

Step Differences in Instructional Video

Tushar Nagarajan, Lorenzo Torresani
FAIR, Meta

Abstract

Comparing a user video to a reference how-to video is a key requirement for AR/VR technology delivering personalized assistance tailored to the user’s progress. However, current approaches for language-based assistance can only answer questions about a single video. We propose an approach that first automatically generates large amounts of visual instruction tuning data involving pairs of videos from HowTo100M by leveraging existing step annotations and accompanying narrations, and then trains a video-conditioned language model to jointly reason across multiple raw videos. Our model achieves state-of-the-art performance at identifying differences between video pairs and ranking videos based on the severity of these differences, and shows promising ability to perform general reasoning over multiple videos.

1. Introduction

Instructional *how-to* videos are an important medium for learning new skills that offer in-depth visual demonstrations of complex procedural activities. In turn, they serve as a valuable resource for building AR/VR assistants that can guide a user through a procedural activity, by aligning user activity to a reference how-to video. Instructional videos have thus been the subject of several recent datasets and benchmarks that are driving new research [6, 9, 41, 45, 48, 59, 68, 69].

A key requirement for such systems is the ability to compare and contrast the user’s execution of a step in the activity with the reference video step, to highlight similarities and differences between them. For example, to let the user know that they used too much detergent (while doing laundry) or that the gravy is too thick (while cooking). This ability has direct value for personalized assistance applications such as progress tracking, mistake detection and surfacing user-activity driven tips.

More generally, reasoning about a video with respect to a *reference video* is a fundamental problem for video understanding that has value for fine-grained video retrieval [11, 53, 54, 57] (e.g., to browse internet videos for “this movie

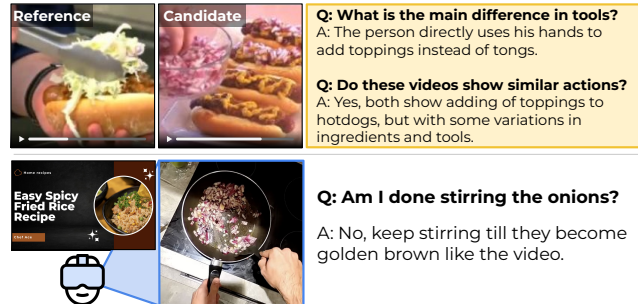


Figure 1. **Main idea.** **Top:** We train models to compare two videos showing the same high-level keystone and to describe their differences (e.g., in tools, ingredients, technique). **Bottom:** Once trained, such models can then help answer questions about a user’s activity compared to a reference (e.g., an internet how-to video) like “did I do this step right?” or “am I done yet?”.

scene, but in a forest”), step detection [6, 45, 46, 48, 69] (e.g., to recognize subtle variations in keysteps) and multi-video question answering and reasoning [5, 36] (e.g., to answer comparative questions like “which video uses the least amount of oil?”).

Despite its importance, there has been limited work on comparing videos. Prior work has explored *change captioning* in images [14, 16, 17, 34, 39, 70], however these works typically consider *pixel-level* differences (e.g., missing or moved objects; changed background objects) in static scenes (e.g., the same parking lot; the same tabletop), or in synthetically generated datasets [34]. They do not consider important semantic differences in activities (e.g., differences in tool use, subtle variations in actions and techniques or visual differences due to state changes), which together with the low-level visual differences, form a complete picture of human-object interactions.

To address these limitations, we propose a video-conditioned language model (VCLM) approach to directly compare two videos of same step in a procedural activity. Specifically, we propose the *difference question answering* task: given a reference and a candidate video, a model must answer a question that involves reasoning across both videos (e.g., what are the differences in tools? techniques?; do the two videos show the same activity?). Such a model that

effectively relates user activity to a reference video, can then provide detailed context to answer more general questions such as “what did I do wrong compared to the reference” or “am I done yet?”. See Fig. 1.

An important practical question is how to source the supervision to train such a model, given that existing video datasets only contain individual videos with captions. Moreover, meaningful differences are not guaranteed to exist between arbitrary pairings of videos. We therefore automatically generate training data from existing large-scale instructional video datasets annotated with keysteps and speech narrations describing what instructors are doing [9, 29]. We pair clips of the *same keystone* (e.g., two videos of a person “stir frying the rice until it is dark yellow”) but from distinct videos to allow for variations between them. For example the first video may use a cast iron pan versus a steel wok or the person may be tossing the food in the pan vs stirring with a spatula. We then leverage recent large language models [50] to generate questions and answers about the similarities and differences between the two videos given their visual descriptions, speech narrations, and visible objects as context. Inspired by work in visual instruction tuning of language models [27, 65], we finally fine-tune a *video-pair* conditioned language model with the collected dataset. The resulting model has the ability to cross-reference videos to compare them, and more generally answer questions that require joint reasoning about both videos simultaneously.

To evaluate our model, we collect a manually annotated dataset of 6292 video pairs with $\sim 36k$ difference captions spanning 5 categories, as well as scores for how severe the differences are. We set up the first benchmark for video comparisons and evaluate models on their ability to describe the differences in specific categories (e.g., “What are the differences in tools? techniques?”) and to rank videos based on their differences (e.g., “Which video shows the least different technique?”). Our models trained with weak-supervision from automatically generated data achieve state-of-the-art results on our benchmark, highlighting its value for personalized assistance applications. Our benchmark will be hosted publicly, to allow the community to make progress towards this under-explored task.

2. Related work

Instructional video understanding Recent large-scale instructional video datasets [6, 9, 41, 45, 48, 59, 68, 69] have facilitated research in step captioning [63, 68], step detection [6, 45, 46, 48, 69], temporal grounding [2, 7, 13, 18, 28], vision-language representation learning [3, 25, 38, 66] and video question answering [56, 58, 60, 62] to name a few. In all these approaches, the goal is to process a single video and then caption, answer questions or temporally localize an action or text within it. While we are also interested in the space of procedural videos in the context of personalized

language-based assistance, in contrast, we develop methods to compare and contrast *multiple videos* — namely a reference video and a candidate video — in order to identify differences and answer comparative questions about them.

Visual differences in images Prior work has studied visual differences in images in the context of attributes [8, 10, 33, 61] (e.g., which shoe is more formal) to facilitate fine-grained recognition. More relevant to our work, *change captioning* [14, 16, 17, 34, 39, 70] involves describing the differences between two images as a text caption. Other work defines differences as 2D bounding boxes [42, 43] or semantic maps [35] for regions that differ. More recently, VCLM models have been trained with “spot-the-difference” data from the above with a similar goal of identifying image differences [20]. In all these cases, the two images typically involve the same scene from multiple viewpoints or over time (e.g., surveillance footage) or are constructed from synthetic images (e.g., 3D geometric shapes re-arranged on a table). The resulting differences therefore focus on simple visual cues like missing or moved objects. More recent approaches use visual differences to retrieve videos [4], however they assume the difference is known (to retrieve a relevant video) rather than identifying and describing it. In contrast, we compare across distinct video clips that show the same high-level keystone. As a result, the difference captions characterize complex variations that arise naturally from the availability of tools and ingredients, differing skill / technique or personal preference.

Visual instruction tuning of language models Given the recent success of large language models (LLMs), several efforts have tried to adapt them for use with various modalities including images, videos, audio etc., typically by aligning captions to modalities or instruction tuning [12, 21, 26, 27, 31, 32, 64, 65]. All these approaches typically use text captions or generate instruction tuning data based on a *single image or video*. In contrast, we generate instruction data for pairs of videos (a reference, and a target video) to allow vision conditioned language models to jointly reason about them both. Some approaches do train on multiple images interleaved with text [1, 19, 51], however they do not support instructions at inference, and instead rely on in-context few-shot prompting to respond. In contrast, our approach can respond to arbitrary questions about a video with respect to a reference clip.

3. Approach

Our goal is to train models to answer questions about a video in the context of a reference video, by jointly reasoning about the two. The problem is two-fold: where do we source data of pairs of videos with relevant questions

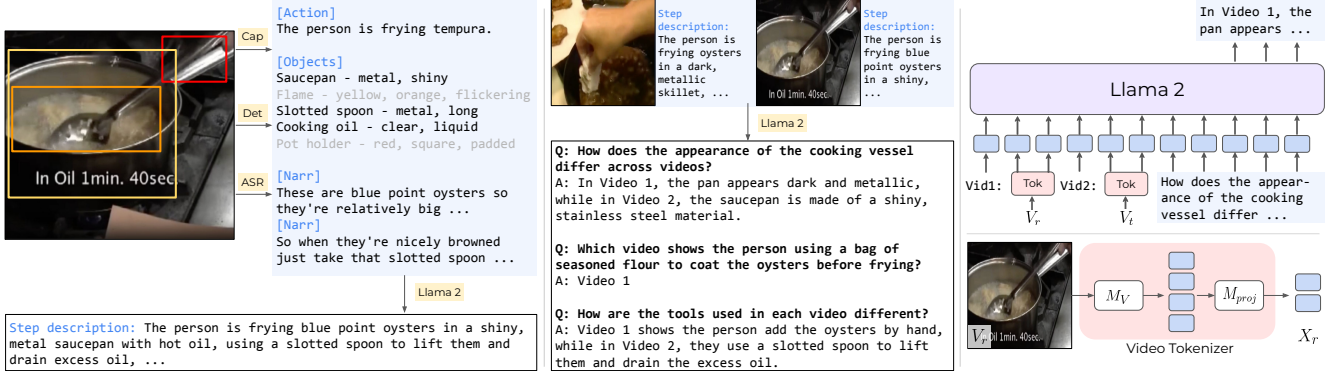


Figure 2. **Step differences framework.** We first generate a comprehensive step description including information from action captions, object detections and ASR narrations (left panel). We then select pairs of clips with similar step descriptions, and automatically generate questions and answers that compare the two (center panel, Sec. 3.2). Finally, we instruction-tune an LLM to generate answers conditioned on the generated questions and encoded representations of both videos (right panel, Sec. 3.3). Once trained, the model directly operates on video clips to compare them, without the need for captions, ASR or object detections.

to train such models and what model architectures support training with multiple videos? For the former, we turn to automatically generating this data using large-language models (LLMs) parsing narrated video from existing datasets. For the latter, we use vision-conditioned language models (VCLMs) — a powerful class of models for single-video question answering — adapted to our multi-video setting. In the following, we first formally define our task (Sec. 3.1). Next, we describe our automatic training data generation pipeline (Sec. 3.2). Finally, we discuss training and downstream inference (Sec. 3.3).

3.1. Task definition

We require models that collectively answer questions about two videos. Formally, given a reference video V_r , a candidate video V_c and a question q , models must produce a corresponding answer a . This formulation is an extension of standard video question answering or captioning [67] with a response that is additionally conditioned on a reference video. The questions can take various forms, for example “How is the dough being prepared differently in Video 2?”; “What is the similarity in mixing techniques between the two videos?”. Critically, these questions all share the assumption that a single video alone (either the reference or the candidate) is insufficient to answer the question — reasoning over both videos is required.

In our experiments, we train models with a diverse set of automatically-generated question-answer pairs. At test time, we focus on *step differences*, where the q is of the form “what is the main difference between these videos in the category g ” and g is the difference category (e.g., ingredients, techniques, etc.). This structure captures a representative range of fine-grained differences, and allows for consistent evaluation of models as we will show.

3.2. Step differences dataset generation

To train our models, we require a dataset of paired videos along with questions and answers (QA) relating the two in the form (V_r, V_c, q, a) . However, current video datasets typically contain individual video clips annotated for actions, narrations, or single-video QA which is incompatible with our task definition. We therefore construct this from existing video datasets using large-language models, inspired by prior work on instruction tuning [12, 20, 27, 37, 49].

Constructing this dataset from existing video datasets is non-trivial. On the one hand, selecting random pairs of videos showing very different content (e.g., sports vs. cooking) or near-identical videos (e.g., from repetitions of the same activity by the same participant) will lead to trivial differences. On the other hand, naively selecting video pairs of the same class in action recognition datasets (e.g., “Bookbinding” or “Mowing the lawn”) will not highlight fine-grained differences of interest, and will instead focus on global differences (e.g., changes in actors or scenes). Moreover, these datasets do not come with text descriptions to construct differences from.

We therefore propose to use videos from the large-scale procedural video dataset HowTo100M [29], specifically cooking-themed videos labeled for keysteps from HT-Step [9]. Instructional videos are an ideal data source as they are narrated and show the same high-level keystone, but with variations that arise naturally from availability of tools and ingredients, differing skill / technique or personal preference.

Specifically, for two videos showing the same keystone (e.g., Slowly pour the sauce over the dumplings), we assume one is the reference V_r and the other is the candidate video V_c , with corresponding speech narrations. First, we generate descriptions of the actions and objects (including their attributes) using off-the-shelf captioning models [31]. These

models often hallucinate details in their generations, so we additionally filter object descriptions based on the scores of a pre-trained detection model [30] and filter action descriptions using visual grounding models [55]. Details about the filtering stage are in Supp. Finally, we aggregate the information from these three sources (ASR narration, filtered objects and actions) to synthesize a detailed step description for each video (Fig. 2, left panel). We then prompt a language model (in our case, Llama 2 [50]) to generate both questions and answers comparing the two videos based on their step descriptions. In short, the prompt takes the form: “Video 1: {description1}. Video 2: {description2}. Summarize the differences and generate 3 question-answer pairs comparing the two videos.” (Fig. 2, center panel). An overview of the data generation pipeline with examples at each stage can be seen in Fig. 2. See Supp. for more examples and full step description prompt details.

The resulting dataset contains QA instances over video pairs across 87740 unique video clips. Note that the LLM-generated data is noisy — they may hallucinate details that are not present in the video, misunderstand the ASR narrations, produce irrelevant questions or incorrect answers to questions. Despite this, they offer valuable *weak supervision* to train our VCLM models, as our experiments will show.

3.3. Paired video instruction tuning

We require a model that can generate natural language responses to video comparison questions in our dataset. To do this, we adapt a vision-conditioned language model (VCLM) to our multi-video setting via visual instruction tuning. In short, visual instruction tuning aligns the outputs of an image (or video) backbone to a powerful LLM to condition its responses on the visual content. This strategy has been successful in prior work for single image/video captioning and question answering [12, 21, 26, 27, 31, 64, 65]. We extend this to support comparisons across multiple videos. In our experiments, we use a Llama2 [50] LLM aligned with an Internvideo [55] backbone following prior work [31]. Note that it is possible to directly provide multiple videos to existing models by adding extra visual tokens to the input prompt, however their performance is degraded as they not trained to support this. We compare against such models.

Specifically, for an instruction-tuning instance (V_r, V_c, q, a) , we generate an instruction prompt in the Llama2 format as follows.

```
<s> [INST] <<SYS>> You are a helpful AI assistant that answers questions about a pair of videos. Answer in a single sentence. Here is the first video: {V_r}. Here is the second video: {V_c}. <</SYS>> {q} [/INST] {a}
```

We encode the text tokens in this prompt using the LLM’s pre-trained text encoder. We encode each video into a sequence of spatiotemporal tokens using a pre-trained video

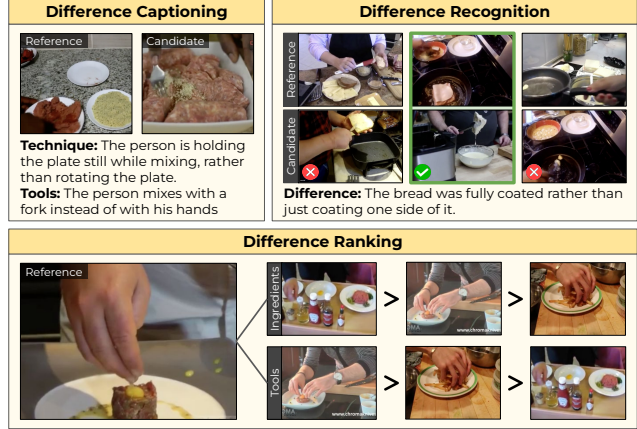


Figure 3. **Evaluation tasks.** We evaluate on describing (DiffCap), recognizing (DiffMCQ) and ranking (DiffRank) differences.

backbone M_V , and then align them to the LLM’s input space using a learnable projection module M_{proj} . The resulting encoded instruction prompt is a sequence of tokens comprising a mix of text and visual tokens, which can then be processed by the LLM (Fig. 2, right panel).

The model is trained using the original auto-regressive objective to maximize the probability of generating the answer tokens, conditioned on the question, reference and candidate video, and is trained using a standard cross-entropy loss.

$$p(X_a | X_r, X_c, X_q) = \prod_{i=1}^{|X_a|} p_{\theta}(X_{a,i} | X_r, X_c, X_q, X_{a,<i}) \quad (1)$$

$$X_r = M_{proj}(M_V(V_r)) \quad (2)$$

$$X_c = M_{proj}(M_V(V_c)), \quad (3)$$

where $X_{a,i}$ is the i -th answer token in the sequence, X_r (X_c) are the visual tokens corresponding to the reference (candidate) videos, X_q, X_a are tokens of the question and answer, $X_{a,<i}$ are answer tokens that occur before $X_{a,i}$ and θ are the learnable parameters in \mathcal{M}_{proj} .

Note that the video encoder and the LLM weights are frozen, and the loss is computed only for answer tokens. Only the projection layer is fine-tuned. Once trained, our model will be able to refer to each video, discuss their similarities and compare them. We evaluate our model by autoregressively generating text in response to various prompts coupled with reference and candidate videos.

3.4. Describing, recognizing and ranking step differences in procedural videos

Finally, we use our trained models to identify and rank fine-grained differences between pairs of video. We cast these tasks into the paired-video QA framework as follows.

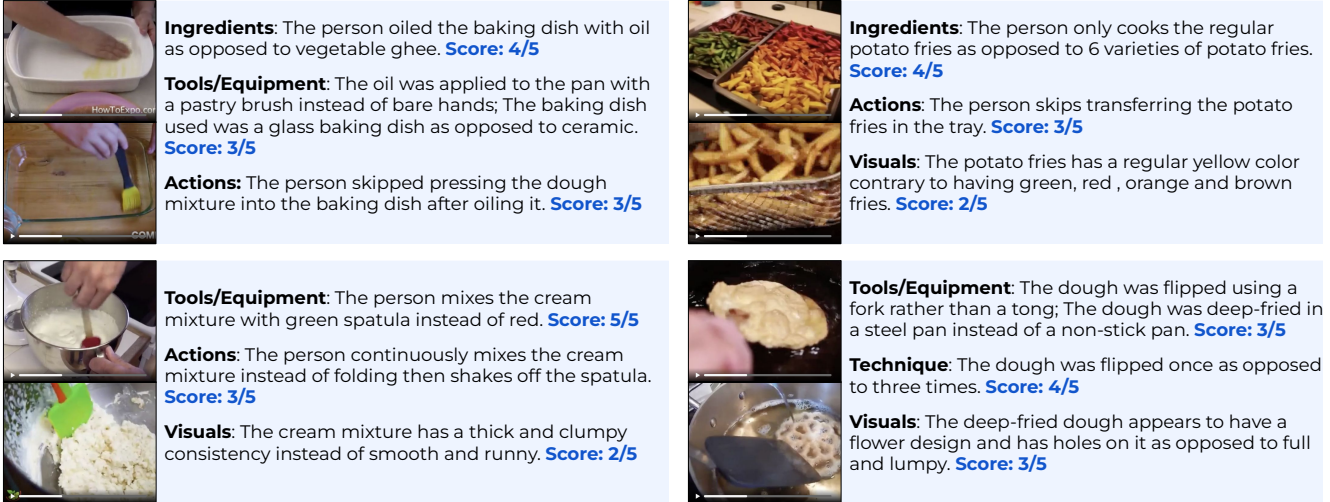


Figure 4. **StepDiff dataset samples** We annotate text describing differences in various categories and scores for *how different* the videos are in each category (1 = very different; 5 = nearly identical). More examples are in Supp.

Difference captioning (DiffCap) The goal is to generate a textual description of the differences between two videos in a specific category g (e.g., ingredients, tools). The question q takes the form “what is the main difference between these videos in the category g ”. The difference caption is generated auto-regressively using the trained model.

Difference recognition (DiffMCQ) The goal is to select the correct video pair that matches the difference caption, from a list of candidate video pairs $\{(V_r^i, V_c^i)\}_{i=1..4}$. This is a discriminative version of the captioning task above inspired by recent work in vision-language feature learning [24]. For this, we compute $p(a|V_r^i, V_c^i, q)$ — the likelihood of generating the difference text given the pair of videos following Eqn. 2 — and then select the pair with the highest score.

Difference ranking (DiffRank) The goal is to rank video instances $\{V_c^i\}_{i=1..4}$ based on how different they are to a common reference video V_r , in terms of a particular category of interest g . For this, we set q to be “do these two videos show the same g ? Answer YES or NO.”, and rank each candidate video based on the likelihood of generating “YES” as the response.

Together, these tasks are a representative suite of problems for instructional video understanding that require comparing videos along various axes. DiffCap tests how accurately a model can describe differences in natural language, DiffMCQ tests how well it can discriminate differences between videos, and DiffRank tests how well the model can assess the severity of these differences to rank them. A model for these tasks can enable applications that guide user action (e.g., to follow a reference video tutorial) or help browse through large collections of videos (e.g., to find the perfect variation of a recipe). Fig. 3 illustrates these tasks.

4. Experiments

We evaluate our VCLM model on the three step difference tasks from Sec. 3.4.

Dataset We construct a test dataset from videos in HT-Step [9]. HTStep contains videos from a large-scale procedural video dataset, HowTo100M [29] (*Cooking & Entertainment*), with temporal segments (clips) annotated for keysteps (e.g., “fry then onions until golden brown”). We manually annotate pairs of clips, where each pair corresponds to instances of the same labeled keystone, but from distinct videos. Annotators are asked to identify the main differences across 5 categories (ingredients, tools/equipment, techniques, visual differences) and write difference captions of a consistent style — what happens in the target clip, compared to what happens in the reference (e.g., “The person uses a deep fryer to fry the potatoes instead of shallow frying it in a pan”). They are then asked to score the difference caption in each category on a scale of 1-5 based on how severe the difference is, where 1 is a significant difference (e.g., swapping out a critical ingredient that would change the dish entirely) and 5 is nearly identical (e.g., minor cosmetic differences that does not affect the activity). A rubric is used to ensure consistency in scoring.

Note that this data is only used for evaluation — we exclude these pairs from the automatic training data generation pipeline described in Sec. 3.2 to ensure that the model has not seen these instances during training. In total, we collect 35988 difference captions across 6292 clip pairs, involving 8396 unique clips. See Fig. 4 for examples. Full collection details and dataset statistics are in Supp.

Baselines We compare several classes of models.

	DIFFCAP			DIFFMCQ	DIFFRANK
	BLEU	CIDER	ROGUE-L	Acc %	τ
VLEmbed (CLIP) [40]	–	–	–	0.396	0.127
VLEmbed (InternVideo) [55]	–	–	–	0.451	0.170
Socratic (BLIP-2) [22]	0.159	0.036	0.169	0.335	0.022
Socratic (LLaVA) [27]	0.151	0.031	0.166	0.332	0.004
Socratic (Step desc.)	0.141	0.020	0.172	0.392	0.013
VCLM (LLaVA) [27]	0.211	0.069	0.199	0.381	0.019
VCLM (AnyMAL) [31]	0.209	0.115	0.196	0.471	0.032
Interleaved (IDEFICS) [19]	0.217	0.080	0.210	0.376	0.081
Interleaved (AnyMAL)	0.207	0.102	0.198	0.497	0.014
StepDiff	0.223	0.104	0.215	0.541	0.181

Table 1. **Results.** Our approach outperforms three classes of baselines built on top of state-of-the-art vision-language embedding and VCLM models. VLEmbed baselines are excluded from DiffCap as they cannot generate text.

- **VLEmbed** is a class of vision-language model that embeds images or video in the same space as text, and then compares their similarity in the shared space. Video *pair* embeddings are calculated as the average of individual video embeddings¹. We use CLIP [40] and InternVideo [55].
- **Socratic** is a class of VCLMs that first converts videos into text using a captioning model, and then prompts a text-only LLM with these captions. These models are powerful, but often require complex, manually engineered prompts. We use state-of-the-art visual captioners (BLIP-2 [22], LLaVA-1.5 [27]) as well our aggregate step descriptions from Sec. 3.2. We use Llama2 to process the captions regardless of which model generated them, for fair comparisons.
- **VCLM** is a class of visual instruction-tuned language model trained for video captioning and question answering (for a single video). We directly add extra tokens for the reference video into the prompt to be consistent with our paired-video QA task. We compare LLaVA-1.5 [27] and AnyMAL [31].
- **Interleaved** is a class of models that are trained with interleaved sequences of images/videos and text, and naturally support multiple videos as inputs, but are not explicitly trained to compare them. We compare the recently proposed IDEFICS [19] and a model we train on sequences of (video, ASR) pairs from HowTo100M (training details in Supp.).

These baselines represent a spectrum of leading strategies for vision-language reasoning, including methods that directly embed video and language in the same space (VLEmbed), ones that explicitly convert videos to text and perform exclusively text-based reasoning (Socratic) and ones that perform joint vision-text reasoning on videos (VCLM, Interleaved). We ensure that each class of baselines include methods that have been trained on in-domain HowTo100M videos, while excluding the evaluation videos, to ensure fair comparisons with our approach. These are InternVideo,

¹We evaluate other aggregation strategies in Supp.

Socratic (Step desc.), VCLM (AnyMAL), and Interleaved (AnyMAL). Additional pretraining and implementation details are in Supp.

Implementation details We use the Llama2-chat-70B [50] as the base LLM for all our experiments. Following prior work [31], M_V is an Internvideo [55] video encoder that inputs 8 uniformly sampled frames from each video clip and generates 2048 spatio-temporal tokens. M_{Proj} is a 2-layer Perceiver [15] module followed by a linear layer head to output 32 tokens in the LLM’s input dimension. During training, all parameters are frozen except for M_{Proj} . StepDiff models are initialized from Interleaved model weights before finetuning (interleaved data is retained during finetuning). For baselines, we use the largest available versions of models — InstructBLIP (Vicuna13B), LLaVA (Vicuna13B), AnyMAL (70B), IDEFICS (80B). Full implementation and training details are in Supp.

4.1. Difference captioning

We first evaluate how well our model can describe differences in video pairs (DiffCap). As mentioned in Sec. 3.1, q is of the form “what is the main difference between these videos in the category g ” where g is the difference category. Since there may be multiple annotated differences in the same category, we group them together and treat them as a ground truth set, resulting in a dataset with 22292 instances. We measure standard text generation metrics including CIDER [52] and ROUGE-L [23]. Outputs are post-processed using simple string matching techniques to ensure difference captions are generated in the correct format (details in Supp). For the socratic baselines, we provide the generated caption instead of the video tokens in the prompt from Sec. 3.3. Table 1 (left) shows our results. The socratic models perform poorly as they are limited by the information contained in the base captions. It is infeasible to generate captions that exhaustively describe every aspect of a video, without knowing what is of interest, and without the risk of model hallucinations. The VCLM models perform better,

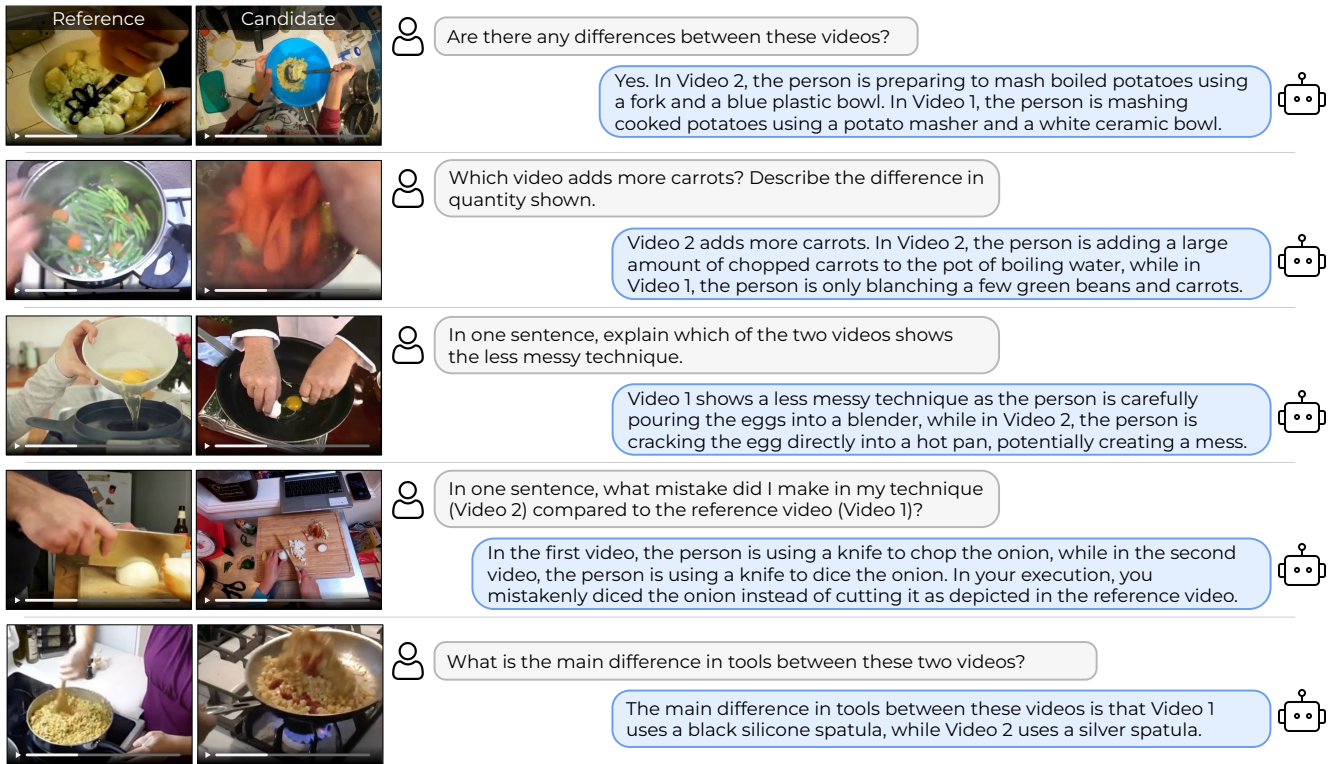


Figure 5. **Extended QA on video pairs.** Our model which can describe differences (row 1) can be prompted (i.e., queried without any form of retraining) for comparative reasoning (e.g., “why are they different?”, “how different are they?” row 2-3), or to bootstrap mistake detection (row 4). A failure case is shown in row 5 due to model hallucination.

especially when trained to process multiple videos (i.e., interleaved models), however they still fall short of our approach that can explicitly compare and contrast videos. The example in Fig. 6 highlights the sensitivity of Socratic models to input captions (e.g., the reference caption did not mention the use of hands), and shows how VCLM models tend to hallucinate details. Our approach can correctly describe the difference. See Supp for more examples.

4.2. Difference recognition

While the captioning metrics are informative, they are based on word overlap statistics, and do not always capture the semantics of the text well. To address this, we evaluate on DiffMCQ – the discriminative version of the captioning task. We adapt the same dataset from DiffCap, except we sample a single difference caption for each category if there are multiple differences present. Further, we sample three *negative* video pairs from other instances in the dataset that involve similar objects and actions (details in Supp). For the VLEmbd baselines, we score each video pair using the cosine similarity between their average visual embeddings and the text embedding of the difference caption. We compare variants of this baseline considering only the reference or target in Supp. For all LLM-based baselines, we compute the

likelihood of generating the difference caption for each video pair, under each model as discussed in Sec. 3.1. We evaluate top-1 accuracy. Table 1 (center) shows our results. The joint feature embedding models capture some semantics, but are insufficient for identifying differences. Socratic models have a similar trend to captioning results, however models trained on step differences show large improvements, highlighting the value of careful curation for generating captions. Among VCLM models, ones that have seen in-domain HowTo100M videos during training have an edge over the others (i.e., LLaVA, IDEFICS), with interleaved models again being superior. Our model outperforms all these approaches with a 5% accuracy improvement over the strongest baseline. Fig. 7 shows performance increases by difference category, over the weakest baseline (Socratic). Our approach shows large relative improvements on most categories especially in technique and tool use (both 46%), which require fine-grained action understanding.

4.3. Difference ranking

Finally, we evaluate how well our model can rank videos based on the severity of differences compared to a common reference (DiffRank). Each reference video in the dataset is paired with four target videos, scored along each category

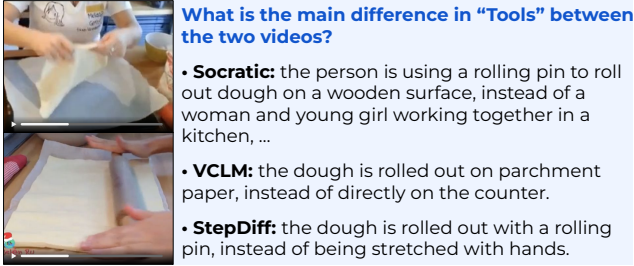


Figure 6. **DiffCap baselines.** Our approach can describe differences without relying on input captions (like Socratic) and is less prone to hallucinating details (like VCLM).




Figure 7. **DiffMCQ performance by category.** Improvements are reported over the weakest baseline (Socratic).

axis. For example, some videos may be very similar in terms of technique, but very different in terms of ingredients. We only retain instances where there is a clear ranking (i.e., no more than one tie in scores). The resulting dataset contains 3932 instances involving 5746 unique clips.

As discussed in Sec. 3.1, we rank each target video candidate based on the likelihood of producing the response “YES” when asked whether it is similar to the reference. We use the Kendall’s τ rank correlation metric to evaluate how well the generated ranking compares to the ground truth ranking annotators provide. Table 1 (right) shows our results. Unlike the previous tasks, the joint embedding models perform better than the LLM based baselines for two reasons. First, similarity in the embedding space directly translates to a score for ranking, rather than relying on computing “YES” token probability as a proxy for this. Second, there is a high correlation between the rankings across categories for the same set of instances ($\tau = 0.63$). For example, videos ranked low in similarity for *tools* when compared to a reference are often also ranked low for *technique*. Despite these issues, our approach is able to outperform all baselines, showcasing its versatility as a retrieval and ranking model.

4.4. Extending QA beyond atomic differences

Next, we show how our model can be prompted to answer questions beyond just “describe the differences”. LLMs have

	DIFFCAP			DIFFMCQ	DIFFRANK
	BLEU	CIDER	ROGUE-L	Acc %	τ
StepDiff	0.223	0.104	0.215	0.541	0.181
w/o interleaved data	0.214	0.094	0.210	0.499	0.185
w/o QA filtering	0.222	0.096	0.212	0.516	0.120
w/ 13B LLM	0.216	0.124	0.205	0.527	0.175

Table 2. **Ablation experiments.** Impact of retaining interleaved training data, careful filtering of QA training data and LLM size.

shown remarkable abilities for complex, multi-step reasoning in text – our training framework unlocks the same kind of reasoning for multiple videos, based on their differences. In Fig. 5, we show some examples of this. Our model is able to naturally describe differences as it was trained for this task (row 1), but also has the ability to perform comparative reasoning (row 2-3) or explain mistakes (row 4). We show a failure case in row 5, where the model hallucinates content – a characteristic feature of the LLM models it is built upon. Moreover, our model works with egocentric video (row 1, 4), despite being trained on largely third-person video content (HowTo100M), which is promising for AR/VR user assistance applications.

4.5. Ablation experiments

Finally, we ablate several design choices in our model in Table 2. As mentioned in Sec. 4, we finetune models on both interleaved ASR data as well as our generated pair QA data. Without the interleaved data, the model performance drops on two tasks, likely due to catastrophic forgetting (*w/o interleaved data*). Next, we show the importance of filtering the generated QA data (*w/o QA filtering*), given the high likelihood of hallucinations produced by the LLM. Finally, we swap out the 70B LLM model for a smaller sized one (*w/ 13B LLM*), causing the performance to drop, though not significantly.

5. Conclusion

We proposed StepDiff, a video-conditioned language model (VCLM) that can compare and contrast videos to reveal fine-grained differences between them. We propose an approach that can automatically generate instruction-following paired-video QA training data from large-scale procedural video data, and a manually curated benchmark to evaluate models. Our experiments on describing and identifying differences, as well on ranking videos based on differences demonstrate the value of our approach for personalized assistance applications. Future work can leverage our work for personalized retrieval (e.g., retrieve content based on user-activity), or multi-video QA beyond instructional videos.

Acknowledgements Thanks to Efi Mavroudi, Huiyu Wang, Triantafyllos Afouras and Yale Song for helpful discussions; Kumar Ashutosh and Suyog Jain for help with annotation tooling and collection; Austin Miller and Honey Manglani for managing the annotator workforce.

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. 2
- [2] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *Proceedings of the IEEE international conference on computer vision*, pages 5803–5812, 2017. 2
- [3] Kumar Ashutosh, Rohit Girdhar, Lorenzo Torresani, and Kristen Grauman. Hiervl: Learning hierarchical video-language embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23066–23078, 2023. 2
- [4] Kumar Ashutosh, Zihui Xue, Tushar Nagarajan, and Kristen Grauman. Detours for navigating instructional videos. In *CVPR*, 2024. 2
- [5] Ankan Bansal, Yuting Zhang, and Rama Chellappa. Visual question answering on image sets. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pages 51–67. Springer, 2020. 1
- [6] Siddhant Bansal, Chetan Arora, and CV Jawahar. My view is the best view: Procedure learning from egocentric videos. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIII*, pages 657–675. Springer, 2022. 1, 2
- [7] Peijun Bao, Qian Zheng, and Yadong Mu. Dense events grounding in video. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 920–928, 2021. 2
- [8] Steven Chen and Kristen Grauman. Compare and contrast: Learning prominent visual differences. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1267–1276, 2018. 2
- [9] Daffy. Htstep. In *NeurIPS (Datasets and Benchmarks)*, 2023. 1, 2, 3, 5
- [10] Maxwell Forbes, Christine Kaeser-Chen, Piyush Sharma, and Serge Belongie. Neural naturalist: generating fine-grained image comparisons. *arXiv preprint arXiv:1909.04101*, 2019. 2
- [11] Sonam Goenka, Zhaoheng Zheng, Ayush Jaiswal, Rakesh Chada, Yue Wu, Varsha Hedau, and Pradeep Natarajan. Fashionvlp: Vision language transformer for fashion retrieval with feedback. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14105–14115, 2022. 1
- [12] Tao Gong, Chengqi Lyu, Shilong Zhang, Yudong Wang, Miao Zheng, Qian Zhao, Kuikun Liu, Wenwei Zhang, Ping Luo, and Kai Chen. Multimodal-gpt: A vision and language model for dialogue with humans. *arXiv preprint arXiv:2305.04790*, 2023. 2, 3, 4
- [13] Tengda Han, Weidi Xie, and Andrew Zisserman. Temporal alignment networks for long-term video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2906–2916, 2022. 2, 3
- [14] Mehrdad Hosseinzadeh and Yang Wang. Image change captioning by learning from an auxiliary task. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2725–2734, 2021. 1, 2
- [15] Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General perception with iterative attention. In *International conference on machine learning*, pages 4651–4664. PMLR, 2021. 6, 3
- [16] Harsh Jhamtani and Taylor Berg-Kirkpatrick. Learning to describe differences between pairs of similar images. *arXiv preprint arXiv:1808.10584*, 2018. 1, 2
- [17] Hoeseong Kim, Jongseok Kim, Hyungseok Lee, Hyunsung Park, and Gunhee Kim. Agnostic change captioning with cycle consistency. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2095–2104, 2021. 1, 2
- [18] Hilde Kuehne, Juergen Gall, and Thomas Serre. An end-to-end generative framework for video segmentation and recognition. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–8. IEEE, 2016. 2
- [19] Hugo Laurençon, Lucile Saulnier, Léo Tronchon, Stas Bekman, Amanpreet Singh, Anton Lozhkov, Thomas Wang, Siddharth Karamcheti, Alexander M Rush, Douwe Kiela, et al. Obelisc: An open web-scale filtered dataset of interleaved image-text documents. *arXiv preprint arXiv:2306.16527*, 2023. 2, 6, 4
- [20] Bo Li, Yuanhan Zhang, Liangyu Chen, Jinghao Wang, Fanyi Pu, Jingkang Yang, Chunyuan Li, and Ziwei Liu. Mimic-it: Multi-modal in-context instruction tuning. *arXiv preprint arXiv:2306.05425*, 2023. 2, 3
- [21] Bo Li, Yuanhan Zhang, Liangyu Chen, Jinghao Wang, Jingkang Yang, and Ziwei Liu. Otter: A multi-modal model with in-context instruction tuning. *arXiv preprint arXiv:2305.03726*, 2023. 2, 4
- [22] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 6, 4
- [23] Chin-Yew Lin and Franz Josef Och. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL-04)*, pages 605–612, 2004. 6
- [24] Kevin Qinghong Lin, Jinpeng Wang, Mattia Soldan, Michael Wray, Rui Yan, Eric Z XU, Difei Gao, Rong-Cheng Tu, Wen-zhe Zhao, Weijie Kong, et al. Egocentric video-language pretraining. *Advances in Neural Information Processing Systems*, 35:7575–7586, 2022. 5
- [25] Xudong Lin, Fabio Petroni, Gedas Bertasius, Marcus Rohrbach, Shih-Fu Chang, and Lorenzo Torresani. Learning to recognize procedural activities with distant supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13853–13863, 2022. 2

- [26] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*, 2023. 2, 4
- [27] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023. 2, 3, 4, 6
- [28] Effrosyni Mavroudi, Triantafyllos Afouras, and Lorenzo Torresani. Learning to ground instructional articles in videos through narrations. *arXiv preprint arXiv:2306.03802*, 2023. 2
- [29] 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 *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2630–2640, 2019. 2, 3, 5, 1
- [30] Matthias Minderer, Alexey Gritsenko, and Neil Houlsby. Scaling open-vocabulary object detection. *arXiv preprint arXiv:2306.09683*, 2023. 4, 1
- [31] Seungwhan Moon, Andrea Madotto, Zhaojiang Lin, Tushar Nagarajan, Matt Smith, Shashank Jain, Chun-Fu Yeh, Prakash Murugesan, Peyman Heidari, Yue Liu, et al. Anymal: An efficient and scalable any-modality augmented language model. *arXiv preprint arXiv:2309.16058*, 2023. 2, 3, 4, 6, 1
- [32] OpenAI. Gpt4v. ???, 2023. 2
- [33] Devi Parikh and Kristen Grauman. Relative attributes. In *2011 International Conference on Computer Vision*, pages 503–510. IEEE, 2011. 2
- [34] Dong Huk Park, Trevor Darrell, and Anna Rohrbach. Robust change captioning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4624–4633, 2019. 1, 2
- [35] Jin-Man Park, Jae-Hyuk Jang, Sahng-Min Yoo, Sun-Kyung Lee, Ue-Hwan Kim, and Jong-Hwan Kim. Changesim: Towards end-to-end online scene change detection in industrial indoor environments. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8578–8585. IEEE, 2021. 2
- [36] Abhirama Subramanyam Penamakuri, Manish Gupta, Mithun Das Gupta, and Anand Mishra. Answer mining from a pool of images: Towards retrieval-based visual question answering. *arXiv preprint arXiv:2306.16713*, 2023. 1
- [37] Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. Instruction tuning with gpt-4. *arXiv preprint arXiv:2304.03277*, 2023. 3
- [38] Shraman Pramanick, Yale Song, Sayan Nag, Kevin Qinghong Lin, Hardik Shah, Mike Zheng Shou, Rama Chellappa, and Pengchuan Zhang. Egovlpv2: Egocentric video-language pre-training with fusion in the backbone. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5285–5297, 2023. 2
- [39] Yue Qiu, Shintaro Yamamoto, Kodai Nakashima, Ryota Suzuki, Kenji Iwata, Hirokatsu Kataoka, and Yutaka Satoh. Describing and localizing multiple changes with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1971–1980, 2021. 1, 2
- [40] 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. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 6, 5
- [41] Michaela Regneri, Marcus Rohrbach, Dominikus Wetzel, Stefan Thater, Bernt Schiele, and Manfred Pinkal. Grounding action descriptions in videos. *Transactions of the Association for Computational Linguistics*, 1:25–36, 2013. 1, 2
- [42] Ragav Sachdeva and Andrew Zisserman. The change you want to see (now in 3d). In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2060–2069, 2023. 2
- [43] Ragav Sachdeva and Andrew Zisserman. The change you want to see. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3993–4002, 2023. 2
- [44] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 3
- [45] Fadime Sener, Dibiyadip Chatterjee, Daniel Shelepov, Kun He, Dipika Singhanian, Robert Wang, and Angela Yao. Assembly101: A large-scale multi-view video dataset for understanding procedural activities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21096–21106, 2022. 1, 2
- [46] Gunnar A Sigurdsson, Abhinav Gupta, Cordelia Schmid, Ali Farhadi, and Karteek Alahari. Actor and observer: Joint modeling of first and third-person videos. In *proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7396–7404, 2018. 1, 2
- [47] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. MpNet: Masked and permuted pre-training for language understanding. *Advances in Neural Information Processing Systems*, 33:16857–16867, 2020. 5
- [48] Yansong Tang, Dajun Ding, Yongming Rao, Yu Zheng, Danyang Zhang, Lili Zhao, Jiwen Lu, and Jie Zhou. Coin: A large-scale dataset for comprehensive instructional video analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1207–1216, 2019. 1, 2
- [49] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. Alpaca: A strong, replicable instruction-following model. *Stanford Center for Research on Foundation Models*. <https://crfm.stanford.edu/2023/03/13/alpaca.html>, 3(6):7, 2023. 3
- [50] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. 2, 4, 6
- [51] Maria Tsimpoukelli, Jacob L Menick, Serkan Cabi, SM Eslami, Oriol Vinyals, and Felix Hill. Multimodal few-shot learning with frozen language models. *Advances in Neural Information Processing Systems*, 34:200–212, 2021. 2

- [52] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575, 2015. 6
- [53] Lucas Ventura, Antoine Yang, Cordelia Schmid, and Gül Varol. Covr: Learning composed video retrieval from web video captions. *arXiv preprint arXiv:2308.14746*, 2023. 1
- [54] Nam Vo, Lu Jiang, Chen Sun, Kevin Murphy, Li-Jia Li, Li Fei-Fei, and James Hays. Composing text and image for image retrieval—an empirical odyssey. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6439–6448, 2019. 1
- [55] Yi Wang, Kunchang Li, Yizhuo Li, Yanan He, Bingkun Huang, Zhiyu Zhao, Hongjie Zhang, Jilan Xu, Yi Liu, Zun Wang, Sen Xing, Guo Chen, Junting Pan, Jiashuo Yu, Yali Wang, Limin Wang, and Yu Qiao. Internvideo: General video foundation models via generative and discriminative learning. *arXiv preprint arXiv:2212.03191*, 2022. 4, 6, 3, 5
- [56] Bo Wu, Shoubin Yu, Zhenfang Chen, Joshua B Tenenbaum, and Chuang Gan. Star: A benchmark for situated reasoning in real-world videos. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. 2
- [57] Hui Wu, Yupeng Gao, Xiaoxiao Guo, Ziad Al-Halah, Steven Rennie, Kristen Grauman, and Rogerio Feris. Fashion iq: A new dataset towards retrieving images by natural language feedback. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 11307–11317, 2021. 1
- [58] Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. Next-qa: Next phase of question-answering to explaining temporal actions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9777–9786, 2021. 2
- [59] Yale. Goalstep. In *NeurIPS (Datasets and Benchmarks)*, 2023. 1, 2
- [60] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Learning to answer visual questions from web videos. *arXiv preprint arXiv:2205.05019*, 2022. 2
- [61] Aron Yu and Kristen Grauman. Fine-grained visual comparisons with local learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 192–199, 2014. 2
- [62] 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 *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9127–9134, 2019. 2
- [63] Abhay Zala, Jaemin Cho, Satwik Kottur, Xilun Chen, Barlas Oguz, Yashar Mehdad, and Mohit Bansal. Hierarchical video-moment retrieval and step-captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23056–23065, 2023. 2
- [64] Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*, 2023. 2, 4
- [65] Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, et al. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792*, 2023. 2, 4
- [66] Yue Zhao, Ishan Misra, Philipp Krähenbühl, and Rohit Girdhar. Learning video representations from large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6586–6597, 2023. 2
- [67] Yaoyao Zhong, Junbin Xiao, Wei Ji, Yicong Li, Weihong Deng, and Tat-Seng Chua. Video question answering: Datasets, algorithms and challenges. *arXiv preprint arXiv:2203.01225*, 2022. 3
- [68] Luowei Zhou, Chenliang Xu, and Jason Corso. Towards automatic learning of procedures from web instructional videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018. 1, 2
- [69] Dimitri Zhukov, Jean-Baptiste Alayrac, Ramazan Gokberk Cinbis, David Fouhey, Ivan Laptev, and Josef Sivic. Cross-task weakly supervised learning from instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3537–3545, 2019. 1, 2
- [70] Yang Zou, Jongheon Jeong, Latha Pemula, Dongqing Zhang, and Onkar Dabeer. Spot-the-difference self-supervised pre-training for anomaly detection and segmentation. In *European Conference on Computer Vision*, pages 392–408. Springer, 2022. 1, 2