

# Track the Answer: Extending TextVQA from Image to Video with Spatio-Temporal Clues

Yan Zhang<sup>1,3\*</sup>, Gangyan Zeng<sup>4\*</sup>, Huawen Shen<sup>1,3</sup>, Daiqing Wu<sup>1,3</sup>, Yu Zhou<sup>2†</sup>, Can Ma<sup>1,3†</sup>

<sup>1</sup>Institute of Information Engineering, Chinese Academy of Sciences, China

<sup>2</sup>VCIP & TMCC & DISec, College of Computer Science, Nankai University, China

<sup>3</sup>School of Cyber Security, University of Chinese Academy of Sciences, China

<sup>4</sup>School of Cyber Science and Engineering, Nanjing University of Science and Technology, China

{zhangyan2022,shenhuawen,wudaiqing,macan}@iie.ac.cn, yzhou@nankai.edu.cn, gyzeng@njjust.edu.cn

## Abstract

Video text-based visual question answering (Video TextVQA) is a practical task that aims to answer questions by jointly reasoning textual and visual information in a given video. Inspired by the development of TextVQA in image domain, existing Video TextVQA approaches leverage a language model (e.g. T5) to process text-rich multiple frames and generate answers auto-regressively. Nevertheless, the spatio-temporal relationships among visual entities (including scene text and objects) will be disrupted and models are susceptible to interference from unrelated information, resulting in irrational reasoning and inaccurate answering. To tackle these challenges, we propose the TEA (stands for “Track thE Answer”) method that better extends the generative TextVQA framework from image to video. TEA recovers the spatio-temporal relationships in a complementary way and incorporates OCR-aware clues to enhance the quality of reasoning questions. Extensive experiments on several public Video TextVQA datasets validate the effectiveness and generalization of our framework. TEA outperforms existing TextVQA methods, video-language pretraining methods and video large language models by great margins.

**Code** — <https://github.com/zhangyan-ucas/TEA>

## Introduction

Understanding scene text in images and videos plays an important role in interpreting the physical world, and also is the key to solving many practical applications, such as self-driving vehicles, automatic news reporting and scene text editing (Zeng et al. 2024b; Li et al. 2024). As a representative scene text understanding task (Zeng et al. 2024a), text-based visual question answering (TextVQA) has attracted much interest from the communities and witnessed tremendous progress (Singh et al. 2019; Yang et al. 2021; Fang et al. 2023; Zeng et al. 2023b; Zhang et al. 2024). A majority of work focuses on TextVQA in image domain (i.e. Image TextVQA), which aims at answering questions related

\*These authors contributed equally.

†Corresponding authors: Yu Zhou and Can Ma.

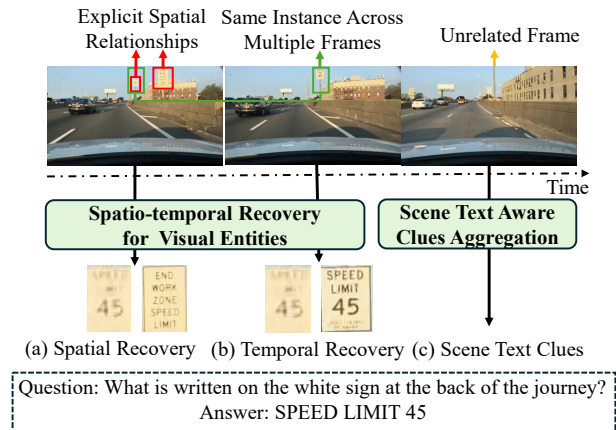


Figure 1: An example to illustrate the challenges in Video TextVQA and the overview of our proposed method. For this given question, the Video TextVQA model requires to first analyze the spatial relationships to locate the answer instance, and then capture the temporal correspondence to track and recognize the answer content. Besides, unrelated frames may interfere with model reasoning. To tackle these challenges, we propose to recover the spatio and temporal information respectively, and aggregate the scene text aware clues for reasonable guidance.

to the text in a given image. Traditional methods (Hu et al. 2019; Zhu et al. 2021) adopt a discriminative framework, which predicts answers from Optical Character Recognition (OCR) tokens (Qin et al. 2023) or vocabulary words. Benefiting from the powerful representation of pretrained language model T5 (Raffel et al. 2019), recent works (Biten et al. 2021; Zeng et al. 2023a) address this task in a generative way, thus significantly boosting the accuracy and generalizability.

In this context, a new task - Video TextVQA goes one step further, which requires reading and reasoning scene text throughout a video. Compared to Image TextVQA, this task is more challenging as the answer can appear anywhere in any frame of the video. To tackle it, an intuitive solution is to adapt the most competitive method for Image

TextVQA to the video domain. Following it, T5-ViteVQA (Zhao et al. 2022) introduces five Transformer blocks to encode modality-specific features across video frames, and then feeds them into a T5-like Transformer encoder-decoder to infer the answer.

Although the applicability of this vanilla pipeline has been demonstrated, there are still some limitations to consider. As illustrated in Fig. 1, answering a Video TextVQA question often requires a comprehensive understanding of contextual information from multiple frames. On the one hand, the spatial relationships between entities within a frame (e.g., the two white signs in the first frame) should be correctly interpreted, and on the other hand, the instance correspondence between different frames (e.g., the sign “SPEED LIMIT 45” appearing in both the first and second frames) should be captured. Following general VideoQA’s practice, current efforts tend to employ absolute-location embedding and temporal index to represent spatio-temporal information respectively. However, they ignore that Video TextVQA deviates from general VideoQA because a large amount of scene text will be involved. Considering scene text is relatively small and exhibits rich layout characteristics, this implicit learning method is hard to obtain satisfactory results. Besides, videos often involve a lot of information that is not related to answering the question (e.g., the unrelated third frame). This redundancy complicates the reasoning process, leading to challenges in analyzing intricate scenarios and filtering useless contents. Previous methods disperse the attention across the entire video and ignore the importance of question-guided reasoning cues, thus tending to predict inaccurate answers.

To overcome the two aforementioned issues, we propose the **TEA** (stands for “Track thE Answer”) method, which promotes the application of the T5-based generative framework to Video TextVQA in two aspects. 1) **Recover the spatio-temporal information.** We contend that previous generative Video TextVQA architectures organizing all visual entities into a 1-D sequence as input disrupt the spatio-temporal relationships that exist in the video. To this end, TEA recovers the spatio-temporal information in a complementary way. First, a temporal convolution module is utilized to fully exploit the continuity characteristics of visual entities in the temporal dimension, so as to capture the dynamic evolution of scene text across video frames. Then, we incorporate multi-granularity relative spatial relationships among scene text and explicitly encode them in the attention computation. In this manner, the composition of words inside each text instance and the spatial interaction between different entity instances can be easily attained. 2) **Aggregate the essential clues.** Understanding the intent of question and following a rational reasoning route are prerequisites for language model to generate correct answers. Following it, we design the scene text aware clues aggregation module, which progressively aggregates useful clues driven by the given question, suppressing the redundant message passing. Extensive experimental results on multiple public datasets underscore the effectiveness of our proposed framework. TEA significantly outperforms existing methods, including image and video TextVQA meth-

ods, video-language pretraining methods and video large language models (LLMs), marking a considerable advancement in the field of Video TextVQA.

The contributions of our work are threefold:

- Motivated by the development of Image TextVQA methods, we propose TEA, a generative Video TextVQA model that enables language model to understand 3-D video features and generate reasonable answers.
- To track the dynamic evolution of answers in the video, a temporal convolution module and an OCR-enhanced relative spatial module are developed to explicitly model spatio-temporal information among visual entities. To avoid interference from irrelevant information, a scene text aware clues aggregation module is utilized to extract the question-related reasoning clues.
- Extensive experiments verify the effectiveness and superiority of our method TEA on several representative Video TextVQA datasets.

## Related Work

### Text-based Visual Question Answering

**Image TextVQA.** Image TextVQA task aims to answer questions by reading and understanding the textual contents in images. Early Image TextVQA approaches (Hu et al. 2019; Kant et al. 2020; Gao et al. 2020) predominantly predict answers from a fixed vocabulary list and OCR tokens (Qiao et al. 2021, 2020; Wang et al. 2022c; Du et al. 2024; Zheng et al. 2024) present in the images. M4C (Hu et al. 2019) leverages a Transformer-based multi-modal architecture and decodes answers auto-regressively with a dynamic pointer network. Built on the top of M4C, SA-M4C (Kant et al. 2020) incorporates the self-attention mask to enhance the spatial relationships among visual entities. Motivated by the success of visual-language multi-modal pre-training, TAP (Yang et al. 2021) designs Image TextVQA-specific pre-training tasks to align visual entities for scene text understanding. Recently, generative Image TextVQA approaches (Biten et al. 2021; Zeng et al. 2023a; Fang et al. 2023) have emerged which enhance the performance by leveraging powerful reasoning abilities and world knowledge from T5 (Raffel et al. 2019). For instance, Biten et al. (2021) utilizes T5 to initialize the proposed LaTr model and address the question-answering problem in a text generation way.

**Video TextVQA.** Video TextVQA requires models to answer questions by jointly reasoning textual and visual information in a given video. Compared to Image TextVQA, Video TextVQA is more challenging as the answer can appear anywhere in any frame of the video. Additionally, the scene text in videos exhibits richer layout characteristics and more varied scales and forms. Motivated by the developments in generative Image TextVQA methods, T5-ViteVQA (Zhao et al. 2022) employs absolute-location embedding and temporal index to represent spatio-temporal among visual entities. However, T5-ViteVQA fails to establish explicit spatio-temporal relationships among visual entities and efficiently process redundant visual entities in videos.

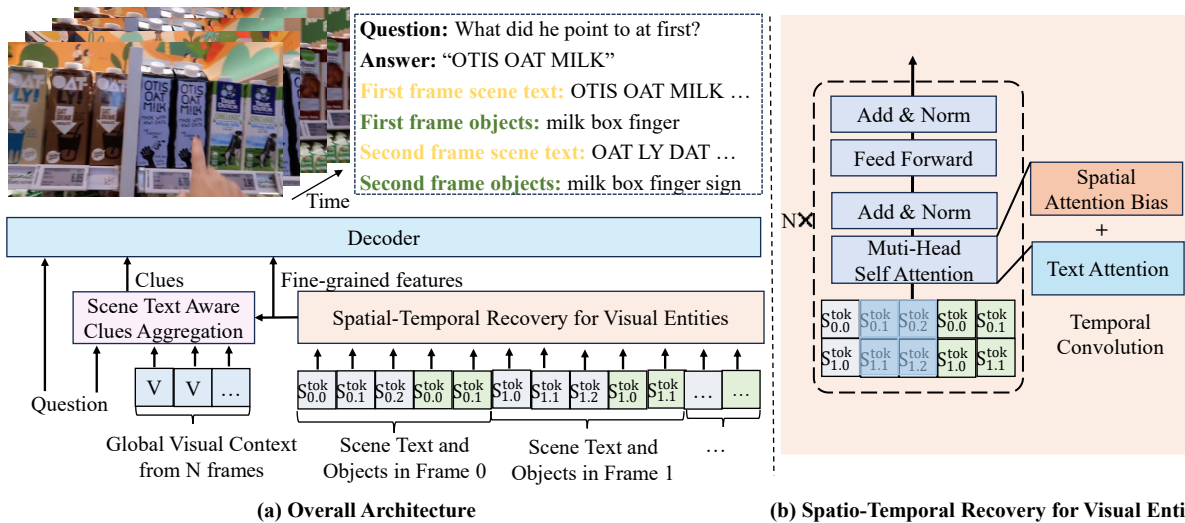


Figure 2: (a) Overall architecture of TEA. Blue and green token  $S_{i,j}^{tok}$  represents the  $j$ th OCR and object token in  $i$ th frame, respectively. (b) Illustration of spatio-temporal recovery for visual entities.

## Video Question Answering

Video question answering (VideoQA) is a crucial task in the visual-language domain, significantly enhancing AI systems' ability to interpret and interact with multimedia content. In recent years, VideoQA has attracted substantial interest from the research community, leading to various methodological advancements. Recent VideoQA methods can be categorized into several prominent approaches: 1) Memory Networks-based Methods: Gao et al. (2018) leverage memory networks to store and process long-term relationships across video frames. 2) Hierarchy Reasoning-based Methods: Le et al. (2020); Xiao et al. (2021) hierarchically decompose questions and videos into sub-questions and video clips as well as derive the final answer based on the reasoning associated with each sub-question. 3) Transformer-based Methods: Gao et al. (2022) take advantage of the Transformer architecture in long-term relationships modeling to facilitate VideoQA research (Wang et al. 2025).

Despite significant advancements in VideoQA, a common limitation persists: the focus predominantly remains on motions and events occurring in videos, disregarding textual information present in videos. Addressing this gap is crucial for developing AI systems capable of fully comprehending and interacting with a broader spectrum of visual content.

## Method

### Overview Framework of TEA

Fig. 2 illustrates the overall architecture of TEA, which primarily comprises visual external systems for extracting fine-grained visual entities and a T5-like Transformer encoder-decoder architecture for reasoning questions. For a given video, we uniformly sample it into multiple  $T$  frames and utilize an OCR system (Wu et al. 2021; Fang et al. 2021; Ye et al. 2024) to determine the location and content of scene text while concurrently employing an object detector (Ren

et al. 2015) to extract objects as auxiliary information to better understand the scene text.

Following Image TextVQA methods (Zeng et al. 2023a), we leverage a pre-trained language model (Raffel et al. 2019) as the foundational model for the Video TextVQA task. To harness the reasoning capabilities and world knowledge of the foundational model, we convert the fine-grained features for Video TextVQA into a textual sequence to achieve input alignment. Specifically, a TEA-Baseline method is proposed, where we concatenate the contents of scene text and the labels of visual entities frame by frame into a sequence, as shown in Fig. 2. Consequently, in TEA-Baseline, the video is represented as  $v \in \mathbb{R}^{T(M+N) \times d}$ , where  $T$  indicates the number of frames,  $d$  denotes the feature dimension of visual entities,  $M$  and  $N$  indicate the number of scene text and objects within each video frame. However, simply feeding a sequence of Video TextVQA features into the foundational model leads to two issues: the disruption of spatio-temporal relationships among visual entities and the inability to effectively process long video sequences that include redundant fine-grained features. Therefore, TEA proposes two significant modules in the following section to recover the spatio-temporal relationships and generate scene text-aware clues to assist in generating answers.

### Spatio-Temporal Recovery for Visual Entities

In this section, we propose two sub-modules to separately recover temporal and spatial relationships. Specifically, a convolution operator is introduced to exploit the continuity characteristics of visual entities in the temporal dimension. Additionally, an OCR-enhanced spatial bias is incorporated into the attention mechanism to emphasize relative spatial relationships.

**Temporal Convolution Module** Typically, T5-based Image TextVQA methods do not account for temporal modeling. To address this limitation, we design a temporal convo-

lution module to reason about the temporal relationships between the same instance across different video frames. This design adheres to two crucial criteria: ensuring interaction between the same instance in different frames, and avoiding modifying the T5 representation space to maintain its reasoning abilities and world knowledge. Based on these principles, as shown in Fig. 2(b), this sub-module employs a disentangled strategy to separately recover temporal and spatial relationships among visual entities and inserts a temporal convolution module between consecutive attention layers.

To be concrete, given the fine-grained features of a video ( $v \in \mathbb{R}^{T(M+N) \times d}$ ) with two dimensions of  $T(M+N)$  and  $d$ , we convert it into  $v \in \mathbb{R}^{T \times (M+N) \times d}$  with three dimensions of  $T$ ,  $(M+N)$ , and  $d$ . A temporal convolution module is proposed to recover temporal relationships for the same instance across different video frames as follows:

$$v' = v + \text{Conv2D}(v \times W_{down}) \times W_{up}, \quad (1)$$

where  $v' \in \mathbb{R}^{T \times (M+N) \times d}$  is the temporal recovered video feature and  $\text{Conv2D}$  denotes 2D-convolution.  $W_{down}$  and  $W_{up}$  indicate the projections for reducing and increasing dimensions. Specifically, a down-projection and up-projection strategy is applied before and after the temporal convolution. To capture temporal relationships for all visual entities across frames, the kernel size of the temporal convolution is designed with consideration of both the length of visual entities and the temporal length of videos. Furthermore, it is worth noting that benefiting from the strategy for disentangled spatio-temporal modeling, the computational complexity is also significantly reduced.<sup>1</sup>

**OCR-enhanced Relative Spatial Module** In the realm of text-rich understanding analysis, relative spatial relationships for visual entities are predominantly central (Hong et al. 2021), particularly for multi-granularity spatial relationships involving scene text (Ye et al. 2024). For instance, a representative question in Video TextVQA, such as “*What is written on the white sign at the back of the journey?*” in Fig. 1, requires the model to locate relative spatial relationships between the last white sign and the text on it. Meanwhile recognizing “SPEED”, “LIMIT”, “45” belong to the same paragraph-level bounding box allows for the establishment of associated semantic information. Therefore, we propose a novel OCR-enhanced relative spatial module that explicitly recovers the multi-granularity spatial relationships during attention calculations.

Following previous Image TextVQA methods built on the top of T5 encoder, we apply for multi-head self-attention mechanism to interact with multi-modal features and introduce an attention bias indicating relative spatial relationships, defined as:

$$e_{ij} = \frac{Q_i * K_j^T}{\sqrt{d_z}} + a_{ij}, \quad (2)$$

where  $e_{ij} \in \mathbb{R}^{head}$  represents the multi-head attention,  $a_{ij}$  indicates the relative spatial bias,  $d_z$  represents the hidden

<sup>1</sup>Joint spatio-temporal modeling complexity  $O(v) = T^2(M + N)^2$  vs. Disentangled spatio-temporal modeling complexity  $O(v) = T(M + N)^2$

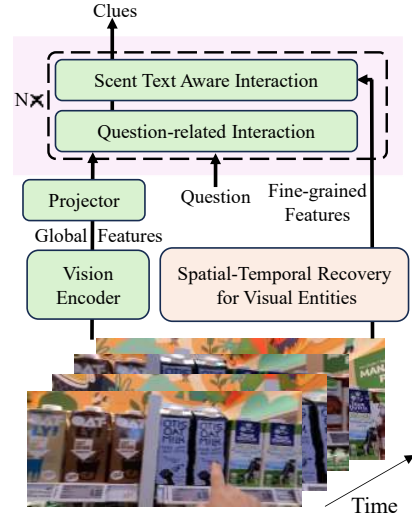


Figure 3: Illustration of scene text aware clues aggregation.

dimension.  $Q$  and  $K$  denote multi-head self-attention query and key for visual entities  $i, j$ , respectively. Note that in Image TextVQA methods  $a_{ij}$  indicates 1-D relative relationships, while in TEA it indicates 2-D relative relationships.

Generally, we leverage visual external systems (Ren et al. 2015; Ye et al. 2024) to obtain spatial locations for scene text and visual objects. For formal description, denoting the normalized bounding box for each visual entity as  $(x_{tl}, y_{tl}, x_{br}, y_{br})$ , inspired by (Hong et al. 2021), we compute relative spatial bias  $a_{ij} \in \mathbb{R}^{head}$  with sinusoidal functions (Vaswani 2017):

$$a_{ij} = [f^{\sin}(x_i - x_j); f^{\cos}(y_i - y_j)]. \quad (3)$$

To enhance the semantic coherence in answering questions, we incorporate multi-granularity spatial relationships for scene text within the relative positional embedding. Specifically, we provide bounding boxes at the word, line, and paragraph levels for calculating the relative positional embeddings. Consequently, the term  $a_{ij}$  in the equation 3 can be expressed as follows:

$$a_{ij} = \begin{cases} a_{ij}^{word} + a_{ij}^{line} + a_{ij}^{para} & \text{if } (i, j) \in \text{scene text,} \\ a_{ij}^{word} & \text{otherwise.} \end{cases} \quad (4)$$

### Scene Text Aware Clues Aggregation

From the previous section, we can derive spatio-temporal features for visual entities. However, directly feeding these features into the decoder causes difficulties in answering certain questions, as it tends to focus too heavily on redundant fine-grained visual entities. Therefore, we propose a module that generates scene text-aware clues for the decoder to enhance the quality of reasoning within questions. In our design, we adhere to three practically crucial criteria: (1) simple and concise clues that indicate global context, (2) relevance to the question, and (3) awareness of fine-grained in-

formation, specifically related to scene text. This module is divided into two steps, as illustrated in Fig. 3.

**Question-related Global Context** In the field of image recognition analysis, a single [CLS] token, positioned at the beginning of the visual encoder, represents the global context of an entire image. Therefore, we utilize the [CLS] token from the CLIP (Radford et al. 2021) vision encoder to capture the global context for visual semantics in each video frame. To guide these global context with Video TextVQA questions, we adopt a self-attention that integrates global features with question features.

In practice, to generalize different sampling strategies, we design a projector to align multiple frames [CLS] token features into a fixed length. A self attention is proposed to extract global features relevant to the question. Specifically, we connect the video global context feature with the corresponding question to form a sequence. After fusion in self-attention, we extract the question-aware video global context as clues:

$$\begin{cases} X = [Global; Ques], \\ X = SelfAttn(X), \\ Clues = X[: len(Global)], \end{cases} \quad (5)$$

where  $[:]$  represents the concatenation of two features, and  $len$  denotes the length of features.

**Scene Text Aware Interaction** Based on question-related interaction, we obtain global features relevant to questions. To supplement the fine-grained details that might be overlooked by the CLIP visual encoder, we integrate them into global contexts with the cross-attention. Specifically, we employ scene text as the fine-grained feature to emphasize its association with the Video TextVQA questions:

$$X = CrossAttn(Q = X; K, V = visual\ entities), \quad (6)$$

where  $Q, K, V$  represent the query, key, and value in cross attention, respectively. Derived from the  $N$ -layer self-cross attention module, scene text-aware clues are then connected with spatio-temporal visual entities and fed into the decoder to generate answers.

## Experiments

### Experiment Settings

**Datasets and Evaluation Metrics** Our method is rigorously evaluated using two public Video TextVQA datasets: M4-ViteVQA (Zhao et al. 2022) and RoadTextVQA (Tom et al. 2023). 1) M4-ViteVQA includes 8,511 video clips with 24,123 question-answer pairs, where the video clips are primarily collected from nine YouTube categories. Specifically, M4-ViteVQA defines two tasks with three settings to evaluate Video TextVQA methods. In Task 1, Video TextVQA methods are trained and tested across all nine categories of M4-ViteVQA. To address different requirements for model robustness, (Zhao et al. 2022) designates *Task1Split1* for regular testing and *Task1Split2* for generalization testing, where the content of videos within the same category varies

significantly (e.g., different shopping venues and sports). In Task 2, Video TextVQA methods are trained using seven categories and tested on the remaining two categories. 2) RoadTextVQA consists of 3,222 driving videos collected from multiple countries, annotated with 10,500 questions, all based on text or road signs present in the driving videos.

For these datasets, we utilize the VQA Accuracy (Acc.) metric and the average normalized Levenshtein similarity (ANLS) to evaluate model performance of TEA.

**Implementation Details** In our experimental setup, following previous Image TextVQA methods, we utilize pre-trained language model as the foundational backbones for processing scene text and objects as well as reasoning answers. Based on the pre-trained T5 (Raffel et al. 2019), we build two models with different capacities, namely TEA-Base and TEA-Large, which have 12+12 and 24+24 encoder-decoder layers, respectively.

### Comparison with Existing Methods

On the M4-ViteVQA *Task1Split1* and RoadTextVQA benchmarks, we evaluate TEA and compare it with existing methods (see Tab. 1). It can be observed that TEA-Large outperforms different types of existing methods. (1) **Image TextVQA Methods.** M4C (Hu et al. 2019) is unable to perform temporal reasoning in videos, achieving an accuracy of 17.91% on M4-ViteVQA test set. (2) **Video-language Pre-training Methods.** These methods (Yang et al. 2020; Wang et al. 2022a,b; Lei, Berg, and Bansal 2023) lack the ability to read and comprehend scene text in videos, resulting in unsatisfactory performance on the Video TextVQA task. All-in-one-B achieves an accuracy of 10.87% on the M4-ViteVQA test set, while GIT achieves 29.58% in accuracy on the RoadTextVQA validation set. (3) **Video-LLMs.** Video-LLMs (Cheng et al. 2024; Lin et al. 2023, 2024) build upon LLMs and leverage a large-scale of video-text pairs pre-training data to align textual and video modalities. To adapt these methods for Video TextVQA, we apply LoRA (Hu et al. 2021) to fine-tune on corresponding dataset. A representative Video-LLMs method, VideoLLaMA2, achieves an accuracy of 20.04% and an ANLS score of 21.73% on the M4-ViteVQA validation set. As reported in the results, Video TextVQA methods are generally superior to Video-LLMs because Video-LLMs lack the capability to capture and understand massive small scene text in videos and do not incorporate architectures specifically designed for Video TextVQA. (4) **Video TextVQA Methods.** Compared with a pioneer Video TextVQA method T5-ViteVQA (Zhao et al. 2022), TEA-Baseline shows a 3.97% accuracy improvement on M4-ViteVQA test set. Although T5 (Raffel et al. 2019) is our common foundation model, TEA-Baseline better aligns multimodal inputs by utilizing textual representation for each visual entity. Compared to TEA-Baseline, TEA-Base recovers the spatio-temporal relationships and provides instructional clues for generating answers, resulting in significant accuracy gains of 5.56% on the M4-ViteVQA test set and 5.23% on the RoadTextVQA validation set. Finally, when upgrading the model capacity to a large size, TEA-Large reaches state-of-the-art perfor-

Method	M4-ViteVQA <i>Task1Split1</i>				RoadTextVQA	
	Validation		Test		Validation	
	Acc.(%)	ANLS(%)	Acc.(%)	ANLS(%)	Acc.(%)	ANLS(%)
M4C (Hu et al. 2019)	18.66	24.20	17.91	23.80	28.92	32.27
JuskAsk (Yang et al. 2020)	10.81	15.40	10.05	14.10	–	–
All-in-one-B (Wang et al. 2022a)	11.47	15.30	10.87	14.80	–	–
GIT (Wang et al. 2022b)	–	–	–	–	29.58	35.23
SINGULARITY (Lei, Berg, and Bansal 2023)	–	–	–	–	24.62	30.79
Video-LLaVA* (Lin et al. 2023)	15.82	17.77	15.43	17.15	30.82	40.92
VideoLLaMA2* (Cheng et al. 2024)	20.04	21.73	20.76	23.55	25.11	36.53
T5-ViteVQA (Zhao et al. 2022)	23.17	30.10	22.17	29.10	–	–
TEA-Baseline	26.79	35.16	26.14	35.14	39.20	45.36
TEA-Base	34.45	42.91	31.70	40.24	44.43	51.69
TEA-Large	<b>37.49</b>	<b>46.38</b>	<b>34.78</b>	<b>43.71</b>	<b>48.14</b>	<b>54.85</b>

Table 1: Accuracy and ANLS comparison with the latest methods on the M4-ViteVQA *Task1Split1* and RoadTextVQA. The best results are highlighted in bold. \* indicates the reproduced results using the open-source code.

Method	M4-ViteVQA <i>Task1Split2</i>				M4-ViteVQA <i>Task2</i>			
	Validation		Test		Validation		Test	
	Acc.(%)	ANLS(%)	Acc.(%)	ANLS(%)	Acc.(%)	ANLS(%)	Acc.(%)	ANLS(%)
M4C (Hu et al. 2019)	13.58	17.20	11.36	16.60	9.16	12.80	7.52	12.50
JuskAsk (Yang et al. 2020)	7.16	10.00	5.47	8.60	4.86	6.70	3.60	6.70
All-in-one-B (Wang et al. 2022a)	6.85	9.20	5.66	7.80	4.20	5.00	3.28	4.60
Video-LLaVA* (Lin et al. 2023)	13.14	14.29	11.19	12.02	10.89	13.23	9.38	11.80
VideoLLaMA2* (Cheng et al. 2024)	18.30	19.63	18.33	20.45	19.68	23.62	16.54	21.80
T5-ViteVQA (Zhao et al. 2022)	17.59	23.10	16.68	23.80	12.30	16.10	9.29	13.60
TEA-Baseline	22.53	31.60	21.20	30.44	17.32	26.32	14.00	21.99
TEA-Base	26.66	36.61	26.29	36.00	20.73	28.18	17.28	26.03
TEA-Large	<b>28.27</b>	<b>36.32</b>	<b>28.43</b>	<b>38.13</b>	<b>22.83</b>	<b>30.21</b>	<b>18.83</b>	<b>28.90</b>

Table 2: Accuracy and ANLS comparison with the latest methods on the M4-ViteVQA *Task1Split2* and *Task2*. The best results are highlighted in bold. \* indicates the reproduced results using the open-source code.

mance, 34.78% and 48.14% on the M4-ViteVQA test set and RoadTextVQA validation set, respectively.

## Analysis of Generalization

In this section, we investigate how TEA contributes to the Video TextVQA from a generalization perspective. Most VideoQA methods tend to overfit specific training sets, achieving high accuracy on these trained video categories while failing to predict answers in other categories. However, in practical applications, it is essential for these methods to generalize their abilities across different domains. Following the experimental settings established by (Zhao et al. 2022), we define two task splits to validate generalization capabilities. As illustrated in Tab. 2, our method achieves consistent improvements in accuracy and ANLS compared to both visual-language pre-training methods and Video-LLMs. We attribute it in two factors. With the incorporation of pretrained language model (Raffel et al. 2019), TEA exploits its world knowledge in scene text descriptions. Additionally, our devised spatio-temporal recovery for vi-

sual entities module explicitly establishes spatio-temporal relationships among visual entities, which are shared across various scenarios, including even in unseen video categories.

## Ablation Studies

**Overall Results** To empirically validate the contribution of key components in TEA, i.e., spatio-temporal recovery for visual entities and scene text aware clues aggregation, we conduct ablation studies on the M4-ViteVQA *Task1Split1* validation set. The corresponding modules are individually added, and their impact on the model’s performance is measured. The overall comparative results are detailed in Tab. 3 and the specific hyperparameters for each module will be discussed in the following sections. Firstly, we denote TEA-Baseline, where multimodal features for scene text and objects across video frames are connected into a sequence and fed into the foundational model T5. This approach disrupts the spatial and temporal relationships between visual entities (#a). When the temporal convolution module, which provides temporal relationships between the same instance

#	Temporal	Spatial	Clue	M4-ViteVQA	
				Acc.(%)	ANLS(%)
a				26.79	35.16
b	✓			28.56	37.51
c		✓		31.91	40.93
d	✓	✓		33.69	42.91
e			✓	29.22	38.06
f	✓	✓	✓	<b>34.45</b>	<b>43.35</b>

Table 3: Overall ablation results on the M4-ViteVQA *Task1Split1* validation set. Temporal, Spatial, Clue represents temporal convolution module, OCR-enhanced relative spatial module, and scene text aware clues aggregation, respectively.

Multi-granularity	Acc.(%)	ANLS(%)
✓	31.06	40.02
	<b>31.91</b>	<b>40.93</b>

Table 4: Ablation experiments about OCR-enhanced relative spatial module on M4-ViteVQA *Task1Split1* validation dataset.

across video frames, is added, there is a 1.77% increase in accuracy (#b). In contrast, the OCR-enhanced relative spatial module results in a more substantial improvement of 5.12% in accuracy (#c). Subsequently, jointly recovering the spatio-temporal relationships leads to an advancement of 6.90% in accuracy and 7.75% in ANLS (#d). To demonstrate the ability of aggregating fine-grained features, we ablate the scene text aware clues aggregation module in two settings: without and with spatio-temporal features. Specifically, there is a 2.43% accuracy increase in the former setting (#e) and a 0.76% accuracy increase in the latter (#f).

### The Effect of OCR-enhanced Relative Spatial Module

The goal of the OCR-enhanced relative spatial module is to establish spatial relationships between visual entities, particularly for scene text. By providing word-level, line-level, and paragraph-level spatial relationships, the model is guided to recognize which scene text contains relevant semantic information. In Tab. 4, replacing these multi-granularity relationships with only word-level ones will lead to a decrease of 0.85% in accuracy.

### The Effect of Scene Text Aware Clues Aggregation

To assess the effectiveness of this module, we explore the features that need to be aggregated in Tab. 5. We can notice that compared to the initial version with only global visual semantic features (#a), integrating the question brings a 0.61% accuracy improvement (#b). This validates the pivotal role of questions in solving the Video TextVQA task, while the perception of fine-grained features is still insufficient. For this reason, by further incorporating the scene text and object features, a significant increase of 1.06% is observed (#c). Furthermore, we find that composing fine-grained features with scene text alone leads to an additional 0.76% improvement (#d). This might suggest that visual ob-

#	Features	Acc.(%)	ANLS(%)
a	–	26.79	35.16
b	Ques	27.40	36.47
c	Ques + OCR + OBJ	28.46	37.21
d	Ques + OCR	<b>29.22</b>	<b>38.06</b>

Table 5: Ablation experiments about scene text aware clues aggregation on M4-ViteVQA *Task1Split1* validation dataset.

Question: What is written on the third line of on the second sign?

Baseline: **Retral**; VideoLLaMA2: **Airport**; Ours/GT: **West Cargo Area**



Question: How much is the cheapest milk?

Baseline: **a2**; VideoLLaMA2: **2.49**; Ours/GT: **\$2**



Figure 4: A case study on M4-ViteVQA.

jects aid in comprehending the relative location and related semantics of scene text, yet they are ineffective in facilitating clues aggregation.

## Qualitative Analysis

**Case Study** In Fig. 4, a representative qualitative example is presented to provide an intuitive analysis of TEA in comparison with TEA-Baseline and VideoLLaMA2. We observe that TEA offers advantages in the following aspects: 1) Spatial-temporal relationships for scene text. TEA determines the relative spatial relationships between the first sign and the scene text associated with the second sign in the first frame. Additionally, since the scene text on the second sign in the first frame is not clearly visible, TEA establishes temporal relationships between the instances of the scene text “West Cargo Area” to enhance its visibility. 2) Aggregation for the essential clues. TEA identifies the third frame is not related to the question and aggregates a clue to assist in answering.

## Conclusion

In this paper, we present an initial investigation into extending Image TextVQA models to Video TextVQA methods through a simple yet highly effective approach, referred to as TEA. TEA aims to recover spatio-temporal relationships between scene texts across different video frames. We aggregate scene text aware clues to guide the model in identifying which frames are relevant to the questions. The capability of TEA to read and comprehend scene texts in videos also contributes to the advancement of understanding dynamic environments within the research community.

## Acknowledgments

Supported by the National Natural Science Foundation of China (Grant NO 62376266 and 62406318), and by the Key Research Program of Frontier Sciences, CAS (Grant NO ZDBS-LY-7024).

## References

- Biten, A. F.; Litman, R.; Xie, Y.; Appalaraju, S.; and Manmatha, R. 2021. LaTr: Layout-aware Transformer for Scene-text VQA. *CVPR*, 16527–16537.
- Cheng, Z.; Leng, S.; Zhang, H.; Xin, Y.; Li, X.; Chen, G.; Zhu, Y.; Zhang, W.; Luo, Z.; Zhao, D.; and Bing, L. 2024. VideoLLaMA2: Advancing Spatial-temporal Modeling and Audio Understanding in Video-LLMs. *ArXiv*, abs/2406.07476.
- Du, Y.; Chen, Z.; Su, Y.; Jia, C.; and Jiang, Y.-G. 2024. Instruction-guided Scene Text Recognition. *arXiv preprint arXiv:2401.17851*.
- Fang, C.; Li, J.; Li, L.; Ma, C.; and Hu, D. 2023. Separate and Locate: Rethink the Text in Text-based Visual Question Answering. *ACM MM*, 4378–4388.
- Fang, S.; Xie, H.; Wang, Y.; Mao, Z.; and Zhang, Y. 2021. Read Like Humans: Autonomous, Bidirectional and Iterative Language Modeling for Scene Text Recognition. *CVPR*, 7094–7103.
- Gao, D.; Li, K.; Wang, R.; Shan, S.; and Chen, X. 2020. Multi-modal Graph Neural Network for Joint Reasoning on Vision and Scene Text. In *CVPR*, 12746–12756.
- Gao, D.; Zhou, L.; Ji, L.; Zhu, L.; Yang, Y.; and Shou, M. Z. 2022. MIST: Multi-modal Iterative Spatial-temporal Transformer for Long-form Video Question Answering. *CVPR*, 14773–14783.
- Gao, J.; Ge, R.; Chen, K.; and Nevatia, R. 2018. Motion-appearance Co-memory Networks for Video Question Answering. *CVPR*, 6576–6585.
- Hong, T.; Kim, D.; Ji, M.; Hwang, W.; Nam, D.; and Park, S. 2021. BROS: A Pre-trained Language Model Focusing on Text and Layout for Better Key Information Extraction from Documents. In *AAAI*, 10767–10775.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. LoRA: Low-rank Adaptation of Large Language Models. In *ICLR*.
- Hu, R.; Singh, A.; Darrell, T.; and Rohrbach, M. 2019. Iterative Answer Prediction with Pointer-augmented Multimodal Transformers for TextVQA. *CVPR*, 9989–9999.
- Kant, Y.; Batra, D.; Anderson, P.; Schwing, A.; Parikh, D.; Lu, J.; and Agrawal, H. 2020. Spatially Aware Multimodal Transformers for TextVQA. In *ECCV*, 715–732.
- Le, T. M.; Le, V.; Venkatesh, S.; and Tran, T. 2020. Hierarchical Conditional Relation Networks for Video Question Answering. *CVPR*, 9969–9978.
- Lei, J.; Berg, T.; and Bansal, M. 2023. Revealing Single Frame Bias for Video-and-language Learning. In *ACL*, 487–507.
- Li, Z.; Shu, Y.; Zeng, W.; Yang, D.; and Zhou, Y. 2024. First Creating Backgrounds then Rendering Texts: A New Paradigm for Visual Text Blending. In *ECAI*.
- Lin, B.; Zhu, B.; Ye, Y.; Ning, M.; Jin, P.; and Yuan, L. 2023. Video-LLaVA: Learning United Visual Representation by Alignment before Projection. *ArXiv*, abs/2311.10122.
- Lin, J.; Yin, H.; Ping, W.; Molchanov, P.; Shoeybi, M.; and Han, S. 2024. VILA: On Pre-training for Visual Language Models. In *CVPR*, 26689–26699.
- Qiao, Z.; Zhou, Y.; Wei, J.; Wang, W.; Zhang, Y.; Jiang, N.; Wang, H.; and Wang, W. 2021. PIMNet: A Parallel, Iterative and Mimicking Network for Scene Text Recognition. In *ACM MM*, 2046–2055.
- Qiao, Z.; Zhou, Y.; Yang, D.; Zhou, Y.; and Wang, W. 2020. SEED: Semantics Enhanced Encoder-decoder Framework for Scene Text Recognition. In *CVPR*, 13528–13537.
- Qin, X.; Lyu, P.; Zhang, C.; Zhou, Y.; Yao, K.; Zhang, P.; Lin, H.; and Wang, W. 2023. Towards Robust Real-time Scene Text Detection: From Semantic to Instance Representation Learning. In *ACM MM*, 2025–2034.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models from Natural Language Supervision. In *ICML*.
- Raffel, C.; Shazeer, N. M.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2019. Exploring the Limits of Transfer Learning with a Unified Text-to-text Transformer. *JMLR*, 21: 140:1–140:67.
- Ren, S.; He, K.; Girshick, R. B.; and Sun, J. 2015. Faster R-CNN: Towards Real-time Object Detection with Region Proposal Networks. *TPAMI*, 39: 1137–1149.
- Singh, A.; Natarajan, V.; Shah, M.; Jiang, Y.; Chen, X.; Batra, D.; Parikh, D.; and Rohrbach, M. 2019. Towards VQA Models that Can Read. In *CVPR*, 8317–8326.
- Tom, G.; Mathew, M.; Garcia, S.; Karatzas, D.; and Jawahar, C. 2023. Reading Between the Lanes: Text VideoQA on the Road. In *ICDAR*, 137–154.
- Vaswani, A. 2017. Attention is All You Need. In *NeurIPS*.
- Wang, A.; Ge, Y.; Yan, R.; Ge, Y.; Lin, X.; Cai, G.; Wu, J.; Shan, Y.; Qie, X.; and Shou, M. Z. 2022a. All in One: Exploring Unified Video-language Pre-training. *CVPR*, 6598–6608.
- Wang, H.; Ye, Y.; Wang, Y.; Nie, Y.; and Huang, C. 2025. Elysium: Exploring Object-level Perception in Videos via MLLM. In *ECCV*, 166–185.
- Wang, J.; Yang, Z.; Hu, X.; Li, L.; Lin, K.; Gan, Z.; Liu, Z.; Liu, C.; and Wang, L. 2022b. GIT: A Generative Image-to-text Transformer for Vision and Language. *TMLR*.
- Wang, W.; Zhou, Y.; Lv, J.; Wu, D.; Zhao, G.; Jiang, N.; and Wang, W. 2022c. TPSNet: Reverse Thinking of Thin Plate Splines for Arbitrary Shape Scene Text Representation. In *ACM MM*, 5014–5025.
- Wu, W.; Cai, Y.; Zhang, D.; Wang, S.; Li, Z.; Li, J.; Tang, Y.; and Zhou, H. 2021. A Bilingual, OpenWorld Video

Text Dataset and End-to-end Video Text Spotter with Transformer. In *NeurIPS*.

Xiao, J.; Yao, A.; Liu, Z.; Li, Y.; Ji, W.; and seng Chua, T. 2021. Video as Conditional Graph Hierarchy for Multi-granular Question Answering. In *AAAI*, 2804–2812.

Yang, A.; Miech, A.; Sivic, J.; Laptev, I.; and Schmid, C. 2020. Just Ask: Learning to Answer Questions from Millions of Narrated Videos. *ICCV*, 1666–1677.

Yang, Z.; Lu, Y.; Wang, J.; Yin, X.; Florencio, D.; Wang, L.; Zhang, C.; Zhang, L.; and Luo, J. 2021. TAP: Text-aware Pre-training for TextVQA and Text Caption. In *CVPR*, 8751–8761.

Ye, M.; Zhang, J.; Liu, J.; Liu, C.; Yin, B.; Liu, C.; Du, B.; and Tao, D. 2024. Hi-SAM: Marrying Segment Anything Model for Hierarchical Text Segmentation. *TPMAI*.

Zeng, G.; Zhang, Y.; Wei, J.; Yang, D.; Zhang, P.; Gao, Y.; Qin, X.; and Zhou, Y. 2024a. Focus, Distinguish, and Prompt: Unleashing CLIP for Efficient and Flexible Scene Text Retrieval. In *ACM MM*, 2525–2534.

Zeng, G.; Zhang, Y.; ZHOU, Y.; Fang, B.; Zhao, G.; Wei, X.; and Wang, W. 2023a. Filling in the Blank: Rationale-augmented Prompt Tuning for TextVQA. *ACM MM*, 1261–1272.

Zeng, G.; Zhang, Y.; Zhou, Y.; Yang, X.; Jiang, N.; Zhao, G.; Wang, W.; and Yin, X.-C. 2023b. Beyond OCR+VQA: Towards End-to-end Reading and Reasoning for Robust and Accurate TextVQA. *PR*, 138: 109337.

Zeng, W.; Shu, Y.; Li, Z.; Yang, D.; and Zhou, Y. 2024b. TextCtrl: Diffusion-based Scene Text Editing with Prior Guidance Control. In *NeurIPS*.

Zhang, Y.; Zeng, G.; Shen, H.; Ma, C.; and Zhou, Y. 2024. Show Exemplars and Tell Me What You See: In-context Learning with Frozen Large Language Models for TextVQA. In *PRCV*, 231–245.

Zhao, M.; Li, B.; Wang, J.; Li, W.; Zhou, W.; Zhang, L.; Xuyang, S.; Yu, Z.; Yu, X.; Li, G.; Dai, A.; and Zhou, S. 2022. Towards Video Text Visual Question Answering: Benchmark and Baseline. In *NeurIPS*, 35549–35562.

Zheng, T.; Chen, Z.; Fang, S.; Xie, H.; and Jiang, Y.-G. 2024. CDistNet: Perceiving Multi-domain Character Distance for Robust Text Recognition. *IJCV*, 132(2): 300–318.

Zhu, Q.; Gao, C.; Wang, P.; and Wu, Q. 2021. Simple is Not Easy: A Simple Strong Baseline for TextVQA and TextCaps. In *AAAI*, volume 35, 3608–3615.