

SOK-Bench: A Situated Video Reasoning Benchmark with Aligned Open-World Knowledge

Andong Wang^{*1}, Bo Wu^{*2}, Sunli Chen³, Zhenfang Chen², Haotian Guan¹
Wei-Ning Lee¹, Li Erran Li⁴, Chuang Gan^{5,2}

¹The University of Hong Kong, ²MIT-IBM Watson AI Lab,

³Tsinghua University, ⁴AWS AI, ⁵UMass Amherst

Abstract

Learning commonsense reasoning from visual contexts and scenes in real-world is a crucial step toward advanced artificial intelligence. However, existing video reasoning benchmarks are still inadequate since they were mainly designed for factual or situated reasoning and rarely involve broader knowledge in the real world. Our work aims to delve deeper into reasoning evaluations, specifically within dynamic, open-world, and structured context knowledge. We propose a new benchmark (SOK-Bench), consisting of 44K questions and 10K situations with instance-level annotations depicted in the videos. The reasoning process is required to understand and apply situated knowledge and general knowledge for problem-solving. To create such a dataset, we propose an automatic and scalable generation method to generate question-answer pairs, knowledge graphs, and rationales by instructing the combinations of LLMs and MLLMs. Concretely, we first extract observable situated entities, relations, and processes from videos for situated knowledge and then extend to open-world knowledge beyond the visible content. The task generation is facilitated through multiple dialogues as iterations and subsequently corrected and refined by our designed self-prompts and demonstrations. With a corpus of both explicit situated facts and implicit commonsense, we generate associated question-answer pairs and reasoning processes, finally followed by manual reviews for quality assurance. We evaluated recent mainstream large vision-language models on the benchmark and found several insightful conclusions. For more information, please refer to our benchmark at www.bobbywu.com/SOKBench.

1. Introduction

The evolutions of Large Language Models (LLMs) [32, 33, 39, 44] and Multimodal Large Language Models

(MLLMs) [21–23, 26, 42, 58, 62] mark a significant milestone in artificial intelligence. These models are making remarkable strides across various domains, exhibiting better capabilities in perceptual, generative, and comprehensive tasks [3, 60]. As models like ChatGPT [32, 33] and successors continue to improve in scale, the possibility to capture commonsense knowledge has also seen obvious advancements. Hence, building models with real-world knowledge is becoming more promising now which was a long-standing but difficult challenge before [5, 43].

Besides novel models, previous efforts have been directed towards evaluating multimodal commonsense reasoning, transitioning from purely language-based models to those incorporating visual comprehension, such as Visual Question Answering (VQA) [1, 15, 19, 28] or Visual Commonsense Reasoning (VCR) [10, 14, 17, 18, 47, 52]. These tasks have an underlying assumption that the given visual and language inputs (e.g., images, videos, questions, etc.) contain most of the required information. However, relying on the given task context does not fully reflect real-world complexities, narrowing the scope of their reasoning evaluation. More recently, the focus has shifted towards developing tasks that require external knowledge for commonsense reasoning [27, 41, 56] and consider the integration of factual [46], evidential [20], or external knowledge sources [11, 27, 41, 48, 56] into evaluations. They primarily provide visual information and external information on static images for reasoning. Thus, this way overlooks several crucial natures such as spatio-temporal, causal, or dynamic processes of real-world activities and events, an aspect crucial for truly understanding and interacting with our environment. Also, the knowledge is written and concluded by free-form descriptions from crowd-source ways [11, 27, 41, 48, 56], presenting its own set of limitations. Consistency in descriptions provided by annotators is difficult to maintain [27, 56], and there are multiple risks of the method being influenced by distinct descriptive or empirical biases [10, 14, 17, 47].

Our research aims to delve deeper into the realm of com-

^{*}The authors contributed equally to the work.




Figure 1. Overview. Instead of using crowd-sourced methods, we design a synthesis pipeline to create the benchmark by leveraging use LLMs and VLLMs, improving efficiency and ensuring consistency. The method helps to automatically generate high-quality question-answers (QAs) and focusing the desirable purposes for evaluating the model’s ability. To generate data aligned with open-world knowledge, we propose to connect situation, general knowledge, and situated commonsense and produced three types of associated knowledge graphs (refer to the subsection 3.2, 3.3, and 3.4). Specifically, it makes more precise inferences based on situational facts and essential commonsense knowledge by aligning with the bottom-up or top-down goals, the reasoning process from Q to A is able to demonstrate explicitly.

monsense reasoning, specifically within dynamic, open-world, and structured contexts. We build Situated Open-World Commonsense Reasoning (SOK-Bench), a novel benchmark consists of over 44K questions with answers and situated commonsense knowledge graphs, covers over 12 types of questions, and sources from about 10K dynamic situations in real-world activities. The models are required to produce appropriate inferences by leveraging both facts within situations and the necessary commonsense or background knowledge.

Thus we propose the Situated Commonsense Graph to connect and integrate the required knowledge of the reasoning process where situated knowledge and general knowledge are structured, compositional, and aligned. Compared with other datasets in Figure 1, our reasoning benchmark has diverse characteristics including instance annotations, compositional generation, and structural alignment of the situated open-world knowledge and rationales. Moreover, we propose a scalable and automatic pipeline for both task generation and annotation. The typical way to generate evaluation task instructions is crowd-source human annotation [1, 56] or adversarial filtering [55]. However, collecting and locating high-quality situated commonsense data are not trivial since considering situation, question, answer and their inner relations requires experienced annotators, huge cost, and high consistency. Instead, we generate evaluation tasks, questions with answers, and reasoning processes, which integrate situated knowledge, general knowledge, and underlying reasoning commonsense goals or logics. We design appropriate instructions and compositions for prompts and define question types, goals to enhance

generation quality, logic, and rationality, as shown in Figure 3 (a). The iterative generation process automatically utilizes the interactions of LLMs or MLLMs in multi-turns, as shown in Figure 4. In experiments, we evaluate the mainstream LLMs and VideoLLMs as representatives for the proposed reasoning benchmark. Although the models have trained on web-scale data or adopted pre-trained foundation models as basis, the experiments indicate significant room for future improvement. We also provide comprehensive ablation comparisons and analysis for generation processes.

- We create SOK-Bench, a novel benchmark to evaluate situated and open-world commonsense reasoning in videos with compositionality, temporality, and causality.
- We propose a scalable method to generate video question-answering pairs for reasoning through iterative conversations with LLMs and MLLMs. We designed diverse prompt compositions and a self-prompting strategy to construct the video knowledge, commonsense knowledge, situated commonsense knowledge graphs, etc.
- We evaluated representative LLMs and MLLMs on the proposed benchmark in different settings and found that existing LLMs and MLLMs still perform inferior on situated open-knowledge reasoning in videos, which further proves our new dataset’s research value.

2. Related Works

2.1. Video Question Answering

Video Question Answering. Our work is closely related to video question answering, which requires a model to watch

Table 1. Comparison between the proposed benchmark with existing datasets.

	Approach	Generation		QA	Knowledge	Situating Knowledge		General	Situating Commonsense Knowledge	
		Real-world	Scalability	Rationale	Representation	Grounding	Dynamic	Knowledge	Existence	Compositionality
CLEVR [17]	Rule	✗	✓	✓	program	✓	✓	single-source	✓	✓
CLEVRER [52]	Rule	✗	✓	✓	program	✓	✓	single-source	✓	✓
EgoTV [12]	Rule	✗	✓	✓	program	✓	✓	single-source	✗	✗
STAR [47]	Rule	✓	✓	✓	graph	✓	✓	single-source	✗	✗
AGQA [10]	Rule	✓	✓	✓	graph	✓	✓	single-source	✗	✗
VCR [57]	Manual	✓	✗	✓	description	✓	✓	multi-source	✓	✗
VIOLIN [24]	Manual	✓	✗	✗	description	✓	✓	single-source	✓	✗
NExT-QA [48]	Manual	✓	✗	✗	description	✗	✓	single-source	✓	✗
Causal-VidQA [20]	Manual	✓	✗	✗	description	✗	✓	single-source	✓	✗
V2C [9]	Manual	✓	✗	✓	graph	✗	✓	multi-source	✓	✗
OK-VQA [27]	Manual	✓	✗	✗	description	✗	✗	multi-source	✓	✗
A-OKVQA [41]	Manual	✓	✗	✓	description	✗	✗	open-world	✓	✗
Ours	Auto, Rule	✓	✓	✓	graph	✓	✓	open-world	✓	✓

the video and answer questions related to the video’s content [7, 10, 19, 20, 47, 48, 52]. Most of these existing benchmarks ask questions about visual attributes, human actions, and activities, or physical intuition [7, 36, 52]. However, none of them has focused on studying events in videos with situated knowledge, open-domain knowledge, and explicit multi-step reasoning rationales.

Commonsense Question-Answering. Our research is aligned with the field of commonsense question answering. It requires models to make use of commonsense to answer questions. It was first studied in the field of natural question answering [29, 51] without any visual context as input. Later, people evaluated commonsense understanding in the context of visual question answering [6, 27, 35, 41]. Given a static image, such benchmarks require models to answer questions based on the visual context of the image and its associated commonsense knowledge. There are some benchmarks studying domain-specific commonsense like physics [2] and social events [54]. However, we are the first benchmark to study reasoning with situated and open-world knowledge in dynamic situations with dense annotations like knowledge graphs and rationales. Table 1 summarizes the differences compared with previous benchmarks.

LLM for Data Annotation. Our work is situated within the domain of using large language models for data annotation [8, 37]. Ding *et al.* [8] examined the effectiveness of GPT-3 as a data annotator for NLP tasks by comparing it with traditional annotation methods and analyzing its performance across various tasks. Pei *et al.* [37] introduced a self-supervised GPT annotation method using a generating-recovering paradigm, leveraging one-shot learning for efficient data-to-summary annotation, and evaluated its performance using alignment scores and human feedback reward networks. Both of these two works only focus on the natural language side. In this work, we propose to prompt LLMs to cooperate with different vision models [21, 62] to annotate videos and generate question-answer pairs related to video content and open knowledge.

Vision-Language Models. Recently, there has been great interest in large pre-trained vision-language models [21, 22, 26, 42, 58, 62] (VLMs). Early models [21, 62] only evaluate their performance on inputs with static images and natural language. There are also some models [26] evaluating performance their performance on videos. However, these VLMs either only showed some qualitative examples [22, 42, 58] or evaluates on traditional video question answering benchmarks [49, 50, 53]. However, [49, 50, 53] are not ideal benchmarks for commonsense reasoning since they mainly focus on the perception side and require no external knowledge to answer the questions. It remains an open question as to how to evaluate VLMs’ performance in understanding situations in videos and reasoning about open-domain knowledge with commonsense.

3. Situated Commonsense Reasoning

3.1. Benchmark Overview

Statistics. Our benchmark consists of 44K questions and answers for reasoning with situated open-world knowledge, accompanied by multiple-choice options, and 10K video clips. Each question-answer (QA) pair links to a hypergraph which is generated from the Situated Knowledge Graph, General Knowledge Graph, and Situated Commonsense Knowledge Graph (refer to the below sections and the algorithms in the Supplementary). The aligned graphs effectively showcase the relations between situated knowledge and general knowledge. Figure 2 presents four QA examples involving different types situated commonsense knowledge. We also present a detailed rationale that elucidates the sequence of reasoning steps that bridge the gap from each question to its corresponding answer, offering a clearer understanding of the underlying thought process (see Figure 1). Figure 3 showed the 12 question types of our dataset. Each question is accompanied by two types of answers: a direct answer and a set of four multiple-choice options. The dual format ensures versatility in response as-

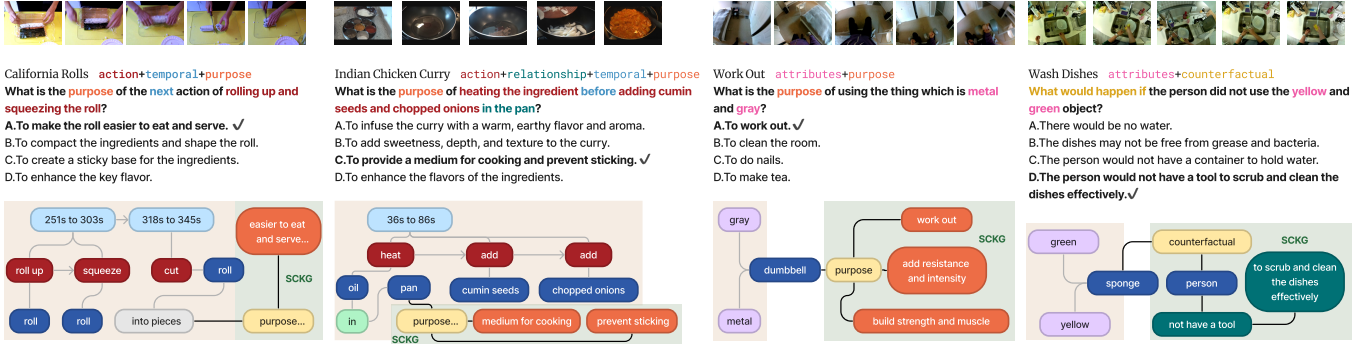


Figure 2. SOK-Bench data examples. Each QA pair corresponds to a video clip (e.g., a video clip showing how to cook *California Rolls*) and a type of situated commonsense knowledge (e.g., *action + temporal + purpose*). For each question, we provide four options, the correct choice, and the associated situated commonsense graphs.




Figure 3. (a) Sankey diagram of the 12 question types. (b) Answer distribution among options for each question type. Meaning of abbreviations: **O**: Object; **A**: Action; **CT**: Counterfactual; **CB**: Contribution; **PU**: Purpose; **I**: Inference; **PO**: Possibility; **ST**: Spatiotemporal; **GK**: General knowledge. Notably, the “Spatiotemporal” includes “obj attributes”, “obj-obj relations”, “obj attribute + obj-obj relation”, and “before/after action” (see Section 3.5).

assessment, as depicted in Figure 3.

Benchmark Generation. We present a new method for creating our benchmark automatically in a structured, controllable, and scalable way. Our approach simplifies manual and cumbersome processes and involves four streamlined stages: 1) Extracting observable content from the videos; 2) Compiling relevant commonsense knowledge exhaustively; 3) Aligning the content of the situations with this commonsense knowledge to reveal underlying logical connections and implications; 4) Formulating questions and answers by integrating the gathered information.

We utilize in-context learning, few-shot demonstrations, and multiple rounds of interactions to prompt LLMs and MLLMs for the generation. They are guided by a set of prompt templates to produce the proposed knowledge graphs. For the final stage, these graphs aid in the creation of question templates, either through manual construction or by directing the LLM to generate question-answer pairs (QAs). This strategy ensures a direct correlation between the QAs, the graphs, and the rationale that maps the reasoning from questions to answers. We’ve taken measures to ensure an unbiased distribution of answer options and have

employed both LLM and human annotators to identify and correct any grammatical errors.

3.2. Situated Knowledge Graph

We derive dynamic situations from videos, which provide insights into event temporality, causality, and dynamics [4]. Recent research has increasingly represented objects and relationships in real-world dynamic situations [10, 45, 47] and visual scenes [14, 16, 56] using graphs. Building on this, we expand upon traditional visual graphs to introduce our Situated Knowledge Graph. This includes more comprehensive information derived from dynamic situations, such as people and objects, their attributes, relationships, and actions.

Our process begins by generating object and action entities, which are identified based on descriptions provided with source videos (e.g., shown in Figure 4). These descriptions segment dynamic situations into distinct activity moments, associating objects $\{o_i\}$ and actions $\{a_i\}$ with specific timestamps $\{y_i\}$. Given original descriptions are human-written and lack a standardized format, we employ Large Language Models (LLMs) to parse the descriptive text. The LLMs systematically convert the sentences into organized lists of objects and actions, along with their corresponding timestamps. This approach streamlines the extraction of key elements from the narrative text. For some objects that are observable in video frames but were not annotated, we use BLIP2 [21] to recognize objects in frames and add them to the initial object list $\{o_i\}$. Next, we extract object attributes and spatial relationships between objects from selected keyframes within the videos. To achieve this, we utilize MiniGPT4 [32] to articulate these object attributes and their spatial interrelations. Using a designated *goal* prompt, our method efficiently generates object attributes for visual, physical, and chemical properties and relationships (e.g., #Prompt#: Please describe the tofu in the image in terms of size, shape, color, and texture in one short sentence within 10 words. #Model#: The tofus are small, white, sliced squares.). Finally, we align the times-

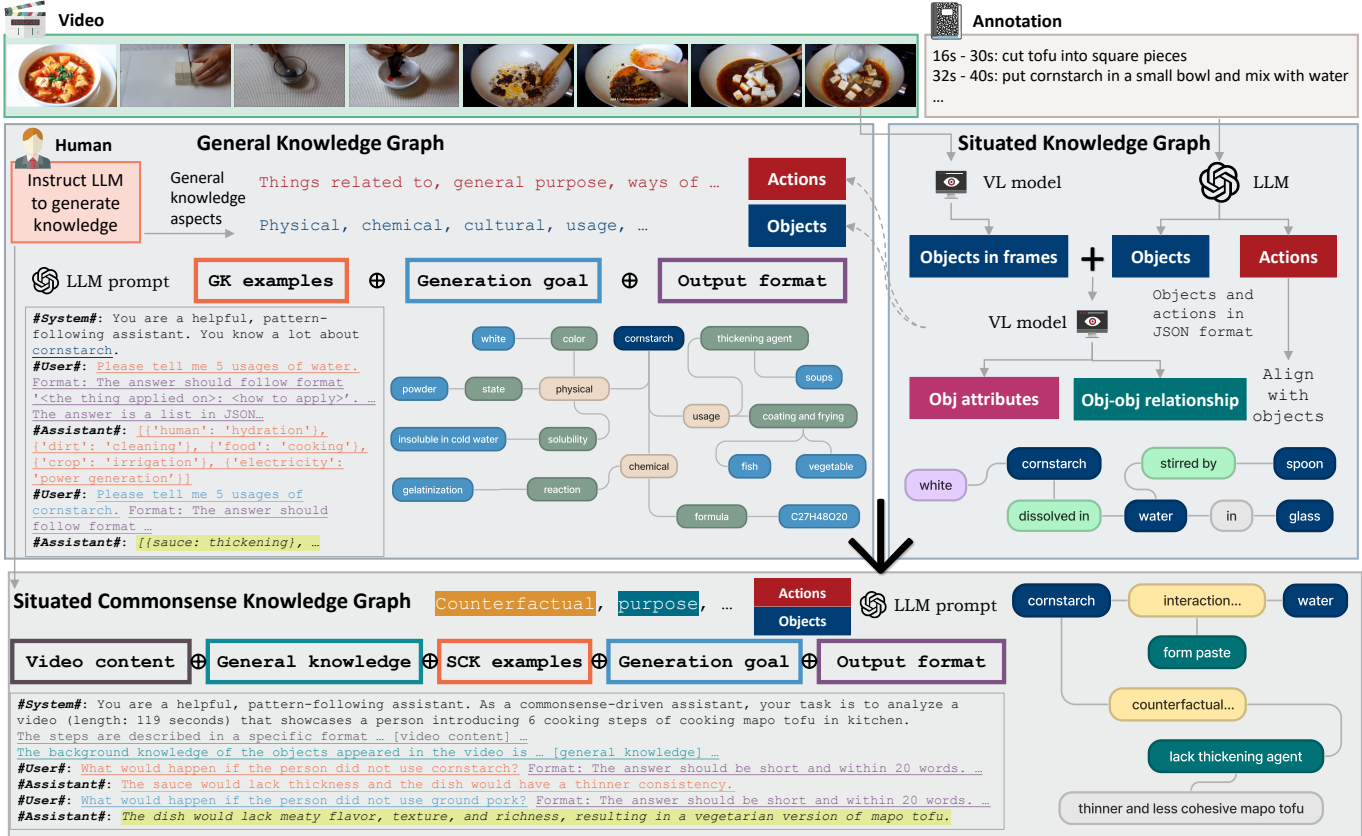


Figure 4. Generation pipeline of Situated Knowledge Graph (SKG, see Section 3.2), General Knowledge Graph (GKG, see Section 3.3), and Situated Commonsense Knowledge Graph (SCKG, see Section 3.4).

tamps of frames with those of actions. When an action’s timestamp coincides with the time at which an object appears in a frame, we link this action and object by adding an $(action, object)$ edge to the situated knowledge graph. Additionally, for each action, we build an additional edge $(timestamp, action)$ to indicate the action sequence.

3.3. General Knowledge Graph

We construct the general knowledge graph beginning from the object o_i or action a_t as initial nodes. Taking “tofu” as a node example (a food product prepared by coagulating soy milk and then pressing the resulting curds into solid white soft tofu blocks), prompting an LLM with a goal query (e.g., “Tell me some common knowledge about tofu”) is a basic approach. While this will generate descriptions with relevant knowledge, the resulting content lacks specificity and completeness, and always misses key points (e.g., the color of tofu). Thus, we propose multiple aspects for knowledge generation (e.g., *usage*, *physical*, and *culture*, etc) and collaborate with “goal” to make it with concrete descriptions. The direct challenges of LLM outputs are verbose, unstructured, or repetitive (e.g., “Tofu is a popular food ...” or “Tofu is a versatile ingredient that ...”). To address

this, we use *general knowledge examples* with the predefined *output format* to extend the knowledge graph edges or nodes. Moreover, relying solely on the format prompt can lead to format inconsistencies in outputs, such as missing quotation marks (e.g., $\{“tofu”: \dots\{“physical”: “color”: “white”\}$) that cause parsing failures. Thus, we utilize a prompt triplet $\langle general\ knowledge\ examples\ \pi, generation\ goal\ \omega, output\ format \rangle$ for GKGs (Figure 4). By employing both *general knowledge examples* and *output format*, we ensure the generated knowledge is not only of high quality but also consistently parseable. Here are the generated outputs for cornstarch usages: $\{“sauce”: “thickening”\}$, $\{“batter”: “making\ crispy”\}$, $\{“body\ powder”: “absorbing\ moisture”\}$, $\{“starch”: “making”\}$, $\{“playdough”: “creating\ homemade”\}$.

3.4. Situated Commonsense Knowledge Graph

We propose Situated Commonsense Knowledge Graph (SCKG) by integrating information from the previously established SKG and GKG for specific dynamic situations. To illustrate the significance and complexity of generating situated commonsense knowledge, consider a straightforward example: thinking about the impact of omitting cornstarch

in the preparation of mapo tofu. Merely knowing that cornstarch is a powdery substance (from general knowledge) is insufficient. We need to combine this with its situated use from the SKG (e.g., “The person mixes cornstarch with water and adds the mixture to the soup.”) and its general properties from the GKG (e.g., “Cornstarch can be used as a thickening agent.”). The result is “The sauce would be less thick, and the dish would have a thinner consistency.”

We initially thought that a four-element prompt, $\langle \text{video content } g_v, \text{ general knowledge } k_g, \text{ generation goal } \omega, \text{ output format } \phi \rangle$, would work for generating situated commonsense knowledge. Our test indicates that the model, while generating correct responses, often provides answers that are too general. For instance, in response to a question about the impact of not using ground pork in mapo tofu, the model might simply reply, “The dish would lack flavor and texture” which, although accurate, lacks detail.

Directly providing examples to instruct the LLM to achieve a specific level of detail isn’t practical due to the dependency of knowledge on the varying video content. To address this issue, we propose a method called Few-Shot Self-Prompting.

Few-Shot Self-Prompting uses previously generated situated commonsense knowledge as examples. The prompt consists of five elements: $\langle \text{video content } g_v, \text{ general knowledge } k_g, \text{ situated commonsense knowledge examples } \pi, \text{ generation goal } \omega \rangle$. The details of the algorithm can be found in the Supplementary Materials.

Figure 4 demonstrates that the situated knowledge examples are constructed using knowledge about the consequences of not using a specific ingredient (e.g., not using *cornstarch* during cooking). Over time, this process leads to more concrete and specific situated commonsense knowledge. The final result, for example, would be, “The dish would lack meaty flavor, texture, and richness, resulting in a vegetarian version of mapo tofu.”

Our Few-Shot Self-Prompting technique enables us to explore various situated commonsense knowledge perspectives, such as considering counterfactual scenarios, understanding the purpose of an action, and recognizing an object’s contribution. Our experiments show that just two iterations of knowledge generation can produce high-quality results, ultimately leading to the creation of the SCKG.

3.5. Question and Answer Generation

Using the three graphs, we create question-answer pairs to test the model’s ability to make accurate inferences based on situational facts and essential commonsense knowledge. We can use two approaches for this: (1) Manually create question-answer templates in a bottom-up manner, designing question templates and providing answers based on the graphs. (2) Automatically generate questions using a LLM in a top-down manner.

3.5.1 Bottom-up QA Generation

We manually design question templates in multi hog way combing both “situation” and “commonsense” (please see the details of template types in Supplementary Materials). For one object / action, we find the connected edges in the three graphs. Next, we design question templates that jump from edges in Situated Knowledge Graph or General Knowledge Graph to those in Situated Commonsense Knowledge Graph. Concretely, for example, the template can be “What would happen if the person did not use the $\langle \text{obj attribute} \rangle$ and $\langle \text{obj-obj relation} \rangle$?” (see one concrete question example concerning “cornstarch” in Figure 1). The model needs to first do spatiotemporal reasoning to identify the object by $\langle \text{obj attribute} \rangle$ and $\langle \text{obj-obj relation} \rangle$ (e.g., “white substance” and “dissolved in water in glass”). Next, the model needs to do situated commonsense reasoning to understand the consequence if that object was not used. In such way, we can comprehensively assess model’s situated commonsense reasoning ability. Note that we also present some easier templates (e.g., “vanilla obj counterfactual”) which are single hog.

Apart from the question and the correct answer, consistent with VCR [57], we generate wrong answers that are relevant but sufficiently dissimilar from the correct answers to minimize shortcut risk (shown in Figure 2). Incorrect options are chosen from the same category of contextual knowledge related to other objects or actions. Also, we ensure an even distribution of the four choices (Figure 3(b)).

3.5.2 Top-down QA Generation

Bottom-up QA generation allows for control over question quality and difficulty but may not create diverse question types. To address this, we suggest a top-down approach for QA generation. We design a structured prompt with five elements: $\langle \text{video content}, \text{ integrated graph}, \text{ QA examples}, \text{ generation goal}, \text{ output format} \rangle$. The integrated graph combines SKG, GKG, and SCKG. The generation goal directs the LLM to create multi-hog questions based on multiple edges from the integrated graph. This prompt design keeps a strong connection between generated QAs and graphs, reducing the chance of model hallucination compared to basing QAs only on video content.

However, the top-down approach has drawbacks, and it’s resource-intensive and slow. With the three graphs involved, it takes a few seconds to generate one question using the top-down method, while the bottom-up method generates the entire benchmark in the same time. In this paper, we mainly use the bottom-up method to create the benchmark, but we discuss top-down generation case studies in the Supplementary.

3.6. Quality Valiation

We invite human helpers to assess the quality by inspecting a subset of the graphs and QAs. We find 93.08% of QA pairs are valid. Please see the Supplementary for details.

3.7. Data Source

We created the benchmark based on videos of the two public video datasets about daily activities.

YouCook2 [61] is a large set of instructional cooking videos from YouTube. It includes over 2,000 videos across 89 recipe categories, with each video annotated with detailed step-by-step instructions.

HOMAGE [40] contains 1,752 video clips featuring 75 daily human activities in home environments. The dataset is thoroughly annotated with action labels, object information, and spatial-temporal relationships

4. Experiments

To evaluate the quality of our dataset, we ask two questions, which are answered in the following sections:

- Do current methods effectively tackle our generated problems? Why do they perform better/worse compared with other QA datasets?
- Is our QA generation procedure efficient and justifiable?

4.1. Baseline Results

To address the first question, we select a range of recent video generation models, including Video-LLaMa [58], PandaGPT [42], Video-ChatGPT [22], AskAnything [22], and Valley [26]. For deployable models, we specifically chose their Vicuna-7B-based versions [38]. Detailed baseline settings are described in the appendix.

Baseline experiment settings. The baseline models are tested in two settings: “multiple-choice” and “direct-answer”. In the multiple-choice setting, models receive a question and four choice candidates; in the direct-answer setting, models are required to provide direct and succinct responses to the questions. To calculate the accuracy in the multiple-choice setting, we need to parse an integer as the choice taken by the model. If the parsing algorithm fails, we calculate the BLEU score between model output and each choice candidate, picking the highest one as the model’s answer. Results are shown in Table 2. We measure the BLEU score [34] and BERT-F1 score [59] between the model’s output and the generated answer in the “direct-answer” setting. Results are shown in Table 3 and 4 respectively.

Result Analysis. From the table, we have the following observations. First, all baseline models perform far from perfect. GPT4v model has the best performance across different settings, which is consistent with its superior performance on standard leaderboards like [25] and [13]. The best performing AskAnything (where ChatGPT is used as

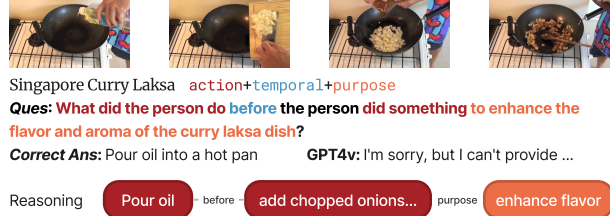


Figure 5. GPT4v’s ability to perform complex combined spatiotemporal and situated commonsense reasoning is limited. The model needs to do two-hog reasoning, *i.e.*, understanding the purpose of “adding chopped onions” is to “enhance the flavor” while knowing the previous action is “pouring oil”.

the language model) only has a BLEU score of 0.110 in table 3. It shows the value of the proposed benchmark to evaluate vision-language models’ capabilities to understand and reason in videos.

Secondly, we notice a significant performance gap between API models (the GPT family) and deployable models (the LLaMa family with Vicuna-7B base model). We remark that this phenomenon can be attributed to the difference in both sizes and instruction-following abilities between GPT LLM family and LLaMa LLM family.

Third, we find that most models perform better on simpler question types (*e.g.*, OCT and APU) compared to more complex ones (*e.g.*, OCB+ST and OCT+ST). For simpler questions, the choices might reveal the answers (*e.g.*, Q: “What did the person use to provide a sticky base to hold ingredients for California Roll?”, Choices: “Rice”, “Avocado”, “Nori”, and “Cucumber”). GPT4v may use this to achieve high multi-choice accuracy on OCB. However, GPT4v might struggle when both spatiotemporal and situated commonsense reasoning are required (see Figure 5).

4.2. Ablation Studies

We answer the second question by conducting ablation studies to evaluate the effectiveness of each component within the generation method for the three graphs.

MLLMs for SKG Generation. To justify the usage of MLLMs which provide spatiotemporal information, we focus on the four type QAs, namely “Object Counterfactual”, “Object Counterfactual + Spatiotemporal”, “Object Contribution”, and “Object Contribution + Spatiotemporal”. In the human evaluation, we ask whether the reasoning questions require to observe the video content, with results showing 91%, 99%, 94%, and 99% respectively. This confirms that MLLMs are critical to generating high-quality situated commonsense questions.

Effectiveness of Prompt Design in GKG Generation. We use LLM to generate KGs in parsable JSON format. Here, we test the effectiveness of specifying output format and providing examples in the prompt structure for GKG generation. When given only the goal of the task, 0% of the

Table 2. Accuracy of baseline models in multiple-choice setting. The question types are represented by numbers from 1 to 12: 1. Object Counterfactual + Spatiotemporal (OCT+ST); 2. Object Contribution + Spatiotemporal (OCB+ST); 3. Object Counterfactual + General Knowledge (OCT+GK); 4. Object Contribution + General Knowledge (OCB+GK); 5. Object Counterfactual (OCT); 6. Object Contribution (OCB); 7. Object Inference (OI); 8. Action Purpose + Spatiotemporal (APU+ST); 9. Action Counterfactual + Spatiotemporal (ACT+ST); 10. Action Purpose (APU); 11. Action Counterfactual (ACT); 12. Action Possibility (APO).

Models	Questions concerning objects							Questions concerning actions					Overall
	1	2	3	4	5	6	7	8	9	10	11	12	
blind ChatGPT [30]	0.224	0.318	0.441	0.275	0.651	0.488	0.206	0.136	0.135	0.574	0.729	0.189	0.365
GPT4v [31]	0.000	0.600	0.600	0.333	0.800	1.000	0.400	0.300	0.222	0.778	0.885	0.600	0.539
Video-LLaMa [58]	0.251	0.265	0.254	0.258	0.301	0.273	0.319	0.267	0.257	0.255	0.257	0.242	0.264
PandaGPT [42]	0.264	0.281	0.286	0.281	0.373	0.335	0.368	0.236	0.223	0.344	0.445	0.350	0.312
Ask Anything [22]	0.256	0.362	0.441	0.404	0.600	0.543	0.276	0.288	0.180	0.713	0.729	0.65	0.455
Video-ChatGPT [22]	0.309	0.206	0.274	0.210	0.394	0.311	0.274	0.277	0.186	0.358	0.458	0.409	0.312
Valley [26]	0.281	0.265	0.322	0.276	0.288	0.290	0.343	0.244	0.202	0.368	0.487	0.353	0.311

Table 3. BLEU score of baseline models in direct-answer setting.

Models	Questions concerning objects							Questions concerning actions					Overall
	1	2	3	4	5	6	7	8	9	10	11	12	
blind ChatGPT [30]	0.032	0.053	0.062	0.038	0.062	0.009	0.010	0.035	0.133	0.046	0.105	0.036	0.052
GPT4v [31]	0.142	0.150	0.078	0.072	0.198	0.015	0.018	0.018	0.124	0.000	0.134	0.092	0.085
Video-LLaMa [58]	0.050	0.051	0.047	0.044	0.051	0.002	0.006	0.026	0.056	0.017	0.058	0.028	0.036
PandaGPT [42]	0.089	0.096	0.085	0.082	0.097	0.008	0.014	0.054	0.103	0.035	0.123	0.052	0.070
Ask Anything [22]	0.141	0.144	0.110	0.111	0.174	0.013	0.030	0.083	0.127	0.194	0.202	0.077	0.110
Video-ChatGPT [22]	0.142	0.125	0.108	0.090	0.152	0.014	0.044	0.079	0.143	0.109	0.164	0.067	0.096
Valley [26]	0.099	0.096	0.095	0.084	0.107	0.010	0.012	0.070	0.112	0.035	0.134	0.048	0.068

Table 4. BERT score of baseline models in direct-answer setting.

Models	Questions concerning objects							Questions concerning actions					Overall
	1	2	3	4	5	6	7	8	9	10	11	12	
blind ChatGPT [30]	0.887	0.878	0.892	0.881	0.884	0.896	0.887	0.886	0.883	0.878	0.892	0.886	0.886
GPT4v [31]	0.954	0.953	0.955	0.959	0.956	0.955	0.953	0.963	0.954	0.957	0.952	0.957	0.956
Video-LLaMa [58]	0.961	0.960	0.961	0.961	0.961	0.961	0.961	0.961	0.961	0.961	0.961	0.961	0.961
PandaGPT [42]	0.952	0.953	0.952	0.952	0.953	0.952	0.953	0.952	0.952	0.952	0.952	0.952	0.952
Ask Anything [22]	0.957	0.957	0.958	0.956	0.961	0.960	0.955	0.962	0.959	0.964	0.966	0.959	0.959
Video-ChatGPT [22]	0.955	0.955	0.957	0.956	0.957	0.963	0.952	0.956	0.958	0.958	0.961	0.956	0.957
Valley [26]	0.955	0.955	0.955	0.955	0.954	0.954	0.955	0.955	0.955	0.955	0.954	0.955	0.955

outputs are valid; specifying output format, this figure increases to 11.1%; with both output format and examples, all outputs can be parsed as JSON files.

Effectiveness of Few-Shot Self-Prompting. We argue the few-shot examples of SCKG generation are necessary. By human testing, we evaluate the proportion of concrete and specific knowledge relations as opposed to un-situated, general knowledge. With Few-Shot Self-Prompting, 97% of the generated knowledge is concrete while the ratio decreases to 61% without the prompt technique.

5. Conclusion

Our research introduces a novel benchmark for Situated Open-World Commonsense Reasoning, advancing AI’s

ability to comprehend and reason in dynamic, real-world contexts. This benchmark includes a wide range of questions and situational analyses that go beyond traditional reasoning paradigms, challenging existing AI systems. Our novel approach in dataset generation offers scalability and enhanced logic in QA pair creation. While current models show promise, our experiment findings highlight the need for significant improvements, pointing towards exciting avenues for future general artificial intelligence.

6. Acknowledgements

We appreciate the contributions of our data quality validation team, Jiajing Zhang, Bingze Dai, Renxian Wang, Yue Xu, Yucong Li, Tuo Zhou, Wei Yi Oon, and Zijian Li.

References

- [1] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *ICCV*, 2015.
- [2] Tayfun Ates, M Samil Atesoglu, Cagatay Yigit, Ilker Kesen, Mert Kobas, Erkut Erdem, Aykut Erdem, Tilbe Goksun, and Deniz Yuret. Craft: A benchmark for causal reasoning about forces and interactions. *arXiv*, 2020.
- [3] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023.
- [4] Shyamal Buch, Cristóbal Eyzaguirre, Adrien Gaidon, Jiajun Wu, Li Fei-Fei, and Juan Carlos Niebles. Revisiting the “video” in video-language understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2917–2927, 2022.
- [5] Aman Chadha and Vinija Jain. ireason: Multimodal commonsense reasoning using videos and natural language with interpretability. *arXiv preprint arXiv:2107.10300*, 2021.
- [6] Yingshan Chang, Mridu Narang, Hisami Suzuki, Guihong Cao, Jianfeng Gao, and Yonatan Bisk. Webqa: Multihop and multimodal qa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16495–16504, 2022.
- [7] Zhenfang Chen, Kexin Yi, Antonio Torralba, Josh Tenenbaum, and Chuang Gan. Comphy: Compositional physical reasoning of objects and events from videos. In *International Conference on Learning Representations*, 2022.
- [8] Bosheng Ding, Chengwei Qin, Linlin Liu, Lidong Bing, Shafiq Joty, and Boyang Li. Is gpt-3 a good data annotator? *arXiv*, 2022.
- [9] Zhiyuan Fang, Tejas Gokhale, Pratyay Banerjee, Chitta Baral, and Yezhou Yang. Video2Commonsense: Generating commonsense descriptions to enrich video captioning. In *EMNLP*, 2020.
- [10] Madeleine Grunde-McLaughlin, Ranjay Krishna, and Maneesh Agrawala. Agqa: A benchmark for compositional spatio-temporal reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11287–11297, 2021.
- [11] Pranay Gupta and Manish Gupta. Newsvqa: Knowledge-aware news video question answering. In *Pacific-asia conference on knowledge discovery and data mining*, pages 3–15. Springer, 2022.
- [12] Rishi Hazra, Brian Chen, Akshara Rai, Nitin Kamra, and Ruta Desai. Egotv: Egocentric task verification from natural language task descriptions. *arXiv*, 2023.
- [13] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *ICLR*, 2021.
- [14] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *CVPR*, 2019.
- [15] Yunseok Jang, Yale Song, Youngjae Yu, Youngjin Kim, and Gunhee Kim. Tgif-qa: Toward spatio-temporal reasoning in visual question answering. In *CVPR*, 2017.
- [16] Jingwei Ji, Ranjay Krishna, Li Fei-Fei, and Juan Carlos Niebles. Action genome: Actions as compositions of spatio-temporal scene graphs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10236–10247, 2020.
- [17] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*, 2017.
- [18] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Judy Hoffman, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Inferring and executing programs for visual reasoning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2989–2998, 2017.
- [19] Jie Lei, Licheng Yu, Mohit Bansal, and Tamara L Berg. Tvqa: Localized, compositional video question answering. In *EMNLP*, 2018.
- [20] Jiangtong Li, Li Niu, and Liqing Zhang. From representation to reasoning: Towards both evidence and commonsense reasoning for video question-answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [21] 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.
- [22] KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, 2023.
- [23] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023.
- [24] Jingzhou Liu, Wenhui Chen, Yu Cheng, Zhe Gan, Licheng Yu, Yiming Yang, and Jingjing Liu. Violin: A

- large-scale dataset for video-and-language inference. In *CVPR*, 2020.
- [25] Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating math reasoning in visual contexts with gpt-4v, bard, and other large multimodal models. *arXiv*, 2023.
- [26] Ruipu Luo, Ziwang Zhao, Min Yang, Junwei Dong, Minghui Qiu, Pengcheng Lu, Tao Wang, and Zhongyu Wei. Valley: Video assistant with large language model enhanced ability. *arXiv preprint arXiv:2306.07207*, 2023.
- [27] Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition*, pages 3195–3204, 2019.
- [28] Jonghwan Mun, Paul Hongsuck Seo, Ilchae Jung, and Bohyung Han. Marioqa: Answering questions by watching gameplay videos. In *ICCV*, 2017.
- [29] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. Ms marco: A human-generated machine reading comprehension dataset. 2016.
- [30] OpenAI. ChatGPT: Optimizing Language Models for Dialogue. <https://openai.com/chatgpt>, 2023. Version 4.0.
- [31] OpenAI. Gpt-4 technical report, 2023.
- [32] R OpenAI. Gpt-4 technical report. *arXiv*, pages 2303–08774, 2023.
- [33] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- [34] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *ACL*, 2002.
- [35] Jae Sung Park, Chandra Bhagavatula, Roozbeh Mottaghi, Ali Farhadi, and Yejin Choi. Visualcomet: Reasoning about the dynamic context of a still image. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pages 508–524. Springer, 2020.
- [36] Maitreya Patel, Tejas Gokhale, Chitta Baral, and Yezhou Yang. Cripp-vqa: Counterfactual reasoning about implicit physical properties via video question answering. *arXiv preprint arXiv:2211.03779*, 2022.
- [37] Xiaohuan Pei, Yanxi Li, and Chang Xu. Gpt self-supervision for a better data annotator. *arXiv*, 2023.
- [38] Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. Instruction tuning with gpt-4, 2023.
- [39] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [40] Nishant Rai, Haofeng Chen, Jingwei Ji, Rishi Desai, Kazuki Kozuka, Shun Ishizaka, Ehsan Adeli, and Juan Carlos Niebles. Home action genome: Cooperative compositional action understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11184–11193, 2021.
- [41] Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. A-okvqa: A benchmark for visual question answering using world knowledge. In *European Conference on Computer Vision*, pages 146–162. Springer, 2022.
- [42] Yixuan Su, Tian Lan, Huayang Li, Jialu Xu, Yan Wang, and Deng Cai. Pandagpt: One model to instruction-follow them all. *arXiv preprint arXiv:2305.16355*, 2023.
- [43] Niket Tandon, Aparna S Varde, and Gerard de Melo. Commonsense knowledge in machine intelligence. *ACM SIGMOD Record*, 46(4):49–52, 2018.
- [44] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [45] Aisha Urooj, Hilde Kuehne, Bo Wu, Kim Chheu, Walid Bouselham, Chuang Gan, Niels Lobo, and Mubarak Shah. Learning situation hyper-graphs for video question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14879–14889, 2023.
- [46] Peng Wang, Qi Wu, Chunhua Shen, Anthony Dick, and Anton Van Den Hengel. Fvqa: Fact-based visual question answering. *IEEE transactions on pattern analysis and machine intelligence*, 40(10):2413–2427, 2017.
- [47] Bo Wu, Shoubin Yu, Zhenfang Chen, Joshua B Tenenbaum, and Chuang Gan. Star: A benchmark for situated reasoning in real-world videos. In *NeurIPS*, 2021.
- [48] Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. Next-qa: Next phase of question-answering to explaining temporal actions. In *CVPR*, 2021.
- [49] Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. Video question answering via gradually refined attention over appearance and motion. In *ACM Multimedia*, 2017.

- [50] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*, pages 5288–5296, 2016.
- [51] Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov, and Christopher D Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. *arXiv*, 2018.
- [52] Kexin Yi, Chuang Gan, Yunzhu Li, Pushmeet Kohli, Jiajun Wu, Antonio Torralba, and Joshua B Tenenbaum. Clevrer: Collision events for video representation and reasoning. In *International Conference on Learning Representations*, 2020.
- [53] 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.
- [54] 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.
- [55] Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. Swag: A large-scale adversarial dataset for grounded commonsense inference. *arXiv preprint arXiv:1808.05326*, 2018.
- [56] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6720–6731, 2019.
- [57] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reasoning. In *CVPR*, 2019.
- [58] 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.
- [59] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert, 2020.
- [60] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging llm-as-a-judge with mt-bench and chatbot arena, 2023.
- [61] 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.
- [62] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023.