA uses all frame features and OCR token features for QA prediction.

MIST [14] is a VideoQA model. Since general VideoQA models struggle to understand scene text in videos, we extend MIST to the Grounded TextVideoQA task by replacing its original answer decoder with the answer decoder proposed in this paper. For spatio-temporal grounding, MIST first divides the video into segments of equal length, then introduces a cascaded segment and region selection module to adaptively select a question-related video segment, and further selects Top K_1 frames and Top K_2 OCR tokens. In implementation, we use the grounded frames and OCR tokens to perform QA.

TranSTR [29] is a VideoQA model. We also apply our answer decoder to TranSTR to extend it to the Grounded TextVideoQA. For spatio-temporal grounding, TranSTR designs a differentiable selection module that adaptively collects Top K_1 critical frames and Top K_2 OCR tokens in each frame as grounded results, and finally uses OCR-enhanced frame features for QA prediction.

Qwen2-VL [44] is a cutting-edge open-source MLLM that supports video inputs, excelling in VQA and visual grounding tasks. By integrating temporal and spatial information, it delivers strong performance in complex spatio-temporal reasoning and grounding scenarios. To explore the performance of Qwen2-VL on the Grounded TextVideoQA task, we design a dedicated Chain-of-Thought prompt (CoT) (Fig. 5) to explore its QA and grounding capabilities. The CoT prompt guides the model in identifying the most relevant video frame and scene text bounding box for answering the question.

GPT-4o-mini [34] is a closed-source, lightweight variant of

TABLE V: Ablation study of temporal-spatial grounding.

Variant	IoU@0.5	IoU@0.3	Acc	ANLS	GQA@0.5	GQA@0.3
T2S-QA	28.14	41.62	10.20	0.1472	2.96	3.62
w/o TSG	-	-	8.64	0.1380	-	-
w/o SG	3.35	5.77	8.94	0.1351	0.54	0.63
w/o TG	14.85	25.55	9.57	0.1430	1.96	3.05

TABLE VI: Ablation study of the input of answer decoder.

Variant	IoU@0.5	IoU@0.3	Acc	ANLS	GQA@0.5	GQA@0.3
T2S-QA	28.14	41.62	10.20	0.1472	2.96	3.62
w/o O	16.98	25.69	5.44	0.0803	1.27	0.63
w/o F	25.68	37.08	9.72	0.1391	2.05	2.87
w/o F & Q	23.93	34.93	8.57	0.1179	1.87	2.69

the GPT-4 family, designed for efficiency while maintaining robust multimodal reasoning capabilities. Although it performs well in VQA, its visual grounding ability is constrained by weaker spatio-temporal reasoning. As with Qwen2-VL, using the designed CoT prompt in Fig. 5, we also evaluate GPT-4o-mini's performance on the Grounded TextVideoQA task.

The main results in Tab. III show that our model T2S-QA consistently outperforms all baselines under different evaluation metrics and settings. Such strength is unique to T2S-QA compared with other baseline methods. More specifically, T2S-QA wins more under the Top-5 setting versus the Top-1, especially for grounding. Examining the performance of all other methods under the Top-1 setting, we find that accurately localizing the precise scene-text region corresponding to a specific answer is highly challenging. This also explains the relatively marginal performance gap among models.

In Tab. IV, we evaluate the performance of MLLMs on ViTXT-GQA through three setups: direct QA ("w/ Standard") in Fig. 5 (a), QA with Chain-of-Thought (CoT) prompt ("w/ CoT") in Fig. 5 (b), and scene-text grounded QA with Chain-of-Thought prompt ("w/ Grounded CoT") in Fig. 5 (c). Our findings are as follows: (1) *MLLMs face severe challenges in scene-text spatio-temporal grounding despite their good QA performance.* The results show that under the "Grounded CoT" strategy, GPT-4o-mini [34] and Qwen2-VL [44] achieve low IoU@0.5 scores of 0.19% and 2.14%, respectively, failing behind T2S-QA. (2) *CoT prompting improves QA accuracy.* Switching from "w/ Standard" to "w/ CoT", GPT-4o-mini [34] shows a 3.80% increase in Acc, demonstrating CoT's effectiveness in enhancing reasoning. To further explore the CoT prompt, we adopt a representative QA pair as the one-shot CoT prompt in Fig. 5 (c). The results demonstrate a modest improvement in both textual answer accuracy and visual grounding performance. (3) *The performance of explicit grounded reasoning and implicit grounded reasoning is inconsistent.* When transitioning from "w/ Standard" to "w/ Grounded CoT", Qwen2-VL [44] suffers a 15.27% decrease in Acc, indicating the inability to effectively integrate grounding in the QA process. We further identify key QA failure cases, including reliance on irrelevant scene text or hallucinated visual content and difficulty recognizing occluded or partially visible text, underscoring the challenges and future directions in scene text-aware video QA.

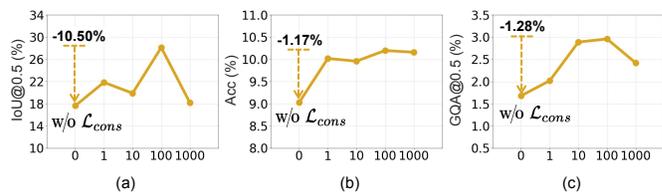


Fig. 6: Study of contrastive loss \mathcal{L}_{cons} with λ .

TABLE VII: Ablation study of different scene text features. *Linguistic* features include FastText feature and PHOC feature. *Spatial* features include bounding box features, *Temporal* features include temporal id features and track id features.

Linguistic	Spatial	Temporal	IoU@0.5	Acc	ANLS	GQA@0.5
✓	-	-	24.64	10.30	0.1414	3.05
✓	✓	-	28.31	10.39	0.1471	2.60
✓	✓	✓	28.60	10.45	0.1520	3.50

C. Ablation Study (RQ2)

To understand the effects of different components in T2S-QA, we dissect the model performance on ViTXT-GQA test set and report results with the Top 5x5 setting.

Effectiveness of Temporal-to-Spatial Grounding. In Tab. V, we respectively study the effectiveness of T2S-QA by removing temporal and spatial grounding ("w/o TSG"), only removing temporal grounding ("w/o TG"), and only removing spatial grounding ("w/o SG"). Specifically, "w/o TG" variant uses the grounded OCR tokens in all video frames as positive OCR tokens and the others as negative OCR tokens; "w/o SG" variant uses all OCR tokens of the positive video frame as the positive OCR token, and all OCR tokens of the negative video frame as the negative OCR token; "w/o TSG" variant represents reserving all video frames and OCR tokens for answering, *i.e.*, without answer grounding.

Tab. V shows that both spatial and temporal grounding are key for better QA performance. We note that the model variant "w/o SG" performs the worst in both answer grounding and grounded QA. This suggests that only learning the differences between video frames is insufficient to capture subtle local changes in scene text within a frame, making it challenging to achieve accurate answer grounding. Moreover, compared with T2S-QA, "w/o TG" also shows a significant decrease of 13.29% on IoU@0.5 and 1% decrease in GQA@0.5. This indicates that directly performing one-stage grounding of scene text across the entire video yields suboptimal results. **Study Inputs of Answer Decoder.** Tab. VI shows that when removing OCR tokens ("w/o O") both T2S-QA's grounding and QA results degenerate significantly by 11.16% on IoU@0.5 and 4.76% in Acc respectively. This demonstrates the significance of OCR token grounding for answer prediction. Additionally, we show that both the frame appearance F feature and question Q help the performance.

Effectiveness of Contrastive Loss. Fig. 6 shows that T2S-QA's performances degenerate remarkably without the contrastive loss, especially the grounding result degrades by over 10%. We also investigate different values for the coefficients λ . The results suggest that $\lambda = 100$ brings the best performance.

TABLE VIII: Performance with different numbers of video frames. Box Cov. (%) is the average coverage percentage of annotated boxes w.r.t the sampled video frames.

T	Box Cov.	IoU@0.5	IoU@0.3	Acc	ANLS	GQA@0.5	GQA@0.3
32	68.75	18.79	27.95	9.55	0.1418	1.87	2.68
48	82.19	18.84	27.55	9.66	0.1413	1.96	2.87
64	92.25	28.14	41.62	10.20	0.1472	2.96	3.62
80	96.62	18.68	27.09	9.93	0.1413	2.51	2.92

TABLE IX: Performance with different numbers of transformer layers.

Variant	IoU@0.5	IoU@0.3	Acc	ANLS	GQA@0.5	GQA@0.3
w/o Trans. Layer	16.36	25.72	9.18	0.1272	2.14	2.42
w/ 1-layer Trans.	25.08	38.18	8.70	0.1262	2.87	4.23
w/ 2-layer Trans.	28.60	41.94	10.45	0.1520	3.50	4.23
w/ 3-layer Trans.	27.00	39.53	8.97	0.1294	3.17	3.89

Study of Scene Text Features. In Tab. VII, the results show that linguistic features obtained by character-level encoding are crucial based on the scene text recognition results. Capturing spatial and temporal features can complement the spatiotemporal characteristics of scene text in videos. Their combination produces the best performance.

Different Numbers of Video Frames. We conduct experiments with four different numbers of sampled video frames, *i.e.*, $T = 32, 48, 64,$ and 80 . Tab. VIII shows that uniformly sampling 80 frames can cover over 96% of the annotated scene-text regions. However, the best results are achieved at frame number of 64. We speculate that learning with too many frames make it harder to identify the very limited positive frames that carry the needed scene texts.

Different Numbers of Transformer Layers. In Tab. IX, we conduct experiments by varying the number of Transformer layers for the architecture introduced at Sec. III-A. It shows that when the Transformer module is removed (“w/o Trans. Layer”), the performance on IoU@0.5 drops by 12.24%, and the performance on GQA@0.5 drops by 1.03%. The results show that the optimal configuration is with two layers.

D. In-Depth Analysis (RQ3)

Our previous experiments show that models obtain low QA accuracy even with relatively higher grounding behavior. We speculate that the reasons are two folds: *ineffective answer evaluation* and *poor scene-text detection and recognition*. The former indicates that the predicted textual answer does not exactly match the ground-truth one (*e.g.* 30 vs. 30 M.P.H in Fig. 1). The latter hints that the OCR system fails to convert the grounded scene-text into textual answers (*e.g.* M vs. M.P.H).

Human Performance and Upper Bound. To discern the error from ineffective answer evaluation, we sample around 20% test data which contains 200 questions and 163 videos and conduct a human study by asking volunteers to answer the questions and ground the related scene text areas. Table X shows that even human (we assume that humans are with perfect scene-text recognition) can only achieve an accuracy of 64.32%. A further investigation of the the wrong answers suggests that an absolute 25% out of 36% of they are acceptable and should be regarded as correct. For example, “35131.86”

TABLE X: Study of scene-text grounding and recognition in Grounded TextVideoQA.

Method	IoU@0.5	IoU@0.3	Acc	ANLS	GQA@0.5	GQA@0.3
Human	77.39	80.90	64.32	0.8012	51.76	54.27
Upper Bound	56.52	76.07	72.62	0.6967	41.53	56.15
T2S-QA (CLIPOCR)	28.60	41.94	10.45	0.1520	3.50	4.23
w/ CLIPOCR GT Box	96.70	96.81	11.96	0.1875	11.24	11.56

TABLE XI: Performance of T2S-QA under different scene-text detection and recognition systems.

OCR Detection	OCR Recognition	IoU@0.5	Acc	ANLS	GQA@0.5
GoMatching [19]	ABINet [12]	25.26	9.89	0.1436	2.15
TransVTSpotter [46]	ABINet [12]	28.14	10.20	0.1472	2.69
TransVTSpotter [46]	CLIPOCR [45]	28.60	10.45	0.1520	3.50

vs. “35,131.86”. The results suggest that stringent string matching is ineffective in assessing TextVQA models. Yet, human’s significantly higher grounding performance indicates that grounding annotation is *less subjective*, highlighting the importance of directly evaluating scene-text regions. The gap between our model and human performance underscores the need for further research into trustworthy TextVideoQA.

We report the “Upper Bound” result with the TransVTSpotter [46] detection and CLIPOCR [45] recognition systems, obtained via brute-force search over all detected text regions and recognized outputs. The result in Table X suggests that even a brute-force search can only achieve a grounding accuracy of 56.52%. Note that T2S-QA achieves 28.60%. This reflects the effectiveness of our grounding approach, though the accuracy is not so high. Additionally, we find that the “Upper Bound” is inferior to human in grounding but exceeds human in QA accuracy. This reveals the inadequacy of the existing OCR system in scene-text detection and the challenge of effectively evaluating textual answers because of human subjectivity.

Impact of Scene Text Detection and Recognition. As shown in Table XI, we evaluate various scene text detection and recognition systems and analyze their impact on the Grounded TextVideoQA task from three perspectives. 1) To assess the impact of OCR detection systems, we test the variant where the ground-truth scene text for each QA pair is directly provided to the answer decoder. Compared to the human-annotated answer grounding result (“w/ CLIPOCR GT Box”), the performance of directly iterating over all OCR detection results (“Upper Bound”) drops by 40.18% on IoU@0.5. This drop illustrates the difficulty of scene text detection in dynamic video environments. 2) We analyze the impact of OCR recognition systems. The GQA@0.5 score of “T2S-QA (CLIPOCR)” is 25.1% which is lower than its IoU@0.5 score, suggesting incorrect textual answers even when the scene text is located accurately. The comparison between “T2S-QA (CLIPOCR)” and “w/ CLIPOCR GT Box” shows only marginal improvement (1.51% in Acc) despite perfect grounding. These results emphasize the limitations of current video-based scene text recognition systems. 3) We evaluate various combinations of OCR detection and recognition systems. Using the same recognition system, TransVTSpotter achieves superior grounding due to its robust detection capabilities. Similarly, employing CLIPOCR as the recognition method significantly

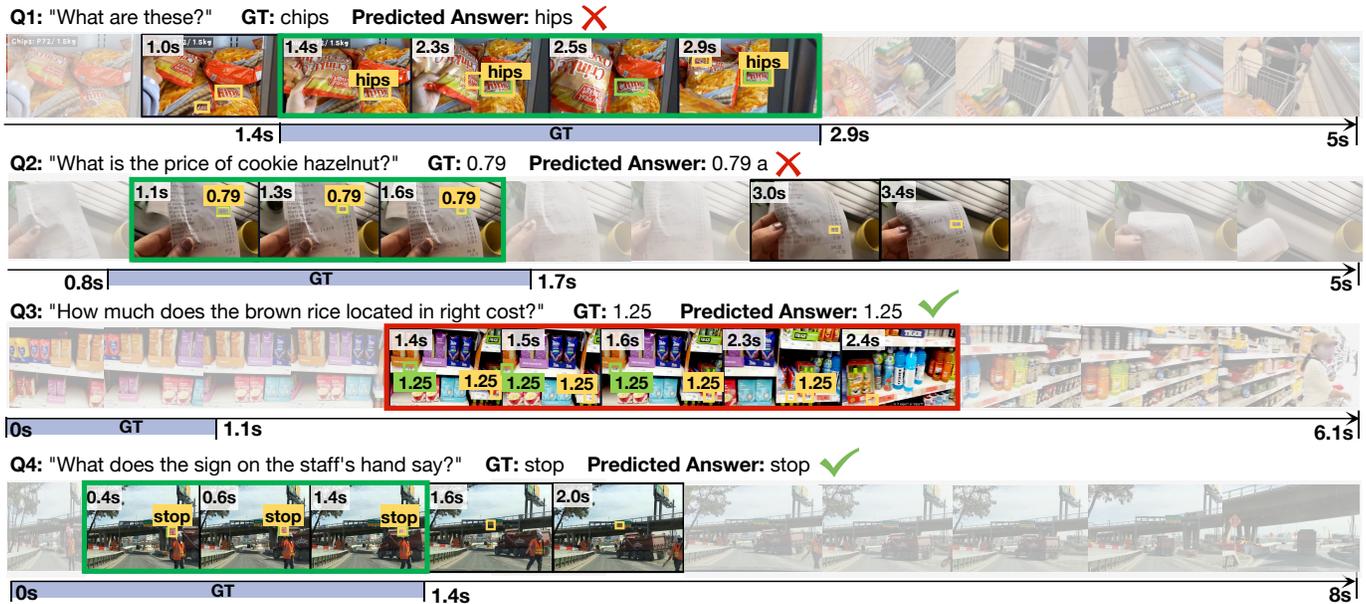


Fig. 7: Case-Study of T2S-QA on ViTXT-GQA test set. The green and yellow boxes in video frames denote the ground-truth boxes and the model-grounded OCR tokens. The blue bar denotes the annotated temporal segment.

improves grounding accuracy when using a fixed detection system. These results demonstrate that high-quality detection-recognition combinations enhance grounding performance and optimize QA accuracy.

Qualitative analysis. As shown in Fig. 7, in **Q1**, T2S-QA successfully identifies four frames at 1.4s, 2.3s, 2.5s, and 2.9s, and then pinpoints the answer region in frames 1.4s, 2.3s, and 2.9s. However, due to insufficient light and partial occlusion, an error occurred in scene text recognition where “chips” is mistakenly recognized as “hips” leading to an incorrect model answer. In **Q2**, T2S-QA correctly locates the key video frames and scene text “0.79” for answering the question. However, due to a one-word difference between the predicted answer and the ground truth answer, with an extra word “a”, the predicted answer is considered incorrect. In **Q1** and **Q2**, the grounded visual answers served as valuable evidence in diagnosing the issues with the model. In **Q3**, although the model answers the question correctly, it does not correctly locate the price of brown rice on the right, but locates other products with the same price. Therefore, the answer may be based on the price of other products “1.25” on the incorrectly identified video frames of 1.4s, 1.5s, 1.6s, 2.3s, and 2.4s. In **Q4**, T2S-QA correctly locates and recognizes the scene text “stop” for answering the question. In conclusion, T2S-QA provides visual evidence for answer prediction, enhancing the transparency of the decision-making process. It assists us in diagnosing the underlying causes of failure prediction more effectively, thereby enabling better model optimization.

VI. CONCLUSION

We explore a novel task called Grounded TextVideoQA. The task is crucial for developing trustworthy VideoQA systems involving scene texts. The task also allows for better evaluation of models by enabling direct assessment of grounded scene text regions instead of just textual answers. To solve the

task, we propose T2S-QA, a visually grounded TextVideoQA model featuring a disentangled temporal-to-spatial contrastive grounding mechanism for scene text grounding and grounded TextVQA. Additionally, we construct the ViTXT-GQA benchmark to facilitate evaluation, which enables direct evaluation of spatio-temporal scene-text grounding. Comparative results with existing models and ablated variants demonstrate the superiority and effectiveness of our approach in both QA and answer grounding. Moreover, we conduct an in-depth analysis of the model’s specific behavior concerning grounding and QA accuracy. The results indicate that the primary reasons for poor TextVideoQA performance are the sub-optimal OCR system and ineffective evaluation metrics. We hope our dataset, baselines, and analyses will contribute to the advancements in faithfully answering scene-text questions in videos.

REFERENCES

- [1] Jon Almazán, Albert Gordo, Alicia Fornés, and Ernest Valveny. Word spotting and recognition with embedded attributes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(12):2552–2566, 2014.
- [2] Ali Furkan Biten, Ruben Tito, Andres Mafla, Lluís Gomez, Marçal Rusinol, CV Jawahar, Ernest Valveny, and Dimosthenis Karatzas. Scene text visual question answering. In *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, pages 4291–4301, 2019.
- [3] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. Enriching word vectors with subword information. In *Trans. Assoc. Comput. Linguist.*, pages 135–146, 2017.
- [4] Chongyan Chen, Samreen Anjum, and Danna Gurari. Grounding answers for visual questions asked by visually impaired people. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 19098–19107, 2022.
- [5] Chongyan Chen, Samreen Anjum, and Danna Gurari. Vqa therapy: Exploring answer differences by visually grounding answers. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 15315–15325, 2023.
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning (ICML)*, pages 1597–1607, 2020.

- [7] Yi Cheng, Hehe Fan, Dongyun Lin, Ying Sun, Mohan Kankanalli, and Joo-Hwee Lim. Keyword-aware relative spatio-temporal graph networks for video question answering. *IEEE Trans. Multimedia*, 26:6131–6141, 2023.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. 2019.
- [9] Yang Ding, Jing Yu, Bang Liu, Yue Hu, Mingxin Cui, and Qi Wu. Mukeya: Multimodal knowledge extraction and accumulation for knowledge-based visual question answering. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 5089–5098, 2022.
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. 2020.
- [11] Chengyang Fang, Jiangnan Li, Liang Li, Can Ma, and Dayong Hu. Separate and locate: Rethink the text in text-based visual question answering. In *Proc. ACM Int. Conf. Multimedia*, pages 4378–4388, 2023.
- [12] Shancheng Fang, Hongtao Xie, Yuxin Wang, Zhendong Mao, and Yongdong Zhang. Read like humans: Autonomous, bidirectional and iterative language modeling for scene text recognition. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 7098–7107, 2021.
- [13] Chuang Gan, Yandong Li, Haoxiang Li, Chen Sun, and Boqing Gong. Vqs: Linking segmentations to questions and answers for supervised attention in vqa and question-focused semantic segmentation. In *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, pages 1811–1820, 2017.
- [14] Difei Gao, Luowei Zhou, Lei Ji, Linchao Zhu, Yi Yang, and Mike Zheng Shou. Mist: Multi-modal iterative spatial-temporal transformer for long-form video question answering. In *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, pages 14773–14783, 2023.
- [15] Xuehao Gao, Yang Yang, Shaoyi Du, Guo-Jun Qi, and Junwei Han. Jointly understand your command and intention: Reciprocal co-evolution between scene-aware 3d human motion synthesis and analysis. *arXiv preprint arXiv:2503.00371*, 2025.
- [16] Xuehao Gao, Yang Yang, Zhenyu Xie, Shaoyi Du, Zhongqian Sun, and Yang Wu. Guess: Gradually enriching synthesis for text-driven human motion generation. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 30(12):7518–7530, 2024.
- [17] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 6904–6913, 2017.
- [18] Dan Guo, Kun Li, Bin Hu, Yan Zhang, and Meng Wang. Benchmarking micro-action recognition: Dataset, method, and application. *IEEE Trans. Circuits Syst. Video Technol.*, 34(7):6238–6252, 2024.
- [19] Haibin He, Maoyuan Ye, Jing Zhang, Juhua Liu, and Dacheng Tao. Gomatching: A simple baseline for video text spotting via long and short term matching. *arXiv preprint arXiv:2401.07080*, 2024.
- [20] Ronghang Hu, Amanpreet Singh, Trevor Darrell, and Marcus Rohrbach. Iterative answer prediction with pointer-augmented multimodal transformers for textvqa. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 9992–10002, 2020.
- [21] Soumya Jahagirdar, Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. Watching the news: Towards videoqa models that can read. In *IEEE/CVF Winter Conf. Appl. Comput. Vision*, pages 4441–4450, 2023.
- [22] Zan-Xia Jin, Heran Wu, Chun Yang, Fang Zhou, Jingyan Qin, Lei Xiao, and Xu-Cheng Yin. Ruart: A novel text-centered solution for text-based visual question answering. *IEEE Trans. Multimedia*, 25:1–12, 2021.
- [23] Gouthaman Kv and Anurag Mittal. Reducing language biases in visual question answering with visually-grounded question encoder. In *Proc. Eur. Conf. Comput. Vis.*, pages 18–34, 2020.
- [24] Jie Lei, Licheng Yu, Mohit Bansal, and Tamara L Berg. Tvqa: Localized, compositional video question answering. In *Proc. Conf. Empir. Methods Nat. Lang. Process.*, 2018.
- [25] Jie Lei, Licheng Yu, Tamara Berg, and Mohit Bansal. Tvqa+: Spatio-temporal grounding for video question answering. 2019.
- [26] Yicong Li, Xiang Wang, Junbin Xiao, and Tat-Seng Chua. Equivariant and invariant grounding for video question answering. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4714–4722, 2022.
- [27] Yicong Li, Xiang Wang, Junbin Xiao, Wei Ji, and Tat-Seng Chua. Invariant grounding for video question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2928–2937, 2022.
- [28] Yicong Li, Xiang Wang, Junbin Xiao, Wei Ji, and Tat-Seng Chua. Transformer-empowered invariant grounding for video question answering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [29] Yicong Li, Junbin Xiao, Chun Feng, Xiang Wang, and Tat-Seng Chua. Discovering spatio-temporal rationales for video question answering. In *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, pages 13869–13878, 2023.
- [30] Yi Liu, Xiaohui Dong, Dingwen Zhang, and Shoukun Xu. Deep unsupervised part-whole relational visual saliency. *Neurocomputing*, 563:126916, 2024.
- [31] Yi Liu, Dingwen Zhang, Nian Liu, Shoukun Xu, and Jungong Han. Disentangled capsule routing for fast part-object relational saliency. *IEEE Trans. Image Process.*, 31:6719–6732, 2022.
- [32] Yulei Niu, Kaihua Tang, Hanwang Zhang, Zhiwu Lu, Xian-Sheng Hua, and Ji-Rong Wen. Counterfactual vqa: A cause-effect look at language bias. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 12700–12710, 2021.
- [33] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2019.
- [34] OpenAI. Gpt-4o mini: advancing cost-efficient intelligence. 2024.
- [35] Tianwen Qian, Jingjing Chen, Shaoxiang Chen, Bo Wu, and Yu-Gang Jiang. Scene graph refinement network for visual question answering. *IEEE Trans. Multimedia*, 25:3950–3961, 2022.
- [36] Tianwen Qian, Ran Cui, Jingjing Chen, Pai Peng, Xiaowei Guo, and Yu-Gang Jiang. Locate before answering: Answer guided question localization for video question answering. *IEEE Trans. Multimedia*, 26:4554–4563, 2024.
- [37] Varun Nagaraj Rao, Xingjian Zhen, Karen Hovsepian, and Mingwei Shen. A first look: Towards explainable textvqa models via visual and textual explanations. 2021.
- [38] Hao Shen, Zhong-Qiu Zhao, Yulun Zhang, and Zhao Zhang. Mutual information-driven triple interaction network for efficient image dehazing. In *Proc. ACM Int. Conf. Multimedia*, pages 7–16, 2023.
- [39] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 8317–8326, 2019.
- [40] Peipei Song, Dan Guo, Xun Yang, Shengeng Tang, and Meng Wang. Emotional video captioning with vision-based emotion interpretation network. *IEEE Transactions on Image Processing*, 33:1122–1135, 2024.
- [41] George Tom, Minesh Mathew, Sergi Garcia-Bordils, Dimosthenis Karatzas, and CV Jawahar. Reading between the lanes: Text videoqa on the road. In *Proc. Int. Conf. Doc. Anal. Recognit.*, pages 137–154, 2023.
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Adv. Neural Inf. Process. Syst.*, 2017.
- [43] Jianyu Wang, Bing-Kun Bao, and Changsheng Xu. Dualvgr: A dual-visual graph reasoning unit for video question answering. *IEEE Trans. Multimedia*, 24:3369–3380, 2021.
- [44] Peng Wang, Shuai Bai, Siman Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024.
- [45] Zixiao Wang, Hongtao Xie, Yuxin Wang, Jianjun Xu, Boqiang Zhang, and Yongdong Zhang. Symmetrical linguistic feature distillation with clip for scene text recognition. In *Proc. ACM Int. Conf. Multimedia*, pages 509–518, 2023.
- [46] Weijia Wu, Yuanqiang Cai, Debing Zhang, Sibao Wang, Zhuang Li, Jiahong Li, Yejun Tang, and Hong Zhou. A bilingual, openworld video text dataset and end-to-end video text spotter with transformer. *arXiv preprint arXiv:2112.04888*, 2021.
- [47] Junbin Xiao, Xindi Shang, Xun Yang, Sheng Tang, and Tat-Seng Chua. Visual relation grounding in videos. In *Proc. Eur. Conf. Comput. Vis.*, pages 447–464. Springer, 2020.
- [48] Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. Next-qa: Next phase of question-answering to explaining temporal actions. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, pages 9777–9786, 2021.
- [49] Junbin Xiao, Angela Yao, Yicong Li, and Tat Seng Chua. Can i trust your answer? visually grounded video question answering. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2024.
- [50] Junbin Xiao, Pan Zhou, Angela Yao, Yicong Li, Richang Hong, Shuicheng Yan, and Tat-Seng Chua. Contrastive video question answering via video graph transformer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(11):13265–13280, 2023.
- [51] Xun Yang, Jianming Zeng, Dan Guo, Shanshan Wang, Jianfeng Dong, and Meng Wang. Robust video question answering via contrastive cross-modality representation learning. *SCIENCE CHINA Information Sciences*, 67:1–16, 2024.
- [52] Shoubin Yu, Jaemin Cho, Prateek Yadav, and Mohit Bansal. Self-chained image-language model for video localization and question answering. In

Adv. Neural Inf. Process. Syst., pages 76749–76771, 2024.

- [53] Jiahao Zhang, Stephen Gould, and Itzik Ben-Shabat. Vidat—ANU CVML video annotation tool. <https://github.com/anucvml/vidat>, 2020.
- [54] Tong Zhang, Hao Fang, Hao Zhang, Jialin Gao, Xiankai Lu, Xiushan Nie, and Yilong Yin. Learning feature semantic matching for spatio-temporal video grounding. *IEEE Trans. Multimedia*, 26:9268–9279, 2024.
- [55] Minyi Zhao, Bingjia Li, Jie Wang, Wanqing Li, Wenjing Zhou, Lan Zhang, Shijie Xuyang, Zhihang Yu, Xinkun Yu, Guangze Li, et al. Towards video text visual question answering: Benchmark and baseline. In *Adv. Neural Inf. Process. Syst.*, pages 35549–35562.
- [56] Sheng Zhou, Dan Guo, Jia Li, Xun Yang, and Meng Wang. Exploring sparse spatial relation in graph inference for text-based vqa. *IEEE Trans. Image Process.*, 32:5060–5074, 2023.
- [57] Sheng Zhou, Dan Guo, Xun Yang, Jianfeng Dong, and Meng Wang. Graph pooling inference network for text-based vqa. *ACM Trans. Multimedia Comput. Commun. Appl.*, 20(4):1–21, 2024.
- [58] Sheng Zhou, Junbin Xiao, Qingyun Li, Yicong Li, Xun Yang, Dan Guo, Meng Wang, Tat-Seng Chua, and Angela Yao. Egotextvqa: Towards egocentric scene-text aware video question answering. *arXiv preprint arXiv:2502.07411*, 2025.
- [59] Wenwu Zhu, Xin Wang, and Wen Gao. Multimedia intelligence: When multimedia meets artificial intelligence. *IEEE Trans. Multimedia*, 22(7):1823–1835, 2020.
- [60] Yongxin Zhu, Zhen Liu, Yukang Liang, Xin Li, Hao Liu, Changcun Bao, and Linli Xu. Locate then generate: Bridging vision and language with bounding box for scene-text vqa. In *Proc. AAAI Conf. Artif. Intell.*, pages 11479–11487, 2023.



Sheng Zhou received the B.E. degree in Internet of Things from Hengyang Normal University, China, in 2020. She is currently pursuing a Ph.D. degree in the School of Computer Science and Technology, Hefei University of Technology, China. Her research interests include multimodal learning, visual and language. She has published relevant papers in conferences and journals: CVPR, IEEE TIP, and ACM ToMM. She also serves as a reviewer for: ACM MM, IJCNN, IEEE TIP, ACM ToMM, and Neurocomputing.



Junbin Xiao received the BE degree from Sichuan University (SCU), the MS Eng degree from the Institute of Computing Technology (ICT), Chinese Academy of Sciences (CAS), and the PhD degree from the Department of Computer Science (CS), National University of Singapore (NUS), and currently serves as a research fellow with the Department of Computer Science (CS), National University of Singapore (NUS). His research focuses on visual relation-oriented VideoQA. He has published relevant papers in the top-tier conferences: CVPR, ECCV, AAAI, ACM MM, and EMNLP. He also serves as a reviewer for: CVPR, ICCV, ECCV, AAAI, IEEE Transactions on Multimedia.



Xun Yang received his Ph.D. degree from the Hefei University of Technology, Hefei, China, in 2017. He is currently a Professor with the Department of Electronic Engineering and Information Science, University of Science and Technology of China (USTC). From 2015 to 2017, he visited the University of Technology Sydney (UTS), Australia as a joint Ph.D student. He was a Research Fellow with the NEXt++ Research Center, National University of Singapore (NUS), from 2018 to 2021. His current research interests include information retrieval, cross-media analysis and reasoning, and computer vision. He regularly serves as the PC member and the invited reviewer for top-tier conferences and prestigious journals in multimedia and artificial intelligence, like the ACM Multimedia, IJCAI, AAAI, CVPR, and ICCV.



Peipei Song received the B.E. degree in electronic information engineering and the Ph.D. degree in signal and information processing from the Hefei University of Technology, China, in 2017 and 2023, respectively. She is currently a Postdoctoral Researcher with the Department of Electronic Engineering and Information Science, University of Science and Technology of China (USTC). Her research interests include computer vision, natural language processing, and video understanding.



Dan Guo received the B.E. degree in computer science and technology from Yangtze University, China, in 2004, and the Ph.D. degree in System Analysis and Integration from Huazhong University of Science and Technology, China, in 2010. She is currently an Associate Professor with the School of Computer Science and Information Engineering, Hefei University of Technology, China. Her research interests include computer vision, machine learning, and intelligent multimedia content analysis.



Angela Yao received the B.A.Sc. degree from the University of Toronto, Canada, and the PhD degree from ETH. She is currently an assistant professor and leads the Computer Vision and Machine Learning (CVML) group with the National University of Singapore. Her research interests include video understanding, 3D pose estimation, activity recognition, sequence and time series modelling. Her research interest lies at the intersection of computer vision, machine learning, and human-computer interaction, with a particular emphasis on developing models and algorithms that can understand human actions and interactions in video data. She serves as the area chair and reviewer for ICLR, ICML, NeurIPS, CVPR, ICCV, IJCAI, and program chair for ICCV, ECCV and 3DGV. She once led the Visual Computing Group with the University of Bonn, Germany and co-founded a startup on smart parking in Zurich, Switzerland.



Meng Wang (SM'17) received the B.E. and Ph.D. degrees in the special class for the gifted young from the Department of Electronic Engineering and Information Science, University of Science and Technology of China, Hefei, China, in 2003 and 2008, respectively. He is currently a Professor with the Hefei University of Technology, China. He has authored over 200 book chapters, journal and conference papers in his research areas. His current research interests include multimedia content analysis, computer vision, and pattern recognition. He was a recipient of the ACM SIGMM Rising Star Award 2014. He is an Associate Editor of the IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, and the IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS.



Tat-Seng Chua received the Ph.D. degree from the University of Leeds, U.K. He is the KITHCT Chair Professor with the School of Computing, National University of Singapore, where he was the Acting and Founding Dean of the School from 1998 to 2000. His main research interests include multimedia information retrieval and social media analytics. In particular, his research focuses on the extraction, retrieval, and question-answering (QA) of text and rich media arising from the web and multiple social networks. He is the Co-Director of NEXt, a joint center between NUS and Tsinghua University, to develop technologies for live social media search. He is the 2015 winner of the prestigious ACM SIGMM Award for Outstanding Technical Contributions to Multimedia Computing, Communications, and Applications. He serves on the editorial boards of four international journals. He is also the co-founder of two technology startup companies in Singapore.