

Understanding Video Scenes through Text: Insights from Text-based Video Question Answering

Soumya Jahagirdar¹Minesh Mathew²Dimosthenis Karatzas³C. V. Jawahar¹

soumya.jahagirdar@research.iiit.ac.in

minesh@wadhwaniai.org

dimos@cvc.uab.es

jawahar@iiit.ac.in

¹ CVIT, IIIT Hyderabad, India² Wadhvani AI³ Computer Vision Center, UAB, Spain

Abstract

Researchers have extensively studied the field of vision and language, discovering that both visual and textual content is crucial for understanding scenes effectively. Particularly, comprehending text in videos holds great significance, requiring both scene text understanding and temporal reasoning. This paper focuses on exploring two recently introduced datasets, NewsVideoQA and M4-ViteVQA, which aim to address video question answering based on textual content. The NewsVideoQA dataset contains question-answer pairs related to the text in news videos, while M4-ViteVQA comprises question-answer pairs from diverse categories like vlogging, traveling, and shopping. We provide an analysis of the formulation of these datasets on various levels, exploring the degree of visual understanding and multi-frame comprehension required for answering the questions. Additionally, the study includes experimentation with BERT-QA, a text-only model, which demonstrates comparable performance to the original methods on both datasets, indicating the shortcomings in the formulation of these datasets. Furthermore, we also look into the domain adaptation aspect by examining the effectiveness of training on M4-ViteVQA and evaluating on NewsVideoQA and vice-versa, thereby shedding light on the challenges and potential benefits of out-of-domain training.

1. Introduction

Multimodal understanding, specifically, VideoQA is a challenging yet crucial problem that involves multimodal and temporal reasoning. Researchers have developed various datasets and methods to facilitate research in this field [8, 10, 13, 17, 19–21]. Xu et al. [20] and Yu et al. [21] propose datasets that contain questions about the events happening in the video but disregard the text. However, works such as Lei et al. [10] and Tapaswi et al. [17] have introduced datasets that use both visual and subtitle information to understand the story. However, these existing datasets



Figure 1. Example illustrating two major concerns of existing text-based VideoQA datasets [7, 22]. Both examples showcase that only **textual information** from a **single frame** is sufficient to obtain answers to the questions.

lacked the ability to handle questions that require reading text in videos. As textual content in man-made environments carries significant semantic information, the need for visual question-answering datasets that involve reading text became evident. These text-based systems have great potential for real-life scenarios, particularly for visually-impaired users and the development of assistive devices. While previous works have explored single-image scene-text and document images [1, 3, 6, 11, 12, 14, 16] there has been limited exploration of works that require extracting information from text present in videos. Recently, Hegde et al. [4] shed light on the bias aspects of TextVQA datasets. Recently, there have been multiple datasets [7, 18, 22] that deal with a new line of research wherein the datasets require models to read and understand the text present in the videos to answer the questions. Jahagirdar et al. [7] proposed a dataset: NewsVideoQA that contains question-answer pairs framed on news videos from multiple news channels, and these questions are formulated such that answering the questions requires an understanding of the embedded text i.e. text occurring in the news videos. Similarly, Zhao et al. [22] introduced a dataset that contains videos from multiple cat-

egories such as shopping, traveling, etc, and requires both temporal reasoning and textual reasoning to answer questions. Tom et al. [18], proposed a dataset for the task of video question answering in the context of driver assistance on road videos. Additionally, a competition ¹ centered on the task of answering questions based on video content using text was introduced.

Table 1. Analysis of 100 random QA pairs from M4-ViteVQA and NewsVideoQA datasets.

Category	M4-ViteVQA (%)	NewsVideoQA (%)
Single Frame	92.0	95.0
Multi Frame	8.0	5.0
Visual Info	33.0	6.0
Textual Info	95.0	100.0
Frame crowded with text	18.0	64.0
Extractive-based	81.0	98.0
Reasoning-based	5.0	2.0
Knowledge-based	1.0	0.0

In this work, we explore the task of text-based video question answering. Firstly, we study and analyze two recently introduced datasets, namely NewsVideoQA [7] and M4-ViteVQA [22], which includes various types of videos, such as news videos, vlogging, traveling and shopping. We conduct an exploratory analysis to examine the level of visual understanding and multi-frame comprehension required for answering the questions in both datasets. Additionally, we conduct experiments using BERT-QA [2], a text-only model, and demonstrate its effectiveness by achieving comparable results to the original methods that consider both visual and textual information. We also analyze domain adaptation by training on M4-ViteVQA and testing on NewsVideoQA, and vice-versa, revealing insights into cross-domain understanding challenges.

2. Benchmarking and Experiments

In this section, we present details of the exploratory analysis and the experiments we conduct. BERT-QA is a transformer-based encoder-only model pre-trained on a large corpus and further finetuned on SQuAD dataset [15] for question answering (Extractive QA). Extractive QA is the task of extracting a short snippet from the document/context on which the question is asked. The answer ‘span’ is determined by its start and end tokens. It is selected for its effective extractive QA performance, implementation ease, and finetuning, despite limitations like no answer generation or handling yes/no questions. Its ability in extracting answers from textual content makes it a suitable choice for tasks where answers are primarily found in the text of the video. To convert both M4-ViteVQA and NewsVideoQA

datasets in SQuAD format, we find the first substring of the answer in the context, which is an approximation of the answer span as followed in [12].

2.1. Exploratory Analysis

For exploratory analysis, we randomly sample 100 QA pairs from both M4-ViteVQA and NewsVideoQA. For each QA pair, we check the following aspects: i) if the question can be answered by a single frame or need multi-frame information, ii) if the question needs visual information and/or textual information to obtain the answer, iii) if the frame which is essential to obtain the answer, is crowded with text (approximately more than 15 OCR tokens). From Table. 1, it can be seen that for both datasets, information from a single frame is sufficient to obtain answers, which is counter-intuitive to the video question-answering task. From Table. 1, it can also be seen that most of the questions in both datasets need textual information to obtain answers. As M4-ViteVQA contains videos from multiple categories, it contains more questions of visual type compared to NewsVideoQA that contains only news videos. Since both datasets are designed for questions that require reading text to answer questions, this has resulted in minimal questions that require multimodal information. We also check for the answer type: i) extractive, ii) reasoning based, and iii) knowledge-based, and combinations of each type. From Table. 1, it can be seen that most of the questions are extractive in nature and have fewer reasoning-based and knowledge-based questions. However, having more reasoning/knowledge-based questions is crucial, thereby creating the need for better methods beyond the scope of text-only models.

2.2. BERT-QA experiments

M4-ViteVQA [22]: The M4-ViteVQA dataset consists of two tasks. The first task is divided into two splits and both splits contain evenly distributed question-answer pairs from all video categories in train-val-test sets. In the second task, the training set comprises videos from seven categories, while the question-answer pairs and videos in validation in test splits are exclusively sourced from the remaining two categories. Zhao et al. [22] also propose a multimodal video question-answering method: T5-ViteVQA, that combines information from multiple modalities including OCR features, question features, and video features.

In our experiments on the BERT-QA model, we first sample frames at 1fps and order the OCR tokens of the frames to the default reading order based on the position of the top-left corner of the OCR token. We further concatenate the ordered OCR tokens which becomes the context of the BERT-QA model. After the training phase, we conduct two types of testing to evaluate the performance of the BERT-QA model. For the first type, we evaluate the model

¹<https://tianchi.aliyun.com/competition/entrance/532050/information>

Table 2. Performance comparison of different methods on the validation set of M4-ViteVQA dataset. It can be seen that a simple text-only model achieves comparable results and beats the scores of the original multimodal method.

Experiment	Finetuning	Task1 Split1		Task1 Split2		Task2	
		ACC.	ANLS	ACC.	ANLS	ACC.	ANLS
T5-ViteVQA	✓	23.17	30.10	17.59	23.10	12.30	16.10
BERT-QA	×	9.03	17.05	8.17	15.81	10.89	18.41
BERT-QA	✓	21.96	32.18	17.10	26.05	16.01	24.08

Table 3. Performance comparison of BERT-QA model on M4-ViteVQA dataset when the answer to questions is present in the concatenated list of OCR tokens from evenly sampled frames.

Split	Acc.	ANLS	No. of QA pairs
Task 1 Split 1	47.42	55.14	911
Task 1 Split 2	42.29	48.90	532
Task 2	38.05	43.82	318

on the entire validation set without checking if the answer is present in the context. This experiment allows us to assess the model’s overall ability to obtain answers. In the second type of testing, we specifically focus on questions that have answers in the context.

NewsVideoQA [7]: This dataset proposes questions on news videos. The dataset has timestamps for each question indicating the frame at which the question was defined. This work also proposes a repurposed baseline: OCR-aware SINGULARITY, which was originally inspired by SINGULARITY [9]. OCR-aware SINGULARITY is a multi-modal transformer-based video question-answering model that combines information from OCR tokens, questions, and visual information from a randomly sampled frame.

In this work, we conduct two types of training on this dataset. In the first approach, we train the BERT-QA model using the OCR tokens of the single frame on which the question was defined (BERT-QA-SF: BERT-QA Single Frame). In the second approach, we concatenate the OCR tokens from frames sampled at 1fps which forms the context of the BERT-QA model. (BERT-QA-MF: Multi-frame). By conducting training in both single-frame (BERT-QA-SF) and multi-frame (BERT-QA-MF) setups, we aim to explore the impact of variations in the length of context on the performance of the BERT-QA model. These two training approaches provide insights into the model’s ability to obtain answers based on either a specific frame or a broader contextual understanding derived from multiple frames.

2.3. Domain Adaptation Experiments

We conduct experiments to determine if the BERT-QA model can perform or generalize well with the out-of-domain context. This evaluation aims to determine if the model can provide accurate answers even in unfa-

Table 4. Performance comparison (Acc.) of M4C, T5-ViteVQA, and BERT-QA on the validation set of Task 1 Split 1.

Set	M4C [5]	T5-ViteVQA [22]	BERT-QA [2]
Easy	19.30	25.09	25.49
Hard	9.02	14.26	16.34
Text	17.26	23.08	31.01
Vision	18.36	24.21	18.82

miliar video categories and their corresponding contexts. To achieve this understanding, we perform several experiments. We check for the performance of the BERT-QA model trained on the **Source dataset** followed by testing on the **Target dataset**. We do this in two settings: i) without finetuning on the target dataset, and ii) with finetuning on the target dataset (Example: Train on NewsVideoQA and test on M4-ViteVQA in two settings i.e. with/without finetuning and vice-versa). By doing these, we try to examine the impact of domain shift and the importance of training the model on videos from diverse categories, where scene text serves as the textual content in one dataset that is M4-ViteVQA, as opposed to embedded text in NewsVideoQA. These experiments help us determine the model’s ability to generalize and adapt to the specific categories of videos.

2.4. Evaluation Metrics and Experimental Setup

Frequently used in the majority of the works on scene-text based visual and video question answering, we use two evaluation metrics — Accuracy (Acc.) and Average Normalized Levenshtein Similarity (ANLS). Accuracy is the percentage of questions for which correct answers and predicted answers match exactly. Whereas ANLS is a similarity-based metric that acts softly on minor answer mismatches. More details can be found in [1]. For all experiments, we train BERT-QA bert-large-uncased-whole-word-masking-finetuned-squad for 15 epochs with a batch-size of 16 on 4 GPUs with a learning rate of 2e-05.

2.5. Quantitative results

In this section, we present the results and analysis of different experiments. In Table. 2, we show the results of the performance of T5-ViteVQA and BERT-QA on dif-

Table 5. We show the performance of OCR-aware SINGULARITY [7] and BERT-QA in different settings. BERT-QA-SF: single frame setup, BERT-QA-MF: multi-frame setup on NewsVideoQA. In the second column, we explain the type of testing. “12 random frames”: considers visual and textual information from 12 random frames, “single random frame”: OCR tokens of a random frame, “single correct frame”: OCR tokens of the correct frame, “1 frame per second”: OCR tokens of frames sampled at 1fps.

Baseline	Type of testing	Acc.	ANLS
OCR-aware SINGULARITY	12 random frames	32.47	35.56
BERT-QA-SF	single random frame	23.71	29.47
BERT-QA-SF	single correct frame	46.55	56.81
BERT-QA-MF	1 frame per second	52.29	61.12

ferent tasks and splits on the validation set of the M4-ViteVQA dataset. It can be seen that a simple text-only model achieves comparable results and beats the scores of T5-ViteVQA for certain splits. The results indicate that we need more datasets that require information from multiple modalities and multiple frames which is a concerning limitation in the current datasets. It can be seen that the BERT-QA relies purely on the OCR output to infer and extract the answer. Therefore, if the OCR output is noisy or if the tokens are incorrectly ordered (errors in default reading order) the model might fail to find the right answer. However, since the ANLS metric acts softly on OCR errors, BERT-QA outperforms T5-ViteVQA on the ANLS metric. In Table. 3, we show the performance of BERT-QA for the questions that contain answers in the context. We create this test set by checking if the answer is a substring of context. For each of the splits, nearly half of the original questions in the validation set have answers in the context. In Table. 4, we show the performance comparison—in terms of Accuracy—of two methods: i) M4C [5]: It uses a multi-modal transformer and an iterative answer prediction module. The model answers questions based on scene-text questions on a single image. ii) T5-ViteVQA: method proposed as a baseline in [22], with BERT-QA on the validation set of Task 1 Split 1. It can be seen that BERT-QA outperforms M4C and T5-ViteVQA on different sets. Here, the “sets” correspond to the type of questions which is provided with the dataset. These sets are: i) easy - answering requires information from a single frame, ii) hard - answering requires information from multiple frames, iii) text - answering requires only reading text, and iv) vision - answering requires both visual and textual information. Only for questions that require visual information, BERT-QA underperforms, yet still manages to obtain decent performance.

In Table. 5, we show results of the performance of different methods on the test set of NewsVideoQA [7] dataset. OCR-aware SINGULARITY is a model trained in a single-frame setup and is tested on a multi-frame setup (by combining visual and textual information from 12 frames - more

Table 6. Out-of-domain training performance for NewsVideoQA and M4-ViteVQA datasets. “Source dataset” corresponds to the dataset on which we train the model, and “Target dataset” corresponds to the dataset we test the model on.

Source dataset	Target dataset	Finetuning on target	Acc.	ANLS
NewsVideoQA	NewsVideoQA	✓	52.29	61.12
M4-ViteVQA	NewsVideoQA	×	40.39	51.86
M4-ViteVQA	NewsVideoQA	✓	50.41	61.04
M4-ViteVQA	M4-ViteVQA	✓	21.96	32.18
NewsVideoQA	M4-ViteVQA	×	7.86	12.68
NewsVideoQA	M4-ViteVQA	✓	22.17	31.95

details in [7]). This is followed by results of BERT-QA-SF i.e. trained on OCR context from a single frame and tested by picking a random frame. In the third row, we show the results of BERT-QA when tested with OCR tokens of the frame on which the question was defined (correct frame). In the fourth row, BERT-QA-MF: BERT-QA is trained and tested on a multi-frame setup. In Table. 6, we show the results of out-of-domain training performance on both [7, 22] datasets. It can be seen that testing a model initially trained on M4-ViteVQA (Source dataset) achieves decent performance on an out-of-domain NewsVideoQA (target dataset) and vice-versa. By further finetuning on the target dataset, the performance of the model increases. This indicates that the BERT-QA model can effectively generalize across domains through out-of-domain training. More details are present in the supplementary.

3. Conclusion

This paper focused on the important task of understanding textual information within videos for question-answering. The study provides insights that current text-based VideoQA datasets majorly focus on extractive answers and provide insights that the degree of visual understanding and multi-frame comprehension in current datasets is limited for better VideoQA using text in videos. Additionally, the paper demonstrates the effectiveness of BERT-QA, a text-only model, in achieving comparable performance to original methods on both datasets and also looks into the domain transfer aspect, by comparing the performances by training on one type of dataset and testing on the other. In future developments, we hope to see datasets that prioritize non-extractive answers and incorporate multimodal questions based on multiple frames to facilitate improved multimodal learning.

Acknowledgements. This work is supported by MeitY, Government of India.

References

- [1] Ali Furkan Biten, Rubèn Tito, Andrés Mafla, Lluís Gómez i Bigorda, Marçal Rusiñol, C. V. Jawahar, Ernest Valveny, and

- Dimosthenis Karatzas. Scene text visual question answering. In *ICCV*, pages 4290–4300. IEEE, 2019.
- [2] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*, pages 4171–4186. Association for Computational Linguistics, 2019.
- [3] Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P. Bigham. Vizwiz grand challenge: Answering visual questions from blind people. In *CVPR*, pages 3608–3617. Computer Vision Foundation / IEEE Computer Society, 2018.
- [4] Shamanthak Hegde, Soumya Jahagirdar, and Shankar Gangisetty. Making the V in text-vqa matter. In *CVPR Workshops*, pages 5580–5588. IEEE, 2023.
- [5] Ronghang Hu, Amanpreet Singh, Trevor Darrell, and Marcus Rohrbach. Iterative answer prediction with pointer-augmented multimodal transformers for textvqa. In *CVPR*, pages 9989–9999. Computer Vision Foundation / IEEE, 2020.
- [6] Soumya Jahagirdar, Shankar Gangisetty, and Anand Mishra. Look, read and ask: Learning to ask questions by reading text in images. In *ICDAR (1)*, volume 12821 of *Lecture Notes in Computer Science*, pages 335–349. Springer, 2021.
- [7] Soumya Jahagirdar, Minesh Mathew, Dimosthenis Karatzas, and C. V. Jawahar. Watching the news: Towards videoqa models that can read. In *WACV*, pages 4430–4439. IEEE, 2023.
- [8] Yunseok Jang, Yale Song, Youngjae Yu, Youngjin Kim, and Gunhee Kim. TGIF-QA: toward spatio-temporal reasoning in visual question answering. In *CVPR*, pages 1359–1367. IEEE Computer Society, 2017.
- [9] Jie Lei, Tamara L. Berg, and Mohit Bansal. Revealing single frame bias for video-and-language learning. In *ACL (1)*, pages 487–507. Association for Computational Linguistics, 2023.
- [10] Jie Lei, Licheng Yu, Mohit Bansal, and Tamara L. Berg. TVQA: localized, compositional video question answering. In *EMNLP*, pages 1369–1379. Association for Computational Linguistics, 2018.
- [11] Minesh Mathew, Viraj Bagal, Rubèn Tito, Dimosthenis Karatzas, Ernest Valveny, and C. V. Jawahar. Infographicvqa. In *WACV*, pages 2582–2591. IEEE, 2022.
- [12] Minesh Mathew, Dimosthenis Karatzas, and C. V. Jawahar. Docvqa: A dataset for VQA on document images. In *WACV*, pages 2199–2208. IEEE, 2021.
- [13] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *ICCV*, pages 2630–2640. IEEE, 2019.
- [14] Anand Mishra, Shashank Shekhar, Ajeet Kumar Singh, and Anirban Chakraborty. OCR-VQA: visual question answering by reading text in images. In *ICDAR*, pages 947–952. IEEE, 2019.
- [15] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100, 000+ questions for machine comprehension of text. In *EMNLP*, pages 2383–2392. The Association for Computational Linguistics, 2016.
- [16] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards VQA models that can read. In *CVPR*, pages 8317–8326. Computer Vision Foundation / IEEE, 2019.
- [17] Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhagen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. Movieqa: Understanding stories in movies through question-answering. In *CVPR*, pages 4631–4640. IEEE Computer Society, 2016.
- [18] George Tom, Minesh Mathew, Sergi Garcia, Dimosthenis Karatzas, and C. V. Jawahar. Reading between the lanes: Text videoqa on the road, 2023.
- [19] Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. Video question answering via gradually refined attention over appearance and motion. In *ACM Multimedia*, pages 1645–1653. ACM, 2017.
- [20] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. MSR-VTT: A large video description dataset for bridging video and language. In *CVPR*, pages 5288–5296. IEEE Computer Society, 2016.
- [21] Zhou Yu, Dejing Xu, Jun Yu, Ting Yu, Zhou Zhao, Yueting Zhuang, and Dacheng Tao. Activitynet-qa: A dataset for understanding complex web videos via question answering. In *AAAI*, pages 9127–9134. AAAI Press, 2019.
- [22] Minyi Zhao, Bingjia Li, Jie Wang, Wanqing Li, Wenjing Zhou, Lan Zhang, Shijie Xuyang, Zhihang Yu, Xinkun Yu, Guangze Li, Aobotao Dai, and Shuigeng Zhou. Towards video text visual question answering: Benchmark and baseline. In *NeurIPS*, 2022.