



**HAL**  
open science

# Learning to Answer Visual Questions from Web Videos

Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, Cordelia Schmid

► **To cite this version:**

Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, Cordelia Schmid. Learning to Answer Visual Questions from Web Videos. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2022, 10.1109/tpami.2022.3173208 . hal-03664182

**HAL Id: hal-03664182**

**<https://inria.hal.science/hal-03664182v1>**

Submitted on 10 May 2022

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Learning to Answer Visual Questions from Web Videos

Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, Cordelia Schmid

**Abstract**—Recent methods for visual question answering rely on large-scale annotated datasets. Manual annotation of questions and answers for videos, however, is tedious, expensive and prevents scalability. In this work, we propose to avoid manual annotation and generate a large-scale training dataset for video question answering making use of automatic cross-modal supervision. We leverage a question generation transformer trained on text data and use it to generate question-answer pairs from transcribed video narrations. Given narrated videos, we then automatically generate the HowToVQA69M dataset with 69M video-question-answer triplets. To handle the open vocabulary of diverse answers in this dataset, we propose a training procedure based on a contrastive loss between a video-question multi-modal transformer and an answer transformer. We introduce the zero-shot VideoQA task and the VideoQA feature probe evaluation setting and show excellent results, in particular for rare answers. Furthermore, our method achieves competitive results on MSRVT-QA, ActivityNet-QA, MSVD-QA and How2QA datasets. We also show that our VideoQA dataset generation approach generalizes to another source of web video and text data. We use our method to generate the WebVidVQA3M dataset from the WebVid dataset, i.e., videos with alt-text annotations, and show its benefits for training VideoQA models. Finally, for a detailed evaluation we introduce iVQA, a new VideoQA dataset with reduced language bias and high-quality manual annotations. Code, datasets and trained models are available on our project webpage<sup>1</sup>.

**Index Terms**—Video Question Answering, Cross-Modal Supervision, Question Generation, Zero-Shot Learning

## 1 INTRODUCTION

Answering questions about videos requires a detailed understanding of the visual content and its association with the natural language. Indeed, given the large diversity of questions, methods for Video Question Answering (VideoQA) should reason about scenes, objects and human actions as well as their complex temporal interactions.

Current approaches to VideoQA rely on deep fully-supervised models trained on manually annotated datasets with question and answer pairs [28], [38], [42], [43], [50], [52], [58]. Collecting and annotating VideoQA datasets, however, is cumbersome, time consuming, expensive and therefore not scalable. As a result, current VideoQA datasets are relatively small (see Figure 2). This limitation hinders the progress in the field as state-of-the-art VideoQA models often require a large amount of training data.

In this work, we address the scale issue with a new approach for automatically generating VideoQA datasets as illustrated in Figure 1. The idea is to leverage cross-modal supervision together with text-only tools for question generation and to automatically annotate VideoQA data from a *large amount of videos with readily-available text annotations* in the form of transcribed narrations or “alt-text” annotations available with the video on the Internet. Inspired by the recent progress in language generation using transformer-based language models [12], we leverage transformers trained on a question-answering text corpus to generate

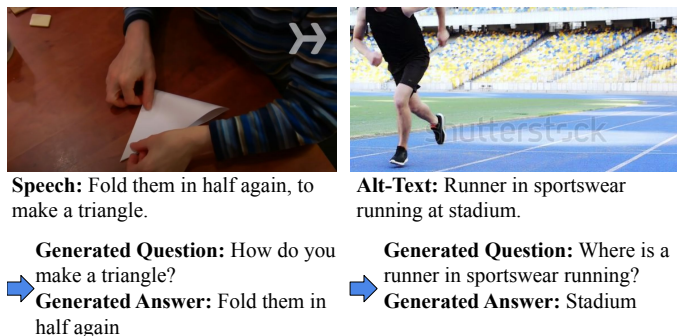


Fig. 1: Given videos with transcribed narration (left) or videos with “alt-text” annotations (right), we leverage language models and cross-modal supervision to obtain large-scale VideoQA data. Top: Example frame with the corresponding text annotation. Bottom: automatically generated question and answer pair.

a diverse set of non-scripted questions and corresponding open-vocabulary answers from text. By applying these transformers to speech transcripts of narrated videos from the large-scale HowTo100M dataset [69] we create HowToVQA69M, an open-ended VideoQA dataset with 69 million video-question-answer triplets and a diverse set of more than 16M unique answers (see Figure 3). We also extend our approach to web videos with readily-available alt-text descriptions and generate the WebVidVQA3M dataset from the WebVid2M dataset [9]. As shown in Figure 2, our HowToVQA69M and WebVidVQA3M datasets are orders of magnitude larger compared to prior VideoQA datasets.

Given the limited diversity of existing datasets, current methods typically reduce VideoQA to a classification problem, where frequent answers are assigned to unique classes. Typically, up to 5K unique possible answers are considered. Such an approach, however, does not scale to the open vocabulary of 16M different

- A. Yang, I. Laptev and C. Schmid are with Inria Paris and Département d’informatique de l’ENS, CNRS, PSL Research University, 75005 Paris, France.
- A. Miech is with DeepMind, London, United Kingdom.
- J. Sivic is with Czech Institute of Informatics, Robotics and Cybernetics, Czech Technical University in Prague.

1. <https://antoyang.github.io/just-ask.html>

answers in HowToVQA69M. To address this problem and to enable video question answering with highly diverse questions and answers, we introduce a training procedure based on contrastive learning between a video-question multi-modal transformer and an answer transformer that can handle free-form answers. This bypasses the need to define a discrete set of answer classes.

The goal of our work is to advance truly open-ended and generic solutions to VideoQA. To evaluate generalization, we propose a new zero-shot VideoQA task where we prohibit any manual supervision of visual data during training, and a new VideoQA feature probe evaluation setting where only the final projection layers of the network are finetuned on the target dataset. Our VideoQA model, trained on HowToVQA69M, demonstrates excellent zero-shot results on multiple existing datasets, especially for rare answers. Additionally, we find that our VideoQA model exhibits strong performance in the VideoQA feature probe evaluation setting. Moreover, when finetuned on target datasets, our model achieves competitive results on MSRVTT-QA [100], ActivityNet-QA [110], MSVD-QA [100] and How2QA [56]. We further show the generalizability of our approach by showing the benefits of WebVidVQA3M for training VideoQA models.

Initial experiments have shown that existing benchmarks for open-ended VideoQA [100], [110] contain a language bias [34], i.e., their questions can often be answered without looking at the video. To better evaluate the impact of visual information in VideoQA, we introduce a new open-ended VideoQA dataset (iVQA) with manually collected questions and answers, where we exclude questions that could be answered without watching the video. Moreover, to account for multiple possible answers, iVQA contains five independently collected answers for each question.

In summary, our work makes the following three contributions:

- (i) We introduce an approach to automatically generate a large-scale VideoQA dataset, HowToVQA69M. Relying on cross-modal supervision, we use transformers trained on an existing text-only question-answering corpus and generate video-question-answer triplets from videos and transcribed narrations. We also apply our method to video alt-text pairs and generate the WebVidVQA3M dataset.
- (ii) We train a VideoQA model on the automatically generated data via contrastive learning between a multi-modal video-question transformer and an answer transformer. We show the efficiency of our model for the new zero-shot VideoQA task and the new VideoQA feature probe task. Our model achieves competitive results in four existing VideoQA benchmarks.
- (iii) Finally, we introduce a new manually annotated open-ended VideoQA benchmark iVQA that excludes non-visual questions and contains multiple possible answers for each question.

Code, datasets and trained models are available at <https://antoyang.github.io/just-ask.html>.

## 2 RELATED WORK

**Visual Question Answering (VQA).** VQA is typically tackled by classifying the image-question (or video-question) representation into a fixed vocabulary of answers. Various approaches to combine spatial image representations and sequential question representations have been proposed [6], [11], [30], [66], [99], [101], [106]. More specifically to the video domain (VideoQA), spatio-temporal video representations in terms of motion and appearance have been

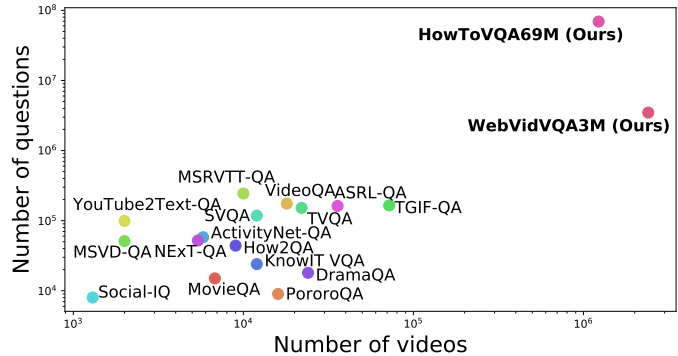


Fig. 2: Comparison of our large-scale HowToVQA69M and WebVidVQA3M datasets with existing VideoQA datasets.

used in [23], [28], [32], [38], [41], [42], [43], [50], [51], [52], [58], [72], [79], [100], [102], [109], [114], [122].

Methods above are limited to pre-defined vocabularies of answers and are difficult to apply outside of specific datasets. To address this problem, Hu *et al.* [37] propose a joint embedding where image-question representations can be matched with free-form answers. Our VideoQA model follows this idea, but instead of relying on manually annotated datasets of limited scale, we train it on a large-scale VideoQA dataset that we automatically generate. In contrast to some previous works using additional video features such as subtitles [16], [45], [46], [48], [53], [54], [56], [62], [89], [95], [105], our video representation is exclusively based on visual information, as we focus on the detailed visual understanding of videos.

To evaluate the generalization of VQA models, Teney and Hengel [90] define zero-shot VQA by answering previously unseen questions, which is a related but less challenging task compared to the zero-shot VQA task we propose in Section 6.2. Vatashsky and Ullman [93] address VQA using COCO image annotations [61], while our zero-shot model is trained with no manual annotations. Our proposed zero-shot VQA task is analogous to zero-shot video retrieval [68] or zero-shot action recognition [73]. We further propose a VQA feature probe evaluation setting where only the final heads of the network are finetuned on the downstream dataset while all other pretrained weights are kept frozen. This setting is analogous to the linear probe evaluation setting commonly used in self-supervised image recognition [13], [14], [19] or self-supervised action recognition [73] but with multiple layers in the head rather than just a single (linear) layer.

Visual question generation (VQG) has been introduced in [70]. The methods in [60] and [81] propose to jointly learn VQG and VQA to improve the image VQA task. However, these works do not generate questions to obtain additional training data, but use visual data annotation for VQG as an additional loss.

**VideoQA datasets.** Manually collecting and annotating video-question-answer triplets is cumbersome, costly and difficult to scale. As a result, current VideoQA datasets [15], [21], [22], [27], [33], [41], [47], [53], [56], [71], [77], [83], [89], [97], [100], [108], [110], [111], [113] are limited in size, as the largest, TGIF-QA [41], contains only 72K annotated clips (see Figure 2 for more details). To address this issue, several works have explored leveraging manually annotated video descriptions [41], [94], [100], [113], [115], [116], [117] for automatic generation of VideoQA datasets, using rule-based [35], [76] approaches. Similarly, in the image domain, Banerjee *et al.* [10] has recently proposed to use annotated image captions from COCO [18] to generate question-

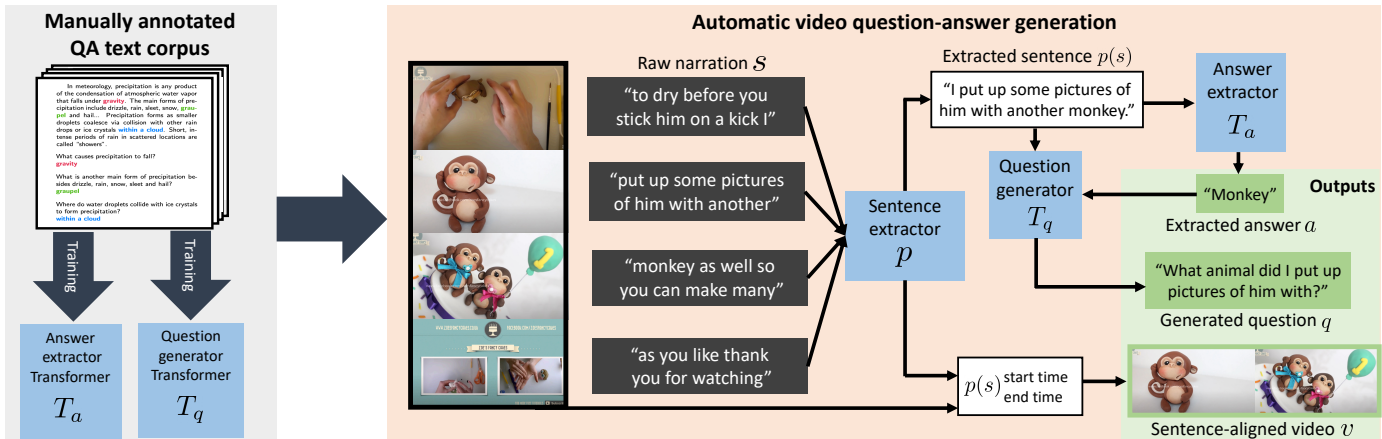


Fig. 3: **Our automatic approach for large-scale generation of video-question-answer triplets from narrated (subtitled) videos.** First, at the language-only training phase (left), the transformer-based answer extractor  $T_a$  and question generator  $T_q$  are trained [74] on a manually annotated text-only question-answer corpus. Then video-question-answer triplets are automatically generated from narrated videos (right). Individual sentences are extracted from the ASR-transcribed narration using a punctuator  $p$ . Each extracted sentence is analyzed with an answer extractor  $T_a$  and a question generator  $T_q$  to produce answer  $a$  and question  $q$ . The timestamps of the narration are used to obtain a video clip  $v$  temporarily aligned to the extracted sentence to form the output video-question-answer triplet  $(v, q, a)$ .

answer pairs using a template-based approach [76].

Instead, we propose to use video annotations in the form of transcribed narrations or alt-text descriptions that are available at large-scale with no manual supervision. Moreover, rule-based generation requires the manual creation of rules by experts which is expensive, and has also been recently outperformed by neural question generation [26], [107], [119] as used in our approach.

**Large-scale pretraining for vision and language.** Several recent methods [4], [20], [24], [39], [40], [55], [57], [59], [64], [65], [85], [88], [104], [118] pretrain multi-modal vision-language representations, such as transformers, using datasets with image captions, e.g., COCO [18], Conceptual Captions [82] and Visual Genome [49]. These methods are often optimized using generic objectives such as masked language losses and losses for text-image matching and image caption generation. In our work, we pretrain models using large amounts of narrated videos. In contrast to task-agnostic pretraining in the previous work, we show the benefits of task-specific pretraining for our target VideoQA task.

**Learning from Web videos.** In this work, we exploit noisy correlations between videos and readily-available text annotations in unlabeled web videos from the recent HowTo100M [69] and WebVid2M [9] datasets. Methods using such readily-available data have shown significant improvements on several tasks including video retrieval [9], [31], action localization [69], action recognition [68] and video captioning [67], [86], [87], [120], sometimes outperforming fully-supervised baselines. Others have used videos with readily available text annotations for the VideoQA task. In detail, Amrani *et al.* [5] propose a text-video pretraining approach and finetune their model for VideoQA. Li *et al.* [56] propose HERO, a pretraining approach restricted to multiple-choice VideoQA, for which questions and answers are treated as a single text stream. Seo *et al.* [80] propose a pretraining approach based on next utterance prediction and finetune their model for VideoQA. Zellers *et al.* [80] propose a pretraining approach based on a mix of frame-level and video-level objectives and finetune for VideoQA. Differently to these methods with task-agnostic pretraining, we propose a pretraining approach specifically dedicated for VideoQA using automatically generated question and answer pairs from readily available text annotations, and show in Section 6

the benefits of our approach.

A preliminary version of this article has appeared in [103].

### 3 LARGE-SCALE GENERATION OF VIDEOQA DATA

This section presents our approach to generate large-scale VideoQA datasets from videos with readily available text annotations. We illustrate the proposed approach on instructional videos with text annotations in the form of transcribed narrations, which in many cases describe the content of the videos. Section 3.1 presents our proposed generation procedures. Section 3.2, then, describes the resulting HowToVQA69M dataset. Our approach can be easily adapted to other type of content, for example, shorter web-videos with with readily text annotations in the form of alt-text, as will be shown in the result section (Section 6.4).

#### 3.1 Generating video-question-answer triplets

We tackle the task of generating video-question-answer triplets from a large-scale instructional video dataset with transcribed spoken narration [69]. This is a challenging task because of transcription errors and lack of punctuation. We also wish to obtain highly diverse data. To address these issues, we propose to leverage powerful language models trained on text data. Our approach is illustrated in Figure 3 and details are given next.

We first present details about the generation procedure. Let  $s$  be the transcribed speech data obtained with automatic speech recognition (ASR). First, we use a recurrent neural network  $p$ , to infer punctuation in the transcribed speech data. We denote the punctuated transcript as  $p(s)$ . We extract video clips  $v$  temporally aligned with the inferred sentences  $p(s)$  using the ASR timestamps. We found that the generation works significantly better when applied to sentences rather than the original sentence fragments from the HowTo100M dataset, see Table 1. Second, for each sentence, we apply a transformer  $T_a$ , to extract a set of potential answers:  $a = T_a(p(s))$ . Third, we use another transformer  $T_q$  to generate a question given each transcript sentence and each extracted answer such that:  $q = T_q(a, p(s))$ . The output is a set of video-question-answer triplets  $(v, q, a)$ .



Punctuation	Generation method	Correct Samples	QA Generation Failure	QA unrelated to video
✓	Heilman <i>et al.</i> [35]	17	54	29
✗	Ours	23	49	28
✓	Ours	<b>30</b>	31	39

TABLE 1: Manual evaluation of our generation method (with and without punctuation) on a random sample of 100 examples compared with a rule-based question-answer generation of [35]. Numbers are obtained with majority voting between 5 annotators.

Question Type	Total	Correct Samples (%)	QA Generation Failure (%)	QA unrelated to video (%)
Attribute	25	28	32	40
Object	17	41	24	35
Action	16	<b>69</b>	19	13
Counting	13	23	15	<b>62</b>
Place	7	0	<b>86</b>	14
People	7	0	43	57
Other	15	13	27	60

TABLE 2: Manual evaluation of our generation method on 100 randomly chosen generated examples split by question type. Results are obtained by majority voting among 5 annotators.

pairs. HowToVQA69M is highly diverse and contains over 16M unique answers, where over 2M unique answers appear more than once and over 300K unique answers appear more than ten times. Examples of  $(v, q, a)$  triplets from the HowToVQA69M dataset are illustrated in Figure 4, showing the diversity and the noise in the automatically obtained annotations in HowToVQA69M.

**Statistical analysis of HowToVQA69M.** Figure 5 shows the statistics of the HowToVQA69M dataset in terms of the question length, answer length and video clip duration. Questions and answers contain 8.7 and 2.4 words on average respectively. Overall, HowToVQA69M contains longer answers than downstream VideoQA datasets like MSRVT-QA, MSVD-QA or ActivityNet-QA, for which answers are on average shorter than 2 words. Each clip lasts 12.1 seconds on average. The distribution of clip duration has a peak at around seven seconds with a long tail of longer clips. These statistics demonstrate the diversity of our HowToVQA69M dataset, in terms of videos, questions and answers.

Word clouds<sup>1</sup> for questions and answers in HowToVQA69M are shown in Figure 6 and illustrate the diverse vocabulary in HowToVQA69M as well as the presence of speech-related words such as *okay*, *right*, *oh*.

**Manual evaluation of HowToVQA69M.** As shown in Figure 4, HowToVQA69M annotations are noisy, which can be attributed to: (i) errors in speech transcription, (ii) speech not describing the video content, or (iii) errors in question-answer generation. We manually evaluate the quality of 100 randomly sampled  $(v, q, a)$  triplets in HowToVQA69M by collecting 5 different annotations for each triplet to reduce variance and report results in Table 1. Among 100 triplets generated by our method we find 30 to be correctly generated and matching well to the video content, 31 are incorrectly generated and 39 are correctly generated but unrelated to the video content. To demonstrate the influence of the different components of our automatic question-answer generation procedure, we compare our results with (i) a variant of our approach that does not split transcribed narrations into sentences using a punctuation, and (ii) a rule-based approach [35] for question-answer generation. Table 1 confirms the importance of punctuation and demonstrates the superior performance of our generation method compared to [35]. Further comparison with [35] is given

1. Word clouds were generated using [github.com/amueller/word\\_cloud](https://github.com/amueller/word_cloud).

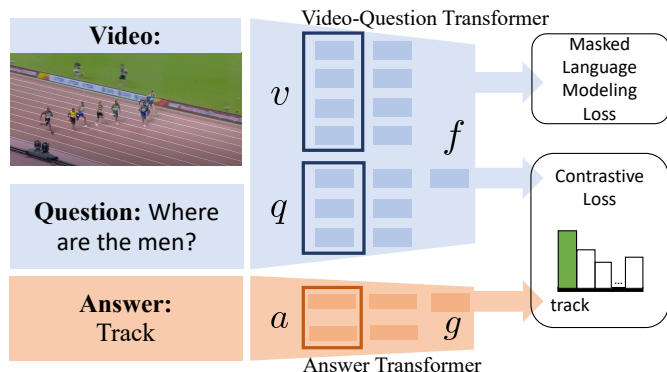


Fig. 7: **Overview of our VideoQA training architecture.** Our model is composed of a video-question module  $f$  based on a multi-modal transformer (top) and an answer module  $g$  based on DistilBERT [78] encoder (bottom). For pretraining, we use a contrastive loss and a masked language modeling loss (right).

in Section 6.6. In terms of inter-rater agreement, for the 300 generated video-question-answer triplets (100 for each generation method), 94 were in an agreement of all 5 annotators, 198 in an agreement of at least 4 annotators, and 299 in an agreement of at least 3 annotators. This high agreement of annotators demonstrates the reliability of the results in Table 1.

We further manually classify the 100 video-question-answer triplets obtained with our method by the question type (“Attribute”, “Object”, “Action”, “Counting”, “Place”, “People”, or “Other”), evaluate the quality of generated triplets for different question types and report results in Table 2. Out of the 6 most common categories, we observe that questions related to “Action” lead to the best annotations, “Counting” questions lead to the highest number of QAs unrelated to the video content, and questions related to “Place” lead to the highest number of QA generation errors. Qualitatively, we found that actions are often depicted in the video, while counted quantities (*e.g.* time, weight, length) mentioned in the speech are hard to guess from the video only. We describe next how we use HowToVQA69M to train our VideoQA model.

## 4 VIDEOQA MODEL AND TRAINING PROCEDURE

This section presents our VideoQA model (Section 4.1) and describes the training procedure (Section 4.2). Figure 7 gives an overview of the model.

### 4.1 VideoQA model

As illustrated in Figure 7, our VideoQA model is composed of two branches: (i) a video-question module  $f$  based on a transformer [92] and a mapping from the CLS token with a linear function. It takes a pair of video  $v$  and question  $q$  as input, models the multi-modal temporal interactions between  $v$  and  $q$  and then outputs an embedding vector  $f(v, q) \in \mathbb{R}^d$ . (ii) The second branch is a text encoder  $g$  that embeds an answer  $a$  as  $g(a) \in \mathbb{R}^d$ . We will denote our model as  $VQA-T$ , standing for VideoQA-Transformer. Note that using the joint (video, question) and answer embeddings allows us to deal with a large open vocabulary of answers present in our new HowToVQA69M dataset as the model can measure similarity between the input video-question embedding and the embedding of any answer. This is in contrast to using a classification answer module [38], [42], [43], [50], [122] that can choose only from a fixed predefined vocabulary

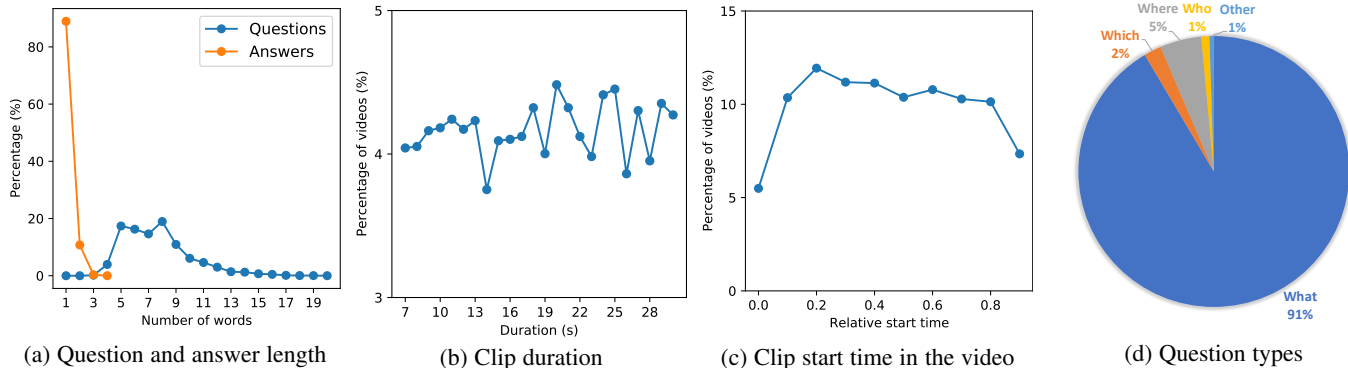


Fig. 8: **Statistics of the iVQA dataset.** (a) Distribution of length of questions and answers. (b) Distribution of video clip duration in seconds. (c) Distribution of video clip relative start time in the original video. (d) Distribution of question types.

of answers. Our embedding can be also easily finetuned on the different downstream VideoQA datasets, which may contain new answers that have not been seen at training. In contrast, the classification answer module has to be retrained when the vocabulary of answers changes. Next, we give details of the language and video representations. Further details about the model are provided in Appendix A.

**Word representation.** The question and answer are separately tokenized with the WordPieces embedding [96] and fed to DistilBERT [78]. DistilBERT is a light version of BERT [25] pretrained in a self-supervised fashion on English Wikipedia and the Toronto Book Corpus [121].

**Video representation.** We use a frozen S3D [98] pretrained on HowTo100M [69] using MIL-NCE [68]. This model is pretrained from scratch on HowTo100M only.

## 4.2 Training procedure

This section describes the training of our VideoQA model on the HowToVQA69M dataset and its finetuning on downstream VideoQA datasets.

**Training on HowToVQA69M.** We wish to make a pair of video and question  $(v, q)$  close to its correct answer  $a$  measured by the dot product of their embeddings,  $f(v, q)^\top g(a)$ . In contrast, the incorrect answers should be far, i.e., the dot product with their embeddings should be small. This can be done by maximizing the following contrastive objective:

$$\max_{f, g} \sum_{i=1}^n \log \left( \frac{e^{f(v_i, q_i)^\top g(a_i)}}{e^{f(v_i, q_i)^\top g(a_i)} + \sum_{(v', q', a') \sim \mathcal{N}_i} e^{f(v', q')^\top g(a')}} \right), \quad (1)$$

where  $(v_i, q_i, a_i)$  represents a generated triplet (video clip, question, answer) from HowToVQA69M. Given a specific positive triplet  $(v_i, q_i, a_i)$ , we construct the set  $\mathcal{N}_i$  of negative triplets by concatenating incorrect answers  $a_j$  within the training batch to the video-question pair  $(v_i, q_i)$  as:  $(v_i, q_i, a_j)$  with  $a_j \neq a_i$ . In particular, if the same negative answer  $a_j$  is present multiple times in a batch, we only count it once. We found that sampling the same negative answer multiple times leads to worse results (see Section 6.9), which we believe is due to different distributions of answers in the pretraining and downstream datasets. Removing duplicate negatives helps to mitigate this difference.

**Finetuning on downstream VideoQA datasets.** We leverage the model pretrained on HowToVQA69M and finetune it on a

downstream VideoQA dataset that typically has a smaller vocabulary of answers  $V$  (e.g.  $|V| \sim 4000$ ). To this end, we adapt the training objective in (1) by constructing the negative set  $\mathcal{N}_i$  from *all* incorrect answers in  $V$ . Note that in such setting (1) becomes equivalent to optimizing the standard cross-entropy objective. In the specific case of multiple-choice VideoQA, the set of negatives  $\mathcal{N}_i$  is the set of incorrect answers for each sample.

**Masked Language Modeling (MLM).** In addition to the contrastive loss (1) we apply the masking loss [25] to question tokens during both pretraining and finetuning. We found this to have a positive regularization effect when finetuning the DistilBERT weights (see Section 6.9).

## 5 iVQA: A NEW VIDEOQA EVALUATION DATASET

In this section we present our **Instructional VQA dataset (iVQA)**. We start from a subset of HowTo100M videos and manually annotate video clips with questions and answers. We aim (i) to provide a well-defined evaluation by including five correct answer annotations per question and (ii) to avoid questions which can be answered without watching the video. The dataset is described below.

**iVQA Data Collection.** iVQA videos are obtained by randomly sampling 7-30 sec. video clips from the HowTo100M dataset [69]. We avoid overlap between datasets and make sure iVQA and HowToVQA69M have no videos in common. Each clip is manually annotated with one question and 5 answers on Amazon Mechanical Turk. We ask workers to annotate questions about objects and scenes in the video and remove videos that could not be annotated. The correctness of annotations is manually verified by the authors. Moreover, we manually reduce the language bias by excluding questions that could be answered without watching the video. To increase diversity, each question is answered by 5 different workers. The answers are restricted to 4 words and are complemented by a confidence level. Questions that receive multiple answers with low confidence are removed. We further describe our data collection interfaces in [103] (Appendix C.1.).

**Statistical analysis of iVQA.** iVQA contains 10,000 video clips with one question and five corresponding answers per clip. We split the dataset into 60%/20%/20% train/validation/test subsets. Figure 8 shows the distributions of question length, answer length, clip duration, clip relative start time in the original video and question types. The average duration of video clips is 18.6 seconds. Clip duration and start time distributions are almost uniform because we randomly sampled both the duration and the start time to obtain the clips, which results in a high video content diversity.



Fig. 9: Word clouds for questions and answers in our iVQA dataset. The frequent occurrence of location and time-specific words (*behind*, *front*, *right*, *left*, *first*, *end*, *beginning*) indicates the presence of the spatial and temporal context within iVQA questions. We can also observe the task-specific vocabulary in iVQA answers related to the domains of cooking, hand crafting and gardening.

Most questions are about objects (*What* questions make up 91% of the data), while some are about places (*Where* questions make up 5% of the data) and people (*Who* questions make up 1% of the data). On average, questions and answers contain 7.6 and 1.1 words, respectively. Answers are in great majority one or two words, which is a result of our collection procedure.

The majority of questions have a consensus between at least 2 annotators, *i.e.* at least 2 annotators providing the same answer. In detail, we observe that 27.0% of questions lead to a perfect consensus among the five answer annotators, 48.4% of questions lead to a consensus among at least four annotators, and 77.3% lead to a consensus among at least three annotators. All but six questions lead to a consensus between at least two annotators. Additionally, 27.5% of questions have two different answers that had a consensus between at least two annotators. Similarly to [7], this motivates us to define the following accuracy measure for a given answer  $a$ :  $acc(a) = \min(\frac{\#ground\ truth\ answers = a}{2}, 1)$ . This metric assigns 100% accuracy to answers confirmed by at least 2 annotators, 50% accuracy to answers confirmed by only 1 annotator and 0% otherwise. Note that this definition is specific to our set-up where we have *multiple* ground truth answers per question.

Word clouds for questions and answers in the iVQA dataset in Figure 9 demonstrate the relation of iVQA to the domains of cooking, hand crafting and gardening. These word clouds also indicate that questions in iVQA often require spatial reasoning (*behind*, *front*, *right*, *left*) and temporal understanding (*first*, *end*, *left*, *beginning*) of the video. The most frequent answer (*spoon*) in iVQA corresponds to 2% of all answers in the dataset. In contrast, the most frequent answers in other existing VideoQA datasets account for more than 9% of all answers in these datasets (we have verified this for MSRVT-QA, MSVD-QA and ActivityNet-QA). As a consequence, the *most frequent answer baseline* is significantly lower for our iVQA dataset compared to other VideoQA datasets. We further evaluate the language bias in iVQA in Section 6.8.

## 6 EXPERIMENTS

This section demonstrates the benefits of training using our generated HowToVQA69M dataset and compares our method to the state of the art. We first outline the used datasets, baseline methods and implementation details in Section 6.1. We then present results for the novel zero-shot VideoQA task in Section 6.2. Next we present results for the novel VideoQA feature probe evaluation setting in Section 6.3. The comparison to the state of the art in VideoQA and alternative training strategies is given in Section 6.4. Section 6.5 presents results for rare answers and split per question type. Then we compare our VideoQA generation approach to previous methods in Section 6.6. We also apply our approach to another video-text datasets in Section 6.7. Finally, we show the importance of the visual modality in iVQA in Section 6.8 and present ablation studies in Section 6.9.

### 6.1 Evaluation Protocol

**Datasets.** We use three datasets for training and five datasets for evaluation as described below. We follow previous evaluation protocols for open-ended settings [50] and use a fixed vocabulary of training answers. Unless stated otherwise, we report top-1 test accuracy and use original splits for training, validation and test.

For training we use our new **HowToVQA69M** dataset introduced in Section 3.2 with 90% and 10% videos in training and validation subsets. For comparison, we also train our model using a large-scale text-video dataset, **HowTo100M** [69], that contains videos with transcribed narrations but *no video-question-answer* triplets. Test and validation videos of downstream datasets are excluded from HowTo100M and HowToVQA69M. To evaluate the general applicability of our approach, we generate another automatic VQA dataset based on **WebVid2M** [9], which consists of 2.5M video-text pairs scraped from the web where video captions are obtained from readily-available alt-text descriptions, see Section 6.7.

We evaluate results on four open-ended VideoQA downstream datasets: **MSRVTT-QA** [100], **MSVD-QA** [100], **ActivityNet-QA** [110] and our new **iVQA** dataset (see Section 5). We also evaluate on a multiple-choice VideoQA dataset **How2QA** [56] where each question is associated with one correct and three incorrect answers. For MSRVT-QA and MSVD-QA, we follow [50] and use a vocabulary of the top 4000 training answers for MSRVT-QA, and all 1852 training answers for MSVD-QA. For our iVQA dataset and ActivityNet-QA, we consider all answers that appear at least twice in the training set, resulting in 2348 answers for iVQA and 1654 answers for ActivityNet-QA.

**Baselines.** To evaluate the contribution of the visual modality, we compare our  $VQA-T$  model with its language-only variant  $QA-T$ .  $QA-T$  does not use video input, *i.e.* we set the input  $v$  of the video-question transformer to zero during both training and testing (see Figure 7). To evaluate our generated dataset, we also compare  $VQA-T$  trained on HowToVQA69M and on HowTo100M. Since HowTo100M has no  $(v, q, a)$  triplets, we only train the  $f$  branch of  $VQA-T$  on HowTo100M using the standard masking and cross-modal matching losses [20], [56], [64], [87], [120]. In the zero-shot setting we evaluate  $VQA-T$  trained on HowTo100M by computing  $f(v, [q, a])$  for concatenated pairs of questions and answers  $[q, a]$ . During finetuning we also initialize



Method	Pretraining Data	iVQA		MSRVTT-QA		MSVD-QA		ActivityNet-QA		How2QA
		Top-1	Top-10	Top-1	Top-10	Top-1	Top-10	Top-1	Top-10	Top-1
Random	$\emptyset$	0.09	0.9	0.02	0.2	0.05	0.5	0.05	0.5	25.0
QA-T	HowToVQA69M	4.4	23.2	2.5	6.5	4.8	15.0	11.6	45.8	38.4
VQA-T	HowTo100M	1.9	11.9	0.3	3.4	1.4	10.4	0.3	1.9	46.2
VQA-T (Ours)	HowToVQA69M	<b>12.2</b>	<b>43.3</b>	<b>2.9</b>	<b>8.8</b>	<b>7.5</b>	<b>22.4</b>	<b>12.2</b>	<b>46.5</b>	<b>51.1</b>

TABLE 3: Comparison with baselines for zero-shot VideoQA. Top-1 and top-10 (for open-ended datasets) accuracy are reported.



**Question:** What design are they making?  
**GT Answer:** rose (4), rose flower (1)  
**QA-T (HowToVQA69M):** pinwheel  
**VQA-T (HowTo100M):** piping bag  
**VQA-T (HowToVQA69M):** rose



**Question:** What is in the man's hand?  
**GT Answer:** shovel (3), spade (2)  
**QA-T (HowToVQA69M):** coin  
**VQA-T (HowTo100M):** planting  
**VQA-T (HowToVQA69M):** shovel



**Question:** What type of material is the man touching?  
**GT Answer:** wood (5)  
**QA-T (HowToVQA69M):** brick  
**VQA-T (HowTo100M):** electric saw  
**VQA-T (HowToVQA69M):** wood



**Question:** What is the woman decorating?  
**GT Answer:** cake (5)  
**QA-T (HowToVQA69M):** cupcakes  
**VQA-T (HowTo100M):** raspberries  
**VQA-T (HowToVQA69M):** cake



**Question:** What fruit is shown at the end?  
**GT Answer:** watermelon (5)  
**QA-T (HowToVQA69M):** pineapple  
**VQA-T (HowTo100M):** slotted spoon  
**VQA-T (HowToVQA69M):** watermelon



**Question:** What object is seen on the left, at the beginning of the video?  
**GT Answer:** teapot (4), pot (1)  
**QA-T (HowToVQA69M):** mirror  
**VQA-T (HowTo100M):** espresso  
**VQA-T (HowToVQA69M):** teapot



**Question:** What type of animal do we see in this video?  
**GT Answer:** shark (3), fish (2)  
**QA-T (HowToVQA69M):** stuffed animal  
**VQA-T (HowTo100M):** filter  
**VQA-T (HowToVQA69M):** fish



**Question:** What tool is being used?  
**GT Answer:** crochet hook (2), needle (2), hook (1)  
**QA-T (HowToVQA69M):** chisel  
**VQA-T (HowTo100M):** crocheting  
**VQA-T (HowToVQA69M):** crochet hook



**Question:** In which room does this scene take place?  
**GT Answer:** kitchen (5)  
**QA-T (HowToVQA69M):** bedroom  
**VQA-T (HowTo100M):** tablespoon  
**VQA-T (HowToVQA69M):** dining room



**Question:** What is the person holding?  
**GT Answer:** phone (4), smartphone (1)  
**QA-T (HowToVQA69M):** camera  
**VQA-T (HowTo100M):** keyboard  
**VQA-T (HowToVQA69M):** tablet

Fig. 10: **Zero-shot VideoQA on iVQA.** The top 4 rows illustrate successful predictions of our model, while the bottom-most row displays failure cases. The values next to the ground truth (GT) answers indicate the number of annotators that gave the answer.

Method	Pretraining data	iVQA	MSRVTT QA	MSVD QA	ActivityNet QA	How2QA
VQA-T	$\emptyset$	3.8	23.2	21.8	22.9	55.3
QA-T	HowToVQA69M	11.4	27.0	29.5	27.6	64.7
VQA-T	HowTo100M	13.8	27.0	32.9	24.7	63.9
VQA-T	HowToVQA69M	<b>24.5</b>	<b>32.9</b>	<b>39.0</b>	<b>30.6</b>	<b>72.9</b>

TABLE 4: Probe evaluation of different pretraining strategies. In each case, only the last projection layers in the model were finetuned on the downstream VideoQA datasets. Top-1 accuracy is reported.

Pretraining data	iVQA	MSRVTT QA	MSVD QA	ActivityNet QA	How2QA
$\emptyset$	23.0	39.6	41.2	36.8	80.8
HowTo100M	28.1	40.4	43.5	38.1	81.9
HowToVQA69M	<b>35.4</b>	<b>41.5</b>	<b>46.3</b>	<b>38.9</b>	<b>84.4</b>

TABLE 5: Benefits of pretraining our *VQA-T* model on our new HowToVQA69M dataset (last row) compared to no pretraining (first row) or pretraining on HowTo100M (second row). In each case our *VQA-T* model was then finetuned on the downstream VideoQA datasets. Top-1 accuracy is reported.

the  $g$  branch of *VQA-T* with parameters of the text encoding obtained from  $f$  (see further details in Appendix A).

**Implementation details.** For the training on HowToVQA69M we use the Adam optimizer and mini-batches with 4096 video clips sampled from 128 random videos. We use a cosine annealing learning rate schedule with initial value of  $5 \times 10^{-5}$ . The optimization over 10 epochs lasts 2 days on 8 Tesla V100 GPUs. Further details are included in Appendix B.

## 6.2 Zero-shot VideoQA

In this section, we address the *zero-shot VideoQA* task where we prohibit any manual supervision of visual data during training. We explore this setup to evaluate the generalization of *VQA-T* trained on HowToVQA69M to unseen downstream datasets. For consistency, we use the vocabulary of answers from downstream datasets during testing (see Section 6.1).

Zero-shot results are presented in Table 3. We first observe that the use of visual cues by *VQA-T* outperforms *QA-T* when both models are trained on HowToVQA69M. This demonstrates the importance of the cross-modality in HowToVQA69M despite the VideoQA annotation being exclusively generated from text-only methods. Since HowToVQA69M has been generated using no manual annotation of visual data, our approach is scalable and can lead to further improvements by increasing the dataset size, as we discuss in Section 6.9.

Training on HowToVQA69M significantly outperforms the training on HowTo100M and the random baseline. This confirms the advantage of our HowToVQA69M dataset for the VideoQA task over other generic text-video datasets that do not contain video-question-answer triplets. We emphasize that our training does not use any information about target VideoQA datasets. Qualitative results for zero-shot VideoQA are presented for our approach and compared with baselines in Figure 10. We observe that *QA-T* (trained on HowToVQA69M) provides plausible but video-unrelated answers to the questions. Moreover, *VQA-T* (trained on

Pretraining Data	Finetuning	iVQA				MSRVTT-QA				MSVD-QA				ActivityNet-QA			
		Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
$\emptyset$	✓	38.4	16.7	5.9	2.6	<b>68.4</b>	44.1	32.9	8.1	71.2	53.7	28.9	8.8	65.6	49.0	25.7	3.9
HowTo100M	✓	46.7	22.0	8.6	3.6	65.2	46.4	34.9	10.6	<b>74.8</b>	58.8	30.6	10.5	<b>67.5</b>	<b>53.3</b>	25.9	4.1
HowToVQA69M	✗	9.0	8.0	9.5	7.7	0.2	6.4	2.4	3.0	9.3	9.0	6.9	4.8	36.3	5.7	3.7	1.5
HowToVQA69M	✓	<b>47.9</b>	<b>28.1</b>	<b>15.6</b>	<b>8.5</b>	66.9	<b>46.9</b>	<b>36.0</b>	<b>11.5</b>	74.7	<b>59.0</b>	<b>35.0</b>	<b>14.1</b>	66.3	53.0	<b>28.0</b>	<b>5.0</b>

TABLE 6: Results of our *VQA-T* model with different training strategies, on subsets of iVQA, MSRVTT-QA, MSVD-QA and ActivityNet-QA, corresponding to four quartiles with Q1 and Q4 corresponding to samples with the most frequent and the least frequent answers, respectively.

Pretraining Data	Finetuning	MSRVTT-QA				MSVD-QA							
		What	Who	Number	Color	When	Where	What	Who	Number	Color	When	Where
$\emptyset$	✓	33.4	49.8	83.1	50.5	78.5	40.2	31.5	54.9	<b>82.7</b>	50.0	74.1	46.4
HowTo100M	✓	34.3	50.2	82.7	<b>51.8</b>	80.0	41.5	34.3	58.6	82.4	<b>62.5</b>	<b>77.6</b>	<b>50.0</b>
HowToVQA69M	✗	1.8	0.7	66.3	0.6	0.6	4.5	7.8	1.7	74.3	18.8	3.5	0.0
HowToVQA69M	✓	<b>35.5</b>	<b>51.1</b>	<b>83.3</b>	49.2	<b>81.0</b>	<b>43.5</b>	<b>37.9</b>	<b>58.0</b>	80.8	<b>62.5</b>	<b>77.6</b>	46.4

TABLE 7: Effect of our pretraining per question type on MSRVTT-QA and MSVD-QA.

Pretraining Data	Finetuning	Motion	Spatial	Temporal	Yes-No	Color	Object	Location	Number	Other
		$\emptyset$	✓	23.4	16.1	3.8	65.6	31.3	26.4	33.7
HowTo100M	✓	26.6	<b>17.7</b>	3.5	<b>67.5</b>	32.8	25.3	34.0	<b>50.5</b>	35.8
HowToVQA69M	✗	2.3	1.1	0.3	36.3	11.3	4.1	6.5	0.2	4.7
HowToVQA69M	✓	<b>28.0</b>	17.5	<b>4.9</b>	66.3	<b>34.3</b>	<b>26.7</b>	<b>35.8</b>	50.2	<b>36.8</b>

TABLE 8: Effect of our pretraining per question type on ActivityNet-QA.

Method	Pretraining data	MSRVTT	MSVD
		QA	QA
E-SA [100]		29.3	27.6
ST-TP [41]		30.9	31.3
AMU [100]		32.5	32.0
Co-mem [32]		32.0	31.7
HME [28]		33.0	33.7
LAGCN [38]		—	34.3
HGA [43]		35.5	34.7
QueST [42]		34.6	36.1
HCRN [50]		35.6	36.1
MASN [79]		35.2	38.0
Bridge to Answer [72]		36.9	37.2
OCRL+LOGNet [23]		36.0	38.2
ClipBERT [52]	COCO [18]+ Visual Genome [49]	37.4	—
Jin <i>et al.</i> [44]	Conceptual Caption [82]	37.6	38.2
SSML [5]	HowTo100M	35.1	35.1
CoMVT [80]	HowTo100M	39.5	42.6
SiaSamRea [109]	COCO [18]+ Visual Genome [49]	41.6	45.5
MERLOT [112]	YT-Temporal-180M	<b>43.1</b>	—
VQA-T	$\emptyset$	39.6	41.2
VQA-T	HowToVQA69M	41.5	<b>46.3</b>
VQA-T	HowToVQA69M+ WebVidVQA3M	41.8	<b>47.5</b>

TABLE 9: Comparison with state of the art on MSRVTT-QA and MSVD-QA (top-1 accuracy).

Pretraining data	ActivityNet	How2QA
	QA	
E-SA [110]	31.8	—
MAR-VQA [122]	34.6	—
HERO [56]	HowTo100M + TV Dataset	—
CoMVT [80]	HowTo100M	74.1
SiaSamRea [109]	HowTo100M	38.8
Visual Genome [49]	COCO [18]+	82.3
MERLOT [112]	YT-Temporal-180M	39.8
VQA-T	$\emptyset$	84.1
VQA-T	HowToVQA69M	<b>41.4</b>
VQA-T	HowToVQA69M+	36.8
WebVidVQA3M	HowToVQA69M+	38.9
WebVidVQA3M	WebVidVQA3M	39.0
		<b>85.3</b>

TABLE 10: Comparison with state of the art on ActivityNet-QA and the public val set of How2QA (top-1 accuracy).

HowTo100M) is able to reply with answers related to the visual content, but doesn’t take into account the question. Our *VQA-T* model trained on HowToVQA69M, on the other hand, correctly understands questions and uses information in the video to provide

correct answers, confirming results in Table 3. We also illustrate some failure cases in Figure 10, showing that our zero-shot *VQA-T* model can fail to understand fine variations in the video or the language, confusing a *kitchen* with a *dining room* or a *phone* with a *tablet*.


### 6.3 VideoQA feature probe evaluation

In this section we further evaluate the generalization capabilities of the multi-modal representation learnt by our pretrained model. To this end, we analyze the effect of *VQA-T* pretraining followed by the *finetuning of the final projection layers* on the target datasets. More precisely, only the final MLP in the video-question module and the final linear layer in the answer module are finetuned. All other weights in the model (notably the video, question, answer representations and the multi-modal transformer) are kept frozen after the large-scale pre-training.


Results for VideoQA feature probe evaluation are reported in Table 4. Similarly as for the zero-shot setting, we observe that the use of visual cues in *VQA-T* outperforms *QA-T* when both models are pretrained on HowToVQA69M. Additionally, pretraining on HowToVQA69M significantly outperforms the pretraining on HowTo100M and the probe baseline trained from scratch. Note that the probe baseline trained from scratch, despite having notably randomly initialized frozen multi-modal transformer weights, achieves reasonable absolute results as it can exploit dataset biases, which are further ablated in Section 6.8. Interestingly, we find that on the iVQA dataset, the probe evaluation of our model pretrained on HowToVQA69M (24.5% accuracy, first line in Table 4) outperforms the fully supervised model trained from scratch (23.0% accuracy, first line in Table 5). These results further confirms the quality of our multi-modal representation learnt from HowToVQA69M.

### 6.4 Benefits of HowToVQA69M pretraining

This section evaluates the effect of *VQA-T* pretraining in combination with finetuning on target datasets. As shown in Table 5, pretraining on HowToVQA69M provides consistent and significant improvements for all datasets when compared to pretraining on HowTo100M and no pretraining. In particular, we observe the largest improvement for our new iVQA dataset which comes



**ASR:** And then just squeeze it through like that.  
**Question (Heilman et al):** What do then just squeeze through like that?  
**Answer (Heilman et al):** it  
**Question (ours):** How do you do it?  
**Answer (ours):** squeeze it through



**ASR:** It is a staple in a lot of asian kitchens.  
**Question (Heilman et al):** What is it?  
**Answer (Heilman et al):** a staple in a lot of asian kitchens  
**Question (ours):** In what type of kitchens is it a staple?  
**Answer (ours):** asian kitchens



**ASR:** And you want it over a very low heat.  
**Question (Heilman et al):** What do you want it over?  
**Answer (Heilman et al):** over a very low heat  
**Question (ours):** What kind of heat do you want it to be over?  
**Answer (ours):** low heat



**ASR:** This is classic premium chicken, grilled sandwich.  
**Question (Heilman et al):** What is classic premium chicken, grilled sandwich?  
**Answer (Heilman et al):** this  
**Question (ours):** What type of sandwich is this?  
**Answer (ours):** classic premium chicken, grilled sandwich



**ASR:** But why do that when you can enjoy the plant for about three months, it'll, keep producing because the leaves grow from the center  
**Question (Heilman et al):** What leaves?  
**Answer (Heilman et al):** the  
**Question (ours):** What part of the plant grows from the center?  
**Answer (ours):** leaves



**ASR:** Next add half a cup of powdered milk and a little shake a quarter teaspoon of salt, which I know, sounds really weird.  
**Question (Heilman et al):** What do I know the quarter teaspoon of?  
**Answer (Heilman et al):** of salt  
**Question (ours):** What is a quarter teaspoon of?  
**Answer (ours):** salt

Fig. 11: Qualitative examples of video-question-answer triplets generated with our trained language models compared to Heilman *et al.* [35], illustrating the higher quality and diversity of triplets obtained with our generation method.

Pretraining Data	Zero-shot				Finetune			
	iVQA	MSRVTT-QA	ActivityNet-QA	How2QA	iVQA	MSRVTT-QA	ActivityNet-QA	How2QA
$\emptyset$	—	—	—	—	23.0	39.6	36.8	80.8
MSRVTT-QA	8.6	—	1.7	42.5	25.2	—	37.5	80.0
ActivityNet-QA	5.5	2.7	—	40.8	24.0	39.9	—	80.7
HowToVQA69M	<b>12.2</b>	<b>2.9</b>	<b>12.2</b>	<b>51.1</b>	<b>35.4</b>	<b>41.5</b>	<b>38.9</b>	<b>84.4</b>

TABLE 11: Comparison of our training on HowToVQA69M with cross-dataset transfer using the previously largest open-ended VideoQA dataset (MSRVTT-QA) and the largest manually annotated open-ended VideoQA dataset (ActivityNet-QA).

Generation Method	Zero-shot			Finetune		
	iVQA	ActivityNet QA	How2QA	iVQA	ActivityNet QA	How2QA
[35]	7.4	1.1	41.7	31.4	38.5	83.0
Ours	<b>12.2</b>	<b>12.2</b>	<b>51.1</b>	<b>35.4</b>	<b>38.9</b>	<b>84.4</b>

TABLE 12: Comparison of our question-answer generation approach with Heilman *et al.* [35], evaluated by downstream performance of the model trained on the generated VideoQA data.

from the same domain as HowToVQA69M. Hence, the automatic generation of training data for other domains using our method can lead to further improvements on other datasets.

We compare our pretrained model to the state-of-the-art in VideoQA in Tables 9-10. Notably, *VQA-T* pretrained on HowToVQA69M outperforms previous methods using comparable pretraining data on all tested datasets. In particular, our method improves over CoMVT [80] that has been pretrained on HowTo100M. We note that the recent SiaSamRea approach [109] improves over our method on MSRVTT-QA (+0.1%) and ActivityNet-QA (+0.9%), but achieves lower results on MSVD-QA (-0.8%) and How2QA (-0.3%). However, SiaSamRea leverages manually annotated visual data for pretraining (COCO [18] and Visual Genome [49]). We also note that MERLOT [112] improves over our method on MSRVTT-QA and ActivityNet-QA, but uses the YT-Temporal-180M dataset for pretraining. This dataset includes HowTo100M but is significantly larger and more diverse (6 millions YouTube videos instead of 1 million).

## 6.5 Analysis of rare answers and question types

**Results for rare answers.** Training on downstream VideoQA datasets typically leads to particularly large improvements for questions with most frequent answers. As shown in Table 6, our approach brings significant improvements both for common and rare answers compared to models trained from scratch or pretrained on HowTo100M. We also find that our pretrained model, in the zero-shot setting, performs similarly across the different quartiles, with the exception of ActivityNet-QA, which includes in its most common answers *yes*, *no*. Interestingly, for the most rare answers in iVQA (Q3 and Q4) our model without finetuning (zero-shot mode) outperforms finetuned models that have not been pretrained on HowToVQA69M. We conclude that VideoQA specific pretraining on additional large-scale, diverse data helps improve generalization of VideoQA models.

Note that in order to have a consistent evaluation with other experiments, we keep the same train vocabulary at test time. This implies that a significant part of answers in the test set is considered wrong because the answer is not in the vocabulary. This represents 16% of answers in iVQA, 3% of answers in MSRVTT-QA, 6% for MSVD-QA and 19% for ActivityNet-QA. Note, however, that our joint embedding framework could allow for different vocabularies to be used at the training and test time.

**Results split per question type.** We also present results per question type for MSRVTT-QA, MSVD-QA and ActivityNet-QA in Tables 7 and 8. Compared to the model trained from scratch or the model pretrained on HowTo100M, we observe consistent improvements by our model for most categories.

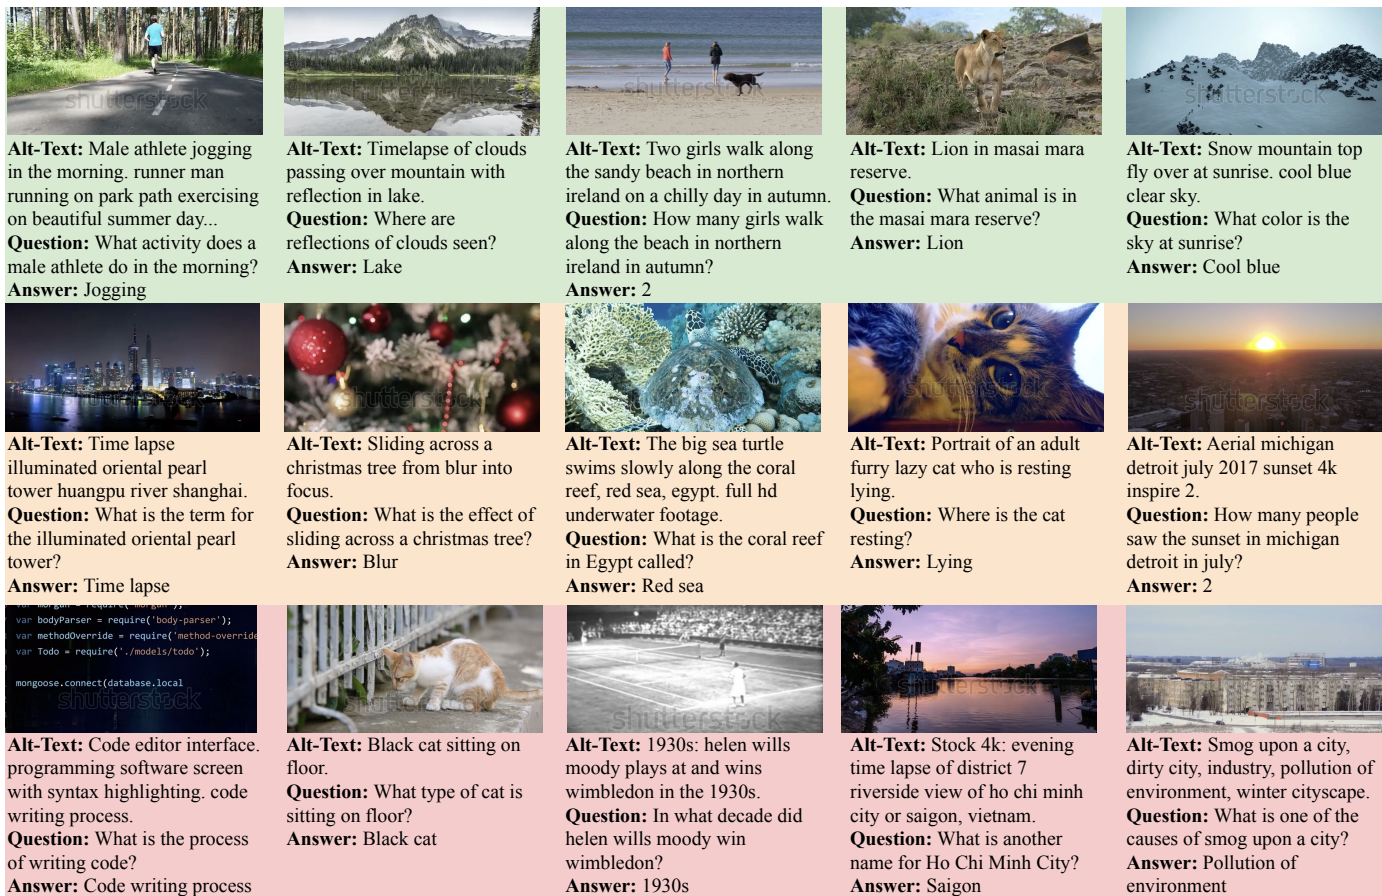


Fig. 12: Examples of questions-answers generated from video alt-text pairs from the WebVid2M dataset [9]. The green color (first row) indicates relevant examples, the orange color (second row) indicates a failure of the question-answer generation, and the red color (third row) indicates that the generated question-answer is unrelated to the visual content.

Pretraining Data	Zero-shot					Finetune				
	iVQA	MSRVTT-QA	MSVD-QA	ActivityNet-QA	How2QA	iVQA	MSRVTT-QA	MSVD-QA	ActivityNet-QA	How2QA
$\emptyset$	—	—	—	—	—	23.0	39.6	41.2	36.8	80.8
WebVidVQA3M	7.3	5.3	12.3	6.2	49.8	28.1	41.2	45.4	38.1	82.4
HowToVQA69M	12.2	2.9	7.5	12.2	51.1	<b>35.4</b>	41.5	46.3	38.9	84.4
HowToVQA69M + WebVidVQA3M	<b>13.3</b>	<b>5.6</b>	<b>13.5</b>	<b>12.3</b>	<b>53.1</b>	35.2	<b>41.8</b>	<b>47.5</b>	<b>39.0</b>	<b>85.3</b>

TABLE 13: Comparison of our VideoQA training datasets generated with different video-text data source, evaluated by downstream performance of the model pretrained on the generated data in zero-shot mode and after finetuning.

## 6.6 Comparison of VideoQA generation methods and VideoQA training datasets

**Comparison of VideoQA generation methods.** We compare our question-answer generation approach to Heilman *et al.* [35], that was notably used in [100], [113], [115], [116], [117] to generate VideoQA data from video descriptions. We run the method of [35] on sentences extracted from HowTo100M, apply our pretraining method on the generated data and show results in Table 12. Note that we do not choose MSRVTT-QA and MSVD-QA as downstream datasets for this comparison because their evaluation sets were automatically generated using Heilman *et al.* [35]. We find that our generation method leads to significantly better performance both in zero-shot and finetuning settings. We supplement this quantitative comparison with a qualitative comparison shown in Figure 11. We found that compared to [35] our generation method provides higher quality as well as higher diversity of question-answer pairs when applied to the uncurated sentences extracted from speech in narrated videos. This further demonstrates the benefit of our transformer-based question-answer

generation approach compared to previous rule-based methods.

**Comparison of VideoQA training datasets.** We also evaluate the importance of our generated HowToVQA69M dataset by comparing our results to cross-dataset transfer using existing VideoQA datasets. We define cross-dataset transfer as a procedure where we pretrain our VideoQA model on a VideoQA dataset and then finetune and test it on another VideoQA dataset. The training follows the procedure described for finetuning in Section 4.2. We report results for cross-dataset transfer in Table 11. Note that we do not use MSVD-QA as downstream dataset as its test set has been automatically generated with the same method [35] as MSRVTT-QA. As can be observed, our approach with pretraining on HowToVQA69M significantly outperforms cross-dataset transfer models using the previously largest VideoQA dataset (MSRVTT-QA), or the largest manually annotated VideoQA dataset (ActivityNet-QA), both for the zero-shot and finetuning settings, on all four downstream datasets. We emphasize that our dataset is generated relying on text-only annotations, while MSRVTT-QA was generated using manually annotated video descriptions and ActivityNet-QA was manually collected. These

Method	iVQA	MSRVTT QA	MSVD QA	ActivityNet QA	How2QA
QA-T	14.1	32.8	32.6	30.4	76.6
VQA-T	23.0	39.6	41.2	36.8	80.8

TABLE 14: Comparison of *QA-T* and *VQA-T* models trained from scratch (without pretraining) on downstream datasets.

results further demonstrate the benefit of our HowToVQA69M dataset for training VideoQA models.

## 6.7 Generalization to other video-text datasets

In this section, we show that our VideoQA generation approach can be generalized to other sources of non-manually annotated video-text paired data. For this, we extend and apply our generation pipeline presented in Section 3.1 to videos with alt-text description, *i.e.* alt-text HTML attribute associated with videos, from the WebVid2M dataset [9].

**WebVidVQA3M dataset.** We first explain how we adapt our generation pipeline detailed in Section 3.1 to video alt-text pairs. As captions in WebVid2M are relatively short, we do not apply the punctuation model but directly apply the question-answer generation models on the captions. Captions in WebVid2M are also not temporally localized, so the generated question-answers are not temporally localized either. They instead refer to the whole videos, which are typically short (4 seconds on average). Applying our generation pipeline to WebVid2M [9], we generate WebVidVQA3M, a dataset of 3,476,610 question-answers associated with 2,404,871 videos. Examples of generated samples are illustrated in Figure 12. These examples show that despite a substantial visual-linguistic domain difference compared to HowTo100M, our approach is able to generate relevant VideoQA data. We believe that qualitatively, the generated QA data from WebVidVQA3M are of better quality than the generated QA data from HowToVQA69M (see Section 3.2). We argue that WebVid2M [9] has a better visual-linguistic correlation and a higher quality of text data compared to HowTo100M [69], which facilitates the VideoQA generation.

**Benefits of training on WebVidVQA3M.** We next apply our pretraining method on the generated data and show results in Table 13. We also explore combining both datasets with a simple curriculum learning strategy, where our model initially pretrained on HowToVQA69M is further trained on WebVidVQA3M. We find that training only on WebVidVQA3M gives competitive performance both in the zero-shot setting and the finetuning setting. Notably, it significantly improves over the variant trained from scratch in the finetuning setting. This shows that our approach can be generalized to other sources of video and text data. Additionally, we find that combining the two datasets for pretraining results in additional improvements both for zero-shot and finetuning. Therefore the difference with previous methods is also increased (see Tables 9-10). Note that as WebVidVQA3M is significantly smaller than HowToVQA69M, our training runs faster on this dataset (20 GPUH instead of 350 GPUH), which gives a practical advantage to WebVidVQA3M. We have open-sourced WebVidVQA3M annotations to facilitate future research.

## 6.8 Importance of the visual modality in iVQA

We show in Table 14 that *QA-T* is a strong baseline compared to *VQA-T* on existing VideoQA datasets, when both are trained from scratch. However, on iVQA, *VQA-T* improves even more over *QA-T* than with other datasets, as measured by absolute

MLM	Sampling without answer repetition	Zero-shot		Finetune	
		iVQA	MSVD-QA	iVQA	MSVD-QA
✗	✗	11.1	6.1	34.7	45.6
✗	✓	12.1	7.0	34.3	45.0
✓	✗	10.9	6.4	34.3	45.1
✓	✓	<b>12.2</b>	<b>7.5</b>	<b>35.4</b>	<b>46.3</b>

TABLE 15: Effect of MLM loss and our negative sampling strategy on HowToVQA69M training.

Pretraining data size	Zero-shot		Finetune	
	iVQA	MSVD-QA	iVQA	MSVD-QA
0%	—	—	23.0	41.2
1%	4.5	3.6	24.2	42.8
10%	9.1	6.2	29.2	44.4
20%	9.5	6.8	31.3	44.8
50%	11.3	7.3	32.8	45.5
100%	<b>12.2</b>	<b>7.5</b>	<b>35.4</b>	<b>46.3</b>

TABLE 16: Effect of the training size of HowToVQA69M.

improvement in top-1 accuracy. This suggests that the visual modality is more important in iVQA than in other VideoQA datasets.

## 6.9 Ablation studies

**Pretraining losses.** As shown in Table 15, removing duplicate negative answers in our contrastive loss, as discussed in Section 4.2, is beneficial notably in the zero-shot setting. Moreover, adding the MLM loss during pretraining improves the downstream results for both zero-shot and finetuning when used in combination with our contrastive learning strategy. These results motivate our proposed pretraining approach.

**Importance of scale.** Results of our method after pretraining on different fractions of HowToVQA69M are shown in Table 16. We construct these subsets such that larger subsets include the smaller ones. These results suggest that the scale is an important factor and that we can expect further improvements with additional pretraining data, both in the zero-shot and finetuning settings.

## 7 CONCLUSION

We propose a novel and scalable approach for training VideoQA models without manually annotated visual data. We automatically generate HowToVQA69M – a large-scale VideoQA training dataset generated from narrated videos with readily-available speech transcripts, significantly exceeding existing datasets by size and diversity. We demonstrate several benefits of pretraining on HowToVQA69M. We are the first to demonstrate zero-shot VideoQA results while using no manually annotated images or videos for training. We also introduce the VideoQA feature probe evaluation setting and show strong generalization capabilities of the multi-modal representation learnt by our pretrained model. Furthermore, finetuning our HowToVQA69M pretrained model on downstream tasks achieves competitive performance on MSRVTT-QA, ActivityNet-QA, MSVD-QA and How2QA. Moreover, we show that our approach generalizes to other sources of web videos by generating the WebVidVQA3M from video alt-text pairs and showing its benefits for VideoQA training. We further validate our approach on our new manually-collected iVQA benchmark.

## ACKNOWLEDGMENTS

This work was granted access to the HPC resources of IDRIS under the allocation 2020-101267 made by GENCI. The work

was funded by a Google gift, the French government under management of Agence Nationale de la Recherche as part of the "Investissements d'avenir" program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute), the Louis Vuitton ENS Chair on Artificial Intelligence and the European Regional Development Fund under project IMPACT (reg. no. CZ.02.1.01/0.0/0.0/15/003/0000468). We thank Pierre-Louis Guhur and Makarand Tapaswi for advice on using Amazon Mechanical Turk, Eloïse Berthier, Quentin Le Lidec and Elliot Chane-Sane for the manual evaluation of generated VideoQA data, and Ignacio Rocco for proofreading.

## REFERENCES

- [1] Punctuator. <https://github.com/ottokart/punctuator2>, 2017.
- [2] Question generation using transformers. [https://github.com/patil-suraj/question\\_generation](https://github.com/patil-suraj/question_generation), 2020.
- [3] Chris Alberti, Daniel Andor, Emily Pitler, Jacob Devlin, and Michael Collins. Synthetic QA corpora generation with roundtrip consistency. In *ACL*, 2019.
- [4] Chris Alberti, Jeffrey Ling, Michael Collins, and David Reitter. Fusion of detected objects in text for visual question answering. In *IJCNLP*, 2019.
- [5] Elad Amrani, Rami Ben-Ari, Daniel Rotman, and Alex Bronstein. Noise estimation using density estimation for self-supervised multimodal learning. In *AAAI*, 2021.
- [6] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 2018.
- [7] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. VQA: Visual question answering. In *ICCV*, 2015.
- [8] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [9] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, 2021.
- [10] Pratyay Banerjee, Tejas Gokhale, Yezhou Yang, and Chitta Baral. WeaQA: Weak supervision via captions for visual question answering. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 2021.
- [11] Hedi Ben-Younes, Rémi Cadene, Matthieu Cord, and Nicolas Thome. MUTAN: Multimodal tucker fusion for visual question answering. In *CVPR*, 2017.
- [12] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- [13] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, 2020.
- [14] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021.
- [15] Santiago Castro, Mahmoud Azab, Jonathan Stroud, Cristina Noujaim, Ruoyao Wang, Jia Deng, and Rada Mihalcea. LifeQA: A real-life dataset for video question answering. In *LREC*, 2020.
- [16] Aman Chadha, Gurmeet Arora, and Navpreet Kaloty. iPerceive: Applying common-sense reasoning to multi-modal dense video captioning and video question answering. In *WACV*, 2021.
- [17] Ying-Hong Chan and Yao-Chung Fan. A recurrent BERT-based model for question generation. In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, 2019.
- [18] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015.
- [19] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *ICCV*, 2021.
- [20] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. UNITER: Universal image-text representation learning. In *ECCV*, 2020.
- [21] Seongho Choi, Kyoung-Woon On, Yu-Jung Heo, Ahjeong Seo, Youwon Jang, Seungchan Lee, Minsu Lee, and Byoung-Tak Zhang. DramaQA: Character-centered video story understanding with hierarchical qa. In *AAAI*, 2021.
- [22] Anthony Colas, Seokhwan Kim, Franck Dernoncourt, Siddhesh Gupte, Daisy Zhe Wang, and Doo Soon Kim. TutorialVQA: Question answering dataset for tutorial videos. In *LREC*, 2020.
- [23] Long Hoang Dang, Thao Minh Le, Vuong Le, and Truyen Tran. Object-centric representation learning for video question answering. In *IJCNN*, 2021.
- [24] Karan Desai and Justin Johnson. VirTex: Learning visual representations from textual annotations. In *CVPR*, 2021.
- [25] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019.
- [26] Xinya Du, Junru Shao, and Claire Cardie. Learning to ask: Neural question generation for reading comprehension. In *ACL*, 2017.
- [27] Chenyou Fan. EgoVQA - an egocentric video question answering benchmark dataset. In *ICCV Workshops*, 2019.
- [28] Chenyou Fan, Xiaofan Zhang, Shu Zhang, Wensheng Wang, Chi Zhang, and Heng Huang. Heterogeneous memory enhanced multimodal attention model for video question answering. In *CVPR*, 2019.
- [29] Marcello Federico, Sebastian Stüker, Luisa Bentivogli, Michael Paul, Mauro Cettolo, Teresa Herrmann, Jan Niehues, and Giovanni Moretti. The IWSLT 2011 evaluation campaign on automatic talk translation. In *LREC*, 2012.
- [30] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *EMNLP*, 2016.
- [31] Valentin Gabeur, Chen Sun, Karteeek Alahari, and Cordelia Schmid. Multi-modal transformer for video retrieval. In *ECCV*, 2020.
- [32] Jiyang Gao, Runzhou Ge, Kan Chen, and Ram Nevatia. Motion-appearance co-memory networks for video question answering. In *CVPR*, 2018.
- [33] Noa Garcia, Mayu Otani, Chenhui Chu, and Yuta Nakashima. KnowIT VQA: Answering knowledge-based questions about videos. In *AAAI*, 2020.
- [34] 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 *CVPR*, 2017.
- [35] Michael Heilman and Noah A Smith. Good question! Statistical ranking for question generation. In *ACL*, 2010.
- [36] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (GELUs). *arXiv preprint arXiv:1606.08415*, 2016.
- [37] Hexiang Hu, Wei-Lun Chao, and Fei Sha. Learning answer embeddings for visual question answering. In *CVPR*, 2018.
- [38] Deng Huang, Peihao Chen, Runhao Zeng, Qing Du, Mingkui Tan, and Chuang Gan. Location-aware graph convolutional networks for video question answering. In *AAAI*, 2020.
- [39] Zhicheng Huang, Zhaoyang Zeng, Yupan Huang, Bei Liu, Dongmei Fu, and Jianlong Fu. Seeing out of the box: End-to-end pre-training for vision-language representation learning. In *CVPR*, 2021.
- [40] Zhicheng Huang, Zhaoyang Zeng, Bei Liu, Dongmei Fu, and Jianlong Fu. Pixel-BERT: Aligning image pixels with text by deep multi-modal transformers. *arXiv preprint arXiv:2004.00849*, 2020.
- [41] Yunseok Jang, Yale Song, Youngjae Yu, Youngjin Kim, and Gunhee Kim. TGIF-QA: Toward spatio-temporal reasoning in visual question answering. In *CVPR*, 2017.
- [42] Jianwen Jiang, Ziqiang Chen, Haojie Lin, Xibin Zhao, and Yue Gao. Divide and conquer: Question-guided spatio-temporal contextual attention for video question answering. In *AAAI*, 2020.
- [43] Pin Jiang and Yahong Han. Reasoning with heterogeneous graph alignment for video question answering. In *AAAI*, 2020.
- [44] WeiKe Jin, Zhou Zhao, Xiaochun Cao, Jieming Zhu, Xiuqiang He, and Yueting Zhuang. Adaptive spatio-temporal graph enhanced vision-language representation for video qa. *IEEE Transactions on Image Processing*, 2021.
- [45] Hyoungun Kim, Zineng Tang, and Mohit Bansal. Dense-caption matching and frame-selection gating for temporal localization in VideoQA. In *ACL*, 2020.
- [46] Junyeong Kim, Minuk Ma, Trung Pham, Kyungsu Kim, and Chang D Yoo. Modality shifting attention network for multi-modal video question answering. In *CVPR*, 2020.
- [47] Kyung-Min Kim, Min-Oh Heo, Seong-Ho Choi, and Byoung-Tak Zhang. Deepstory: Video story qa by deep embedded memory networks. In *IJCAI*, 2017.
- [48] Seonhoo Kim, Seohyeon Jeong, Eunbyul Kim, Inho Kang, and Nojun Kwak. Self-supervised pre-training and contrastive representation learning for multiple-choice video qa. In *AAAI*, 2021.
- [49] Ranjaya Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, Michael Bernstein, and Li Fei-Fei. Visual Genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 2016.
- [50] Thao Minh Le, Vuong Le, Svetha Venkatesh, and Truyen Tran. Hier-

- archical conditional relation networks for video question answering. In *CVPR*, 2020.
- [51] Thao Minh Le, Vuong Le, Svetha Venkatesh, and Truyen Tran. Neural reasoning, fast and slow, for video question answering. In *JCNN*, 2020.
- [52] Jie Lei, Linjie Li, Luwei Zhou, Zhe Gan, Tamara L Berg, Mohit Bansal, and Jingjing Liu. Less is more: Clipbert for video-and-language learning via sparse sampling. In *CVPR*, 2021.
- [53] Jie Lei, Licheng Yu, Mohit Bansal, and Tamara L Berg. TVQA: Localized, compositional video question answering. In *EMNLP*, 2018.
- [54] Jie Lei, Licheng Yu, Tamara L Berg, and Mohit Bansal. TVQA+: Spatio-temporal grounding for video question answering. In *ACL*, 2020.
- [55] Gen Li, Nan Duan, Yuejian Fang, Ming Gong, Daxin Jiang, and Ming Zhou. Unicoder-VL: A universal encoder for vision and language by cross-modal pre-training. In *AAAI*, 2020.
- [56] Linjie Li, Yen-Chun Chen, Yu Cheng, Zhe Gan, Licheng Yu, and Jingjing Liu. HERO: Hierarchical encoder for video+language omni-representation pre-training. In *EMNLP*, 2020.
- [57] Liunan Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. VisualBERT: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019.
- [58] Xiangpeng Li, Jingkuan Song, Lianli Gao, Xianglong Liu, Wenbing Huang, Xiangnan He, and Chuang Gan. Beyond RNNs: Positional self-attention with co-attention for video question answering. In *AAAI*, 2019.
- [59] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *ECCV*, 2020.
- [60] Yikang Li, Nan Duan, Bolei Zhou, Xiao Chu, Wanli Ouyang, Xiaogang Wang, and Ming Zhou. Visual question generation as dual task of visual question answering. In *CVPR*, 2018.
- [61] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common Objects in Context. In *ECCV*, 2014.
- [62] Xudong Lin, Gedas Bertasius, Jue Wang, Shih-Fu Chang, Devi Parikh, and Lorenzo Torresani. VX2TEXT: End-to-end learning of video-based text generation from multimodal inputs. In *CVPR*, 2021.
- [63] Luis Enrico Lopez, Diane Kathryn Cruz, Jan Christian Blaise Cruz, and Charibeth Cheng. Transformer-based end-to-end question generation. *arXiv preprint arXiv:2005.01107*, 2020.
- [64] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, 2019.
- [65] Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 12-in-1: Multi-task vision and language representation learning. In *CVPR*, 2020.
- [66] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Hierarchical question-image co-attention for visual question answering. In *NeurIPS*, 2016.
- [67] Huaishao Luo, Lei Ji, Botian Shi, Haoyang Huang, Nan Duan, Tianrui Li, Xilin Chen, and Ming Zhou. UniViLM: A unified video and language pre-training model for multimodal understanding and generation. *arXiv preprint arXiv:2002.06353*, 2020.
- [68] Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-end learning of visual representations from uncurated instructional videos. In *CVPR*, 2020.
- [69] 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*, 2019.
- [70] Nasrin Mostafazadeh, Ishan Misra, Jacob Devlin, Margaret Mitchell, Xiaodong He, and Lucy Vanderwende. Generating natural questions about an image. In *ACL*, 2016.
- [71] Jonghwan Mun, Paul Hongsuck Seo, Ilchae Jung, and Bohyung Han. MarioQA: Answering questions by watching gameplay videos. In *CVPR*, 2017.
- [72] Jungin Park, Jiyoung Lee, and Kwanghoon Sohn. Bridge to answer: Structure-aware graph interaction network for video question answering. In *CVPR*, 2021.
- [73] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021.
- [74] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*, 2020.
- [75] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. SQuAD: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*, 2016.
- [76] Mengye Ren, Ryan Kiros, and Richard Zemel. Exploring models and data for image question answering. In *NeurIPS*, 2015.
- [77] Arka Sadhu, Kan Chen, and Ram Nevatia. Video question answering with phrases via semantic roles. In *NAACL*, 2021.
- [78] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.
- [79] Ahjeong Seo, Gi-Cheon Kang, Joonhan Park, and Byoung-Tak Zhang. Attend what you need: Motion-appearance synergistic networks for video question answering. In *ACL*, 2021.
- [80] Paul Hongsuck Seo, Arsha Nagrani, and Cordelia Schmid. Look before you speak: Visually contextualized utterances. In *CVPR*, 2021.
- [81] Meet Shah, Xinlei Chen, Marcus Rohrbach, and Devi Parikh. Cycle-consistency for robust visual question answering. In *CVPR*, 2019.
- [82] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual Captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*, 2018.
- [83] Xiaomeng Song, Yucheng Shi, Xin Chen, and Yahong Han. Explore multi-step reasoning in video question answering. In *ACM international conference on Multimedia*, 2018.
- [84] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *JMLR*, 2014.
- [85] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. VL-BERT: Pre-training of generic visual-linguistic representations. In *ICLR*, 2019.
- [86] Chen Sun, Fabien Baradel, Kevin Murphy, and Cordelia Schmid. Contrastive bidirectional transformer for temporal representation learning. *arXiv preprint arXiv:1906.05743*, 2019.
- [87] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. VideoBERT: A joint model for video and language representation learning. In *JCCV*, 2019.
- [88] Hao Tan and Mohit Bansal. LXMERT: Learning cross-modality encoder representations from transformers. In *EMNLP*, 2019.
- [89] Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhagen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. MovieQA: Understanding stories in movies through question-answering. In *CVPR*, 2016.
- [90] Damien Teney and Anton van den Hengel. Zero-shot visual question answering. *arXiv preprint arXiv:1611.05546*, 2016.
- [91] Otokar Tilk and Tanel Alumäe. Bidirectional recurrent neural network with attention mechanism for punctuation restoration. In *Interspeech 2016*, 2016.
- [92] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [93] Ben-Zion Vatashsky and Shimon Ullman. VQA with no questions-answers training. In *CVPR*, 2020.
- [94] Weining Wang, Yan Huang, and Liang Wang. Long video question answering: A matching-guided attention model. *Pattern Recognition*, 2020.
- [95] Thomas Winterbottom, Sarah Xiao, Alistair McLean, and Noura Al Moubayed. On modality bias in the TVQA dataset. *arXiv preprint arXiv:2012.10210*, 2020.
- [96] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016.
- [97] Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. NEX-TQA: Next phase of question-answering to explaining temporal actions. In *CVPR*, 2021.
- [98] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*, 2018.
- [99] Caiming Xiong, Stephen Merity, and Richard Socher. Dynamic memory networks for visual and textual question answering. In *ICML*, 2016.
- [100] 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 international conference on Multimedia*, 2017.
- [101] Huijuan Xu and Kate Saenko. Ask, attend and answer: Exploring question-guided spatial attention for visual question answering. In *ECCV*, 2016.
- [102] Hongyang Xue, Wenqing Chu, Zhou Zhao, and Deng Cai. A better way to attend: Attention with trees for video question answering. *IEEE Transactions on Image Processing*, 2018.
- [103] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Just ask: Learning to answer questions from millions of narrated videos. In *ICCV*, 2021. <https://arxiv.org/abs/2012.00451>.
- [104] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. TubeDETR: Spatio-Temporal Video Grounding with Trans-

- formers. In *CVPR*, 2022.
- [105] Zekun Yang, Noa Garcia, Chenhui Chu, Mayu Otani, Yuta Nakashima, and Haruo Takemura. BERT representations for video question answering. In *WACV*, 2020.
- [106] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. In *CVPR*, 2016.
- [107] Kaichun Yao, Libo Zhang, Tiejian Luo, Lili Tao, and Yanjun Wu. Teaching machines to ask questions. In *IJCAI*, 2018.
- [108] Yunan Ye, Zhou Zhao, Yimeng Li, Long Chen, Jun Xiao, and Yueting Zhuang. Video question answering via attribute-augmented attention network learning. In *ACM SIGIR*, 2017.
- [109] Weijiang Yu, Haoteng Zheng, Mengfei Li, Lei Ji, Lijun Wu, Nong Xiao, and Nan Duan. Learning from inside: Self-driven siamese sampling and reasoning for video question answering. In *NeurIPS*, 2021.
- [110] 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*, 2019.
- [111] Amir Zadeh, Michael Chan, Paul Pu Liang, Edmund Tong, and Louis-Philippe Morency. Social-IQ: A question answering benchmark for artificial social intelligence. In *CVPR*, 2019.
- [112] Rowan Zellers, Ximing Lu, Jack Hessel, Youngjae Yu, Jae Sung Park, Jize Cao, Ali Farhadi, and Yejin Choi. MERLOT: Multimodal neural script knowledge models. In *NeurIPS*, 2021.
- [113] Kuo-Hao Zeng, Tseng-Hung Chen, Ching-Yao Chuang, Yuan-Hong Liao, Juan Carlos Niebles, and Min Sun. Leveraging video descriptions to learn video question answering. In *AAAI*, 2017.
- [114] Zheng-Jun Zha, Jiawei Liu, Tianhao Yang, and Yongdong Zhang. Spatiotemporal-textual co-attention network for video question answering. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 2019.
- [115] Zhou Zhao, Shuwen Xiao, Zehan Song, Chujie Lu, Jun Xiao, and Yueting Zhuang. Open-ended video question answering via multimodal conditional adversarial networks. *IEEE Transactions on Image Processing*, 2020.
- [116] Zhou Zhao, Qifan Yang, Deng Cai, Xiaofei He, Yueting Zhuang, Zhou Zhao, Qifan Yang, Deng Cai, Xiaofei He, and Yueting Zhuang. Video question answering via hierarchical spatio-temporal attention networks. In *IJCAI*, 2017.
- [117] Zhou Zhao, Zhu Zhang, Shuwen Xiao, Zhou Yu, Jun Yu, Deng Cai, Fei Wu, and Yueting Zhuang. Open-ended long-form video question answering via adaptive hierarchical reinforced networks. In *IJCAI*, 2018.
- [118] Luwei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason J Corso, and Jianfeng Gao. Unified vision-language pre-training for image captioning and VQA. In *AAAI*, 2020.
- [119] Qingyu Zhou, Nan Yang, Furu Wei, Chuanqi Tan, Hangbo Bao, and Ming Zhou. Neural question generation from text: A preliminary study. In *National CCF Conference on Natural Language Processing and Chinese Computing*, 2017.
- [120] Linchao Zhu and Yi Yang. ActBERT: Learning global-local video-text representations. In *CVPR*, 2020.
- [121] Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *ICCV*, 2015.
- [122] Yueting Zhuang, Dejing Xu, Xin Yan, Wenzhuo Cheng, Zhou Zhao, Shiliang Pu, and Jun Xiao. Multichannel attention refinement for video question answering. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 2020.



## APPENDIX

In this Appendix, we give additional architecture details for our VideoQA model in Section A and additional implementation details in Section B.

### APPENDIX A VIDEOQA ARCHITECTURE

Our architecture, detailed in Figure 13, has two main modules: (i) a video-question multi-modal transformer (top) and (ii) an answer transformer (bottom). Details are given next.

**Video-question multi-modal transformer.** The input video representation, obtained from a fixed S3D model [98], is composed of  $t$  features denoted  $v = [v_1, \dots, v_t] \in \mathbb{R}^{d_v \times t}$  where  $d_v$  is the dimension of the video features, and  $t$  is the number of extracted features, one per second. The contextualized representation of the question, provided by the DistilBERT model [78], is composed of  $l$  token embeddings denoted as  $q = [q_1, \dots, q_l] \in \mathbb{R}^{d_q \times l}$  where  $d_q$  is the dimension of the DistilBERT embedding and  $l$  is the number of tokens in the question. The inputs to our video-question multi-modal transformer are then defined as a concatenation of question token embeddings and video features

$$u(v, q) = [\tilde{q}_1, \dots, \tilde{q}_l, \tilde{v}_1, \dots, \tilde{v}_t] \in \mathbb{R}^{d \times (l+t)}, \quad (2)$$

with

$$\tilde{q}_s = dp(\sigma(W_q q_s + b_q) + pos_s + mod_q), \quad (3)$$

and

$$\tilde{v}_s = dp(\sigma(W_v v_s + b_v) + pos_s + mod_v), \quad (4)$$

where  $W_q \in \mathbb{R}^{d_q \times d}$ ,  $b_q \in \mathbb{R}^d$ ,  $W_v \in \mathbb{R}^{d_v \times d}$ ,  $b_v \in \mathbb{R}^d$  and learnable parameters,  $mod_q \in \mathbb{R}^d$  and  $mod_v \in \mathbb{R}^d$  are learnt modality encodings for video and question, respectively, and  $[pos_1, \dots, pos_{l+t}] \in \mathbb{R}^{d \times (l+t)}$  are fixed sinusoidal positional encodings.  $\sigma$  is a Gaussian Error Linear Unit [36] followed by a Layer Normalization [8] and  $dp$  refers to Dropout [84].

The multi-modal transformer is a transformer with  $N$  layers,  $h$  heads, dropout probability  $p_d$ , and hidden dimension  $d_h$ . The outputs of the multi-modal transformer  $[Q_1, \dots, Q_l, V_1 \dots V_t] \in \mathbb{R}^{d \times (l+t)}$  are contextualized representations over tokens in the question and temporal video representations. Finally, the fused video-question embedding  $f(v, q)$  is obtained as

$$F(Q_1) = W_{vq} dp(Q_1) + b_{vq}, \quad (5)$$

where  $W_{vq} \in \mathbb{R}^{d \times d}$ ,  $b_{vq} \in \mathbb{R}^d$  are learnable parameters and  $Q_1$  is the multi-modal contextualized embedding of the [CLS] token in the question, as shown in Figure 7.

**Answer transformer.** The contextualized representation of the answer, provided by the DistilBERT model [78], is composed of  $m$  token embeddings denoted as  $a = [a_1, \dots, a_m] \in \mathbb{R}^{d_a \times m}$  where  $d_a$  is the dimension of the DistilBERT embedding and  $m$  is the number of tokens in the answer. Our answer embedding  $g(a)$  is then obtained as

$$G(a_1) = W_a a_1 + b_a, \quad (6)$$

where  $W_a \in \mathbb{R}^{d_a \times d}$ ,  $b_a \in \mathbb{R}^d$  are learnable parameters and  $a_1$  is the contextualized embedding of the [CLS] token in the answer, as shown in Figure 13.

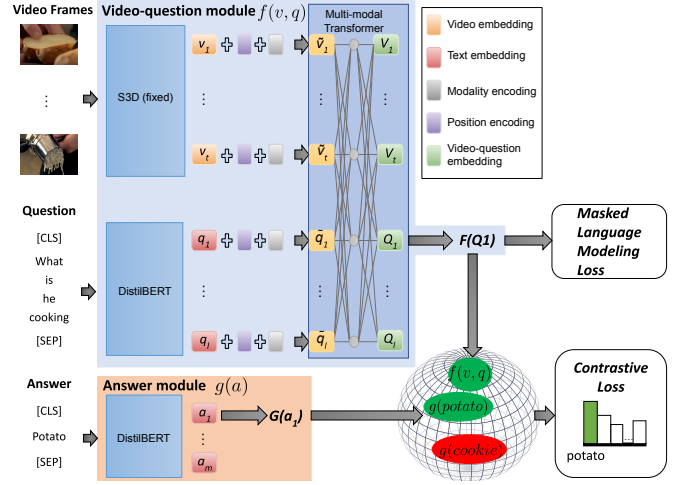


Fig. 13: **Overview of our VideoQA training architecture.** Our model is composed of a video-question module  $f$  based on a multi-modal transformer (top) and an answer module  $g$  based on DistilBERT [78] encoder (bottom). For pretraining, we use a contrastive loss and a masked language modeling loss (right).

### APPENDIX B ADDITIONAL EXPERIMENTAL DETAILS

Below we include additional details regarding the VideoQA generation, hyperparameter settings, training, masked language modeling and pretraining on HowTo100M.

**VideoQA generation.** The input sequence to the answer extractor and question generation transformers are truncated and padded up to a maximum of 32 tokens. The question decoding is done with beam search keeping track of the 4 most probable states at each level of the search tree. We have used the original captions (including stop words) from the HowTo100M dataset [69] and removed word repetitions from adjacent clips.

**VideoQA model.** We use the following hyperparameters:  $l = 20$ ,  $t = 20$ ,  $m = 10$ ,  $d = 512$ ,  $d_h = 2048$ ,  $N = 2$ ,  $H = 8$ ,  $p_d = 0.1$ ,  $d_q = d_a = 768$ ,  $d_v = 1024$ . The video features are sampled at equally spaced timestamps, and padded to length  $t$ . Sequences of question and answer tokens are truncated and padded to length  $l$  and  $m$ , respectively. Attention is computed only on non-padded sequential video and question features.

**Training.** For finetuning, we use a cosine annealing learning rate schedule with initial value of  $1 \times 10^{-5}$ . We use the Adam optimizer with batch size of 256 and training runs for 20 epochs. The final model is selected by the best performance on the validation set.

**Masked Language Modeling.** For the masked language modeling objective, a token is corrupted with a probability 15%, and replaced 80% of the time with [MASK], 10% of the time with the same token and 10% of the time with a randomly sampled token. To guess which token is masked, each sequential question output  $Q_i$  of the multi-modal transformer is classified in a vocabulary of 30,522 tokens, and we use a cross-entropy loss.

**Pretraining on HowTo100M.** For video-text cross-modal matching, we sample one video negative and one text negative per (positive) video-text pair, and use a binary cross-entropy loss. The cross-modal matching module is used to perform zero-shot VideoQA for the variant  $VQA-T$  trained on HowTo100M, by computing scores for  $f(v, [q, a])$  for all possible answers  $a$ , for each video-question pair  $(v, q)$ .