

REVEALING THE ILLUSION OF JOINT MULTIMODAL UNDERSTANDING IN VIDEOQA MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

While VideoQA Transformer models demonstrate competitive performance on standard benchmarks, the reasons behind their success are not fully understood. Do these models jointly capture and leverage the rich multimodal structures and dynamics from video and text? Or are they merely exploiting shortcuts to achieve high scores? Hence, we design *QUAG* (QUadrant AveraGe), a lightweight and non-parametric probe, to critically analyze multimodal representations. *QUAG* facilitates combined dataset-model study by systematic ablation of model’s coupled multimodal understanding during inference. Surprisingly, it demonstrates that the models manage to maintain high performance even under multimodal impairment. We extend *QUAG* to design “*QUAG*-attention”, a simplistic and less-expressive replacement of self-attention. We find that the models with *QUAG*-attention achieve similar performance with significantly less mulops without any finetuning. These findings indicate that the current VideoQA benchmarks and metrics do not penalize models that find shortcuts and discount joint multimodal understanding. Motivated by this, we propose *CLAVI* (Counterfactual in LAnguage and VIdeo), a diagnostic dataset for coupled multimodal understanding in VideoQA. *CLAVI* consists of temporal questions and videos that are augmented to curate balanced counterfactuals in language and video domains. We evaluate models on *CLAVI* and find that all models achieve high performance on multimodal shortcut instances, but most of them have very poor performance on the counterfactual instances that necessitate joint multimodal understanding. Overall, with the multimodal representation analysis using *QUAG* and diagnostic analysis using *CLAVI*, we show that many VideoQA models are incapable of learning multimodal representations and that their success on standard datasets is an illusion of joint multimodal understanding.

1 INTRODUCTION

Multimodal learning with videos and language is challenging, despite the shared sequential nature of these modalities, due to their distinct underlying structures. That is, videos exhibit spatio-temporal dynamics in the pixel space, whereas language representation is composed of the syntax and semantics of word sequences. Hence, tasks like Video Question Answering (VideoQA) (Zhong et al., 2022) are difficult as they necessitate the model to acquire accurate representations of both the modalities and establish meaningful connections between them. Transformers have demonstrated exceptional performance on VideoQA benchmarks (Zhong et al., 2022). However, since they lack the intrinsic inductive biases for these representation, they must learn it from the data (Xu et al., 2021; Patrick et al., 2021). But does the good performance of Transformers on current VideoQA benchmarks necessarily mean that they learn to faithfully represent, leverage, understand and reason the modalities? Or do the current benchmarks and metrics fail to robustly evaluate the models for their multimodal understanding?

This is a valid concern because deep learning models can learn shortcuts to achieve good performance without faithfully representing underlying modalities (Geirhos et al., 2020). For example, seemingly spatio-temporal tasks, like some action classification problems, are shown to be solved without focusing much on temporal representations (Kowal et al., 2022; Sevilla-Lara et al., 2021). Similarly, in VideoQA, recent works report that the datasets contain specific biases (Buch et al.,

2022; Lei et al., 2023). However, these works are restricted to isolated analyses of either the models or the datasets. This raises questions: **Are the models actually learning to jointly leverage and understand the modalities, or is the performance on the current benchmarks an illusion of joint multimodal learning?**

To answer these questions, we propose QUadrant AveraGe (QUAG), a lightweight and non-parametric probe to systematically gauge the reliance of a finetuned model’s performance on joint multimodal representations. We posit that joint multimodal understanding is enabled in the fusion layers by progressively attending to the informative tokens within and between the modalities. QUAG impairs the components of modality fusion by block-averaging attention weights. We apply QUAG on multiple dataset-model combinations, and consistently find that the models manage to achieve high performance on the benchmarks without relying specific multimodal interactions.

This finding is concerning because high performance on established benchmarks should be ideally indicative of coupled multimodal understanding. We establish that these models learn sub-optimal representations; that is, the modality fusion doesn’t effectively capture the information within each modality along with the complementary information in the other modality. We validate the sub-optimality in multimodal representations by replacing self-attention in the pretrained models with simple and less-expressive QUAG-attention. Even though QUAG-attention impairs multimodal capabilities, the models augmented with QUAG-attention manage to maintain the high performance on standard benchmarks without any finetuning. This raises a follow-up question – **How then can we diagnose coupled multimodal understanding?**

Thus, we create Counterfactual in LAnguage and VIsion (CLAVI), a diagnostic benchmark to robustly assess joint multimodal understanding in VideoQA models. Temporal understanding ideally requires coupled multimodal understanding. However, the standard benchmarks do not contain or assess performance on counterfactual instances. CLAVI contains automatically generated balanced temporal counterfactuals in both question and video domains to accurately test if the models can jointly understand temporal cues in the question (temporal prepositions and adverbs) and the video (order of frames) domains (Figure 2). We develop consistent-accuracy metrics to precisely assess the contributions of shortcuts to circumvent joint multimodal understanding. We find that finetuned models have high-accuracy on shortcut instances in CLAVI, but have poor performance on the counterfactual instances that require coupled multimodal understanding. Hence, we position CLAVI as a litmus test to diagnose joint multimodal understanding which is overlooked by the existing datasets.

In summary, our contributions are (i) we develop QUAG, a systematic method to identify sub-optimality in joint multimodal representations, (ii) using QUAG and QUAG-attention, we demonstrate that high performance on established VideoQA benchmarks is not representative of faithful coupled multimodal understanding, and (iii) we develop CLAVI, a new diagnostic benchmark that contains balanced temporal counterfactuals in videos and questions to confidently disambiguate the contributions of shortcuts in joint multimodal learning to benchmark the models. Overall, QUAG and CLAVI provide holistic dataset-model insights that reveal the illusion of multimodal understanding in VideoQA models.

2 DO VIDEOQA MODELS LEARN TO JOINTLY LEVERAGE THE MODALITIES?

We posit that coupled multimodal understanding is enabled in the fusion layers by progressively attending to the informative tokens within and between the modalities. Hence, we design QUAG to systematically ablate the effects of multimodal attention. It impairs the joint multimodal representations in the pretrained model by systematically block-averaging the attention weights to attend to all tokens uniformly at inference time. Based on the targeted modality-interactions, we define special cases of QUAG, collectively called short-circuit operations and analyze the performance drop.

2.1 VIDEO QUESTION ANSWERING SETUP

In VideoQA, the task is to predict the correct answer given a video-question tuple, $(\mathcal{V}, \mathcal{T})$. A VideoQA model consists of a vision encoder $F_{\mathcal{V}} : \mathcal{V} \rightarrow \mathbb{R}^{l_{\mathcal{V}} \times d}$, text encoder $F_{\mathcal{T}} : \mathcal{T} \rightarrow \mathbb{R}^{l_{\mathcal{T}} \times d}$, and a multimodal fusion module $M : (F_{\mathcal{V}}(\mathcal{V}), F_{\mathcal{T}}(\mathcal{T})) \rightarrow \mathbb{R}^{(l_{\mathcal{V}}+l_{\mathcal{T}}) \times d}$, where $l_{\mathcal{V}}$ and $l_{\mathcal{T}}$ are the maximum input sequence lengths of video and text modalities respectively and d is the dimensionality of the fusion model.

Consider M as a composition of n attention-based multimodal fusion blocks, $M = M_n \circ M_{n-1} \circ \dots \circ M_1$. Each fusion block consists of attention, normalization, and token-mixing modules. For our analysis, we consider M to be composed of self-attention transformer blocks. That is, query, key, and value are the transformations of the same input sequence. $\mathbf{X}_{\mathcal{V}\mathcal{T}} = [\mathbf{F}_{\mathcal{V}}(\mathcal{V}) \parallel \mathbf{F}_{\mathcal{T}}(\mathcal{T})] \in \mathbb{R}^{(l_{\mathcal{V}}+l_{\mathcal{T}}) \times d}$ is the input for M , where \parallel is concatenation operator. Since QUAG operates at inference time, we assume the VideoQA model to be finetuned and frozen.

2.2 QUAG: ABLATION OF MODALITY INTERACTIONS

Shortcuts are the spurious features learned by a given model on a given dataset (Murali et al., 2023). Along this axis, we use QUAG to pinpoint the exact failure modes in the dataset representations learned by the models.

Let \mathbf{X}_{i-1} denote the input of the fusion block M_i and let $(\mathbf{Q}_i, \mathbf{K}_i, \mathbf{V}_i)$ be its query, key, and value transformations and $\mathbf{X}_0 = \mathbf{X}_{\mathcal{V}\mathcal{T}}$. Then, the token-mixing operation is given by $\mathbf{T}_i = \mathbf{A}_i \mathbf{V}_i$, where $\mathbf{A}_i = \text{softmax}(\mathbf{Q}_i \mathbf{K}_i^\top)$ is the attention matrix (we omit the scaling factor \sqrt{d} for readability). For \mathbf{Q}_{1u} , \mathbf{K}_{1u} , and \mathbf{V}_{1u} to denote the query, key, and value projections of modality u for the first fusion block, M_1 , we can simplify, \mathbf{A}_1 and \mathbf{T}_1 in terms of their partition blocks, referred to as quadrants henceforth, as:

$$\mathbf{A}_1 = \text{softmax} \left(\left[\begin{array}{c|c} \mathbf{Q}_{1\mathcal{V}} & \mathbf{K}_{1\mathcal{V}}^\top \\ \hline \mathbf{Q}_{1\mathcal{T}} & \mathbf{K}_{1\mathcal{T}}^\top \end{array} \right] \right) \quad \text{and} \quad \mathbf{T}_1 = \left[\begin{array}{c|c} \mathbf{A}_{\mathcal{V}\mathcal{V}}^1 & \mathbf{A}_{\mathcal{V}\mathcal{T}}^1 \\ \hline \mathbf{A}_{\mathcal{T}\mathcal{V}}^1 & \mathbf{A}_{\mathcal{T}\mathcal{T}}^1 \end{array} \right] \left[\begin{array}{c} \mathbf{V}_{1\mathcal{V}} \\ \hline \mathbf{V}_{1\mathcal{T}} \end{array} \right]$$

where $\mathbf{A}_{u_1 u_2}^1$ represents the quadrant of \mathbf{A}_1 corresponding to $(\mathbf{Q}_{1u_1} \mathbf{K}_{1u_2}^\top)$. Note that we skip layer normalization layers in the discussion for simplicity. Hence, we can simplify and write \mathbf{T}_1 as:

$$\mathbf{T}_1 = \left[\begin{array}{c} \mathbf{A}_{\mathcal{V}\mathcal{V}}^1 \mathbf{V}_{1\mathcal{V}} + \mathbf{A}_{\mathcal{V}\mathcal{T}}^1 \mathbf{V}_{1\mathcal{T}} \\ \hline \mathbf{A}_{\mathcal{T}\mathcal{V}}^1 \mathbf{V}_{1\mathcal{V}} + \mathbf{A}_{\mathcal{T}\mathcal{T}}^1 \mathbf{V}_{1\mathcal{T}} \end{array} \right] \quad (1)$$

We follow the same partition quadrants, as defined for \mathbf{A}_1 in M_1 , for \mathbf{A}_j in the downstream fusion layer M_j and denote the quadrants as $\mathbf{A}_{u_1 u_2}^j$. Next, we define row-wise average and replace operator \mathcal{R} that operates on a quadrant of a matrix to replace the values in the quadrant with the mean value of the respective partitioned-row. Note that the values in the other quadrants are unaffected. Given a matrix \mathbf{Z} of size $p \times q$ and let W denote the location of the quadrant of \mathbf{Z} with indices $(p_1^W \dots p_2^W) \times (q_1^W \dots q_2^W)$. We use $[\cdot]_{ij}$ to index the element in row i and column j . Then,

$$[\mathcal{R}(\mathbf{Z}, W)]_{ij} = \begin{cases} \sum_{k=q_1^W}^{q_2^W} \frac{[\mathbf{Z}]_{ik}}{q_2^W - q_1^W + 1} & i \in \{p_1^W, \dots, p_2^W\} \text{ and } j \in \{q_1^W, \dots, q_2^W\} \\ [\mathbf{Z}]_{ij} & \text{otherwise} \end{cases}$$

We can now formally define the QUAG operator, ϕ , as:

$$\phi(\mathbf{A}_i, \mathbf{V}_i, [s_1, s_2, \dots, s_m]) = (\mathcal{R}_{s_1} \circ \mathcal{R}_{s_2} \dots \circ \mathcal{R}_{s_m}(\mathbf{A}_i)) \mathbf{V}_i$$

where $S = [s_1, s_2, \dots, s_m]$ is a list of quadrants such that $\forall s \in S : s \in \{\mathcal{T}\mathcal{T}, \mathcal{T}\mathcal{V}, \mathcal{V}\mathcal{T}, \mathcal{V}\mathcal{V}\}$, $\mathcal{R}_{s_i}(\mathbf{Z})$ is short-hand for $\mathcal{R}(\mathbf{Z}, s_i)$, \mathbf{A}_i and \mathbf{V}_i are the attention and value matrices of M_i respectively. Note that $\mathcal{T}\mathcal{T}$ refers to the quadrant corresponding to $\mathbf{A}_{\mathcal{T}\mathcal{T}}^i$ (independent of the index $1 \leq i \leq n$ of A), similarly $\mathcal{T}\mathcal{V}$ refers to the quadrant corresponding to $\mathbf{A}_{\mathcal{T}\mathcal{V}}^i$, and so on. In implementation, we re-adjust the quadrant boundaries to ignore the padded elements. Refer Figure 1 for an illustrative example. Incorporating QUAG in the existing model pipeline is very easy and we provide the code in the Appendix A.2.2. Since we will be applying the QUAG operator successively on all the layers of M , for brevity, we denote $\Phi(M, S) = \forall_{1 \leq i \leq n} \phi(\mathbf{A}_i, \mathbf{V}_i, S)$. Note that ϕ , and hence, Φ is independent of the order of elements in S .

2.3 SHORT-CIRCUIT OPERATIONS

As QUAG is a generic method of probing multimodal fusion, we consider some special cases based on the value of S below. We call these operations collectively as short-circuiting operations:




Figure 1: Illustrative toy example of unimodal short-circuiting or $\phi(\mathbf{Z}, [\mathcal{TT}, \mathcal{VV}])$, where \mathbf{Z} is the input attention matrix (left-most in the figure), \mathcal{R} is the row-wise average and replace operator and hatching denotes padding. The quadrants that are operated on are highlighted in bright yellow box. Note that $l_{\mathcal{V}} = 3$ and $l_{\mathcal{T}} = 2$ for the model and the video embeddings are pre-concatenated with the question embeddings. As shown in the figure, we apply \mathcal{R} successively to replace the values in the quadrant with the respective row-wise average value. The cells are colored as per their quadrants (\mathcal{VV} : red, \mathcal{VT} : yellow, \mathcal{TV} : blue, \mathcal{TT} : green).

1) $\mathcal{S} = [\mathcal{VV}, \mathcal{TT}]$: $\phi(\mathbf{A}_1, \mathbf{V}_1, [\mathcal{VV}, \mathcal{TT}])$ is equivalent to scaling the average values of $\mathbf{V}_{1\mathcal{V}}$ and $\mathbf{V}_{1\mathcal{T}}$ in the upper and lower blocks of \mathbf{T}_1 respectively (as evident from Eqn. 1). Hence, in the upper block, video queries faithfully attend over text keys but uniformly over video keys. Likewise, text queries attend faithfully over video queries but uniformly over text queries in the lower block. We illustrate this operation in Figure 1 and call such a fusion block to be unimodal average conformable.

Having understood the trivial case, we prove by induction that $\Phi(M, [\mathcal{VV}, \mathcal{TT}])$ leads to unimodal average conformability of all the component fusion blocks in M . Consider a block $M_j \in M$ such that $j > 1$. We want to show that unimodal average conformability of first $\{M_0, M_1, \dots, M_{j-1}\}$ blocks using $\forall_{1 \leq i \leq j-1} \phi(\mathbf{A}_i, \mathbf{V}_i, [\mathcal{VV}, \mathcal{TT}])$ implies $\phi(\mathbf{A}_j, \mathbf{V}_j, [\mathcal{VV}, \mathcal{TT}])$ will make M_j unimodal average conformable. The input of M_j can be decomposed into non-linear and linear (from the residual connection that skips the feed-forward layer of M_{j-1}) projections of $\mathbf{T}_{j-1} + M_{j-2} \circ M_{j-3} \cdots \circ M_1(\mathbf{X}_{\mathcal{VT}}) + \mathbf{X}_{\mathcal{VT}}$. Hence, when $\{M_0, M_1, \dots, M_{j-1}\}$ are unimodal average conformable, $\mathbf{X}_{\mathcal{VT}}$ is the only non-conformable component. And we have shown in the trivial case that $\phi(\mathbf{A}_1, \mathbf{V}_1, [\mathcal{VV}, \mathcal{TT}])$ makes M_1 conformable, hence M_j is also unimodal average conformable under ϕ .

Ultimately, $\Phi(M, [\mathcal{VV}, \mathcal{TT}])$ bypasses the effect of video-video attention and text-text attention. We prove that unimodal token-mixing is reduced to scaling the average of the modalities. We term this as **unimodal short-circuiting**. It ablates unimodal representations to analyze their dependence on the performance of the models. Since the following cases can be proved similarly using induction, we skip the proofs for conciseness.

2) $\mathcal{S} = [\mathcal{VT}, \mathcal{TV}]$: Parallel to unimodal short-circuiting, $\phi(\mathbf{A}_1, \mathbf{V}_1, [\mathcal{VT}, \mathcal{TV}])$ is equivalent to scaling the average values of $\mathbf{V}_{1\mathcal{T}}$ and $\mathbf{V}_{1\mathcal{V}}$ in the upper and lower blocks of \mathbf{T}_1 respectively. Video and text queries faithfully attend to video and text keys respectively while crossmodal attention in video-text is reduced to uniform attention. We term this effect as **crossmodal short-circuiting**. It is complementary to unimodal short-circuiting and assesses the importance of inter-modality token-mixing. It probes if the models actually learn by fusing the information between the two modalities or is it largely driven by unimodal biases within the modalities.

3) $\mathcal{S} = [\mathcal{VV}, \mathcal{TV}]$: This is equivalent to removing the effect of individual of video keys, resulting in averaging the components of video modality in the upper and lower blocks of all \mathbf{T}_i . We call this **video short-circuiting**. Similarly, $\mathcal{S} = [\mathcal{TT}, \mathcal{VT}]$ leads to **text short-circuiting**.

2.4 QUAG-ATTENTION

Along with an assessment of multimodal understanding, QUAG enables a detailed analysis of token mixing for identifying the sub-optimality of learned representations. Sub-optimality occurs if the fusion process doesn't effectively capture the information within each modality along with the complementary information in the other modality. Hence, we use QUAG as an inspiration to propose QUAG-attention, a replacement of self-attention in fusion module that calculates similarities on-

ready short-circuited sequences. That is, QUAG-attention intrinsically induces the models to infer using sub-optimal representations.

Let us consider the case such that the performance of M under video short-circuiting operation is comparable to its performance without any perturbation. If the input of M is $\mathbf{X}_0 = [\mathbf{F}_V(\mathcal{V}) \parallel \mathbf{F}_T(\mathcal{T})]$, then during token-mixing we effectively average and scale the components in the upper-partition $([1, \dots, l_V] \times d)$ of the value matrix in all the fusion blocks. This can be efficiently approximated by replacing the entire upper block with a single row-wise average token using \mathcal{R} before projecting to key and value domains. Note that the query remains unchanged. Similar to QUAG, we perform no finetuning and only modify the calculation of self-attention.

We can generalize it to present new variants of self-attention: collectively known as QUAG-attention. QUAG-attention operates by consistently averaging the corresponding modality blocks within the input of each fusion block. The averaging process occurs prior to the transformation of the input into keys and values. Depending on the sub-optimality in representation, QUAG-attention can be applied to only text, video or both the modalities. It reduces the number of keys and values tokens from $(l_V + l_T)$ to either $(l_T + 1)$ (text-average), $(l_V + 1)$ (video-average) or 2 (text-video-average).

The number of tokens in video and text modalities are generally different. However, due to block averaging, QUAG-attention reduces the effective number of tokens of the modality in key and value domains to one. The token-length mismatch would interfere with softmax operation in attention. Hence, we scale the components of dot-product similarity scores of the averaged keys by the logarithm of the number constituting tokens (that is, the original number of tokens in the block). This is similar to proportional attention used by Bolya et al. (2023) for token-merging.

2.5 EXPERIMENTAL SETTING

Models and Datasets: We evaluate QUAG and QUAG-attention on JustAsk (Yang et al., 2021a) and FrozenBiLM (Yang et al., 2022b) models. We evaluate it on the following datasets (i) **ActivityNet-QA** (Yu et al., 2019): contains 58K open-ended questions on 5.8K sampled videos from ActivityNet (ii) **MSRVTT-QA** (Xu et al., 2017): contains 244K open-ended questions on 10K MSRVTT videos (iii) **NeXT-QA** (Xiao et al., 2021): contains 47K 5-way multiple choice questions with one-correct answer from 5.4K videos. We also report results on the **ATP-Hard** subset of NeXT-QA (Buch et al., 2022) that contains a higher concentration of temporally challenging data requiring multi-frame understanding. **Implementation Details:** All our experiments were performed on 4 NVIDIA A5000 GPUs. We use the official open-source code of the models on GitHub and modify only the self-attention modules. We use the official evaluation code and checkpoints. For NeXT-QA, we use the official dataset and finetune the models with the default parameters. More details in Appendix A.2.3.

2.6 ANALYSIS

The results are shown in Table 1. For comparison to the unperturbed model, we specify the baseline, language-only (without video input) and video-only (without text input) accuracies. The high performance in language-only setting relative to the baseline is indicative of strong unimodal bias towards language. However, these metrics do not provide any information about the exact nature and degree of the sub-optimal representations learned by the models, hence we use QUAG.

The performance of FrozenBiLM on ActivityNet-QA and MSRVTT-QA drops by over 10% (43.6% to 32.3%; 46.6% to 32.8%) with crossmodal short-circuiting, and by 40% with both unimodal (43.6% to 2.4%; 46.6% to 1.0%) and text short-circuiting (43.6% to 1.4%; 46.6% to 1.0%). Furthermore, the drop is less than 1% under video short-circuiting (43.6% to 43.1%; 46.6% to 45.7%). However, for NeXT-QA and ATP-Hard, the performance of FrozenBiLM drops to chance level (20%) under text and unimodal short-circuiting operations but hardly drops with video and text short-circuiting. Parallely, the performance of JustAsk model does not drop by more than 1% for any of the datasets under any short-circuiting operation.

This means that FrozenBiLM consistently does not rely on the core features of the video modality and has a strong reliance on text-modality. Further, for NeXT-QA and ATP-Hard, the model does not leverage any crossmodal interactions. However, for ActivityNet-QA and MSRVTT-QA, it leverages

Table 1: Short-circuit (SC) and QUAG-attention accuracies for JustAsk and FrozenBiLM models on ActivityNet-QA (A-QA), MSRVT-TQA (M-QA), NeXT-QA (N-QA) and ATP-Hard (ATP-H) datasets (*video-average for FrozenBiLM and video-text-average for JustAsk; † percentage decrease in the number of multiplication operations due to QUAG-attention).

	FrozenBiLM				JustAsk			
	A-QA	M-QA	N-QA	ATP-H	A-QA	M-QA	N-QA	ATP-H
Baseline	43.6	46.6	55.8	55.7	38.7	41.8	53.8	44.0
Language-only	32.2	33.2	55.7	55.8	28.2	29.9	42.2	42.0
Video-only	0.1	0.0	20.2	20.1	2.6	6.7	39.1	23.0
SC: unimodal	2.4	1.0	19.8	21.4	38.5	41.5	53.6	43.6
SC: crossmodal	32.3	32.8	56.0	55.6	38.3	41.3	53.5	44.3
SC: video	43.1	45.7	55.8	55.7	38.2	41.3	53.4	44.3
SC: text	1.4	1.0	20.5	21.1	38.6	41.5	53.7	43.6
QUAG-atten*	43.0	45.8	55.6	55.9	38.0	41.0	53.5	44.1
Δ MulOps†		13.6%				68.0%		

some crossmodal interactions (video (query) and text (key) only). On the other hand, JustAsk model does not learn to fuse the modalities across the datasets and relies largely on the text-modality. Note that while the relative performance drop in the *classical* language-only and video-only settings for JustAsk and FrozenBiLM models on ActivityNet-QA and MSRVT-TQA is similar, QUAG points out the differences in their sub-optimal representations.

We use the results from QUAG to apply QUAG-attention on FrozenBiLM and JustAsk that reduce the number of multiplication operations by **13.6%** and **68.0%** respectively, for a less than 1% drop in performance consistently for all the datasets. However, this raises serious concerns because models can learn to *hack* their way around the accuracy metrics for leveraging shortcuts. The supposedly multimodal datasets contain biases and the evaluation metrics do not penalize shortcut learning and provide a false confidence about the abilities of the model. This raises the follow-up question – **How can we confidently benchmark multimodal understanding in VideoQA models?**

3 DOES MULTIMODAL SUB-OPTIMALITY STEMS FROM DATASET BIASES?

Sub-optimality in model representations and shortcut learning can stem from a combination of facets like dataset biases, model architecture (Jelassi et al., 2022), optimization method (Gunasekar et al., 2018), learning paradigm (Liu et al., 2022) etc. Hence, to ablate the effect of dataset biases we curate CLAVI, a diagnostic dataset with temporal counterfactuals in questions and videos that necessitates joint multimodal understanding and penalizes simple shortcut learning. CLAVI is not positioned to replace existing datasets but rather to supplement them, enhancing the understanding of VideoQA models. We finetune the VideoQA models on CLAVI with the prescribed model architecture and training recipe to study and diagnose the representational prowess of the pretrained models.

3.1 CLAVI: DIAGNOSING THROUGH COUNTERFACTUALS

CLAVI consists of 6,018 videos and 114,342 questions (72,770 train and 41,572 test). It contains simple yes-no questions to probe the absolute temporal location of a single action (beginning/end) or the occurrence sequence for a pair of non-overlapping actions (before/after). CLAVI allows for systematic benchmarking and diagnosis of joint multimodal understanding through the lens of balanced video and question temporal counterfactuals. We use question templates to automatically curate the question-answer pairs from the temporal grounding annotations of Charades-STA (Gao et al., 2017). To create temporal **counterfactuals** in the question domain, we replace *before* with *after* and *beginning* with *end* and vice versa. Further, we create temporal **counterfactuals** in the video domain by swapping only the action-segments in the video as shown in Figure 2. We exhaustively consider all the compositions of temporal **counterfactuals** in video and question domains to create balanced **counterfactuals** instances for systematic assessment of multimodal understanding in videos.




Figure 2: Illustrative example of the creation of CLAVI. In the original video (V), the action “turning on a light” (Event A; blue pane) follows “holding clothes” (Event B; brown pane). To create a counterfactual video (V’), we swap the action segments without manipulating the segment separating them. The questions (Q), along with their counterfactual (Q’), are curated for each of the videos. Note that the color of the question panel reflects the correct answer (green for “yes”, pink for “no”). We provide the list of questions in Table 2.

Table 2: List of questions and their counterfactuals in CLAVI for the illustrative example in Fig. 2. For brevity, we present 4 (out of 8) negative control (NC) questions for BA type; comprehensive list in Appendix A.3.3.

Question (Q)	Counterfactual Question (Q')
Existence (E) type	
Was someone turning on light?	
Was someone holding clothes?	
Existence (E) type – (NC)	
Was someone washing mirror?	
Beginning/End (BE) type	
Was the person turning on light at the beginning ?	Was the person turning on light at the end ?
Was the person holding clothes at the end ?	Was the person holding clothes at the beginning ?
Before/After (BA) type	
Did turning on light happen before holding clothes?	Did turning on light happen after holding clothes?
Did holding clothes happen after turning on a light?	Did holding clothes happen before turning on light?
Before/After (BA) type – (NC)	
Did turning on light happen before washing mirror?	Did turning on a light happen after washing mirror?
Did holding clothes happen after washing mirror?	Did holding clothes happen before washing mirror?

We briefly explain the design principle of CLAVI. We choose temporal sequence counterfactuals to benchmark joint multimodal understanding because it requires **unimodal understanding** within the modalities (sensitive to the sequence of (i) frames in the video; (ii) objects, verbs and temporal phrases in the question) as well as **crossmodal understanding** (relating the sequence of actions in the video with that of the question). This also makes temporal ordering as one of the **fundamental** elements of VideoQA. Using yes-no questions with balanced negative instances allows us to have questions that are **unambiguous**, and answers that are **mutually exclusive** and **equally informative** to not be eliminated by prior biased knowledge. We deliberately maintain a simple design for question templates and answer vocabulary that excludes other abilities such as language comprehension, commonsense reasoning, and long-term memory to facilitate **isolated diagnostic analysis** of joint multimodal understanding. Also, we ensure that the dataset size is sufficiently large, as compared to the existing datasets, so that the models do not overfit (Appendix A.3.2).

Based on the temporal cue in the question, CLAVI contains three question types – **Existence (E)**, **Beginning/End (BE)** and **Before/After (BA)**. Further, we define negative control questions containing actions that do not occur in the video (that is, the answer is always “no”) for E and BA types as shown in Table 2. Answering the negative control does not require understanding temporal cues in language and video. Hence, it serves the dual purpose of sanity check of learning and a baseline for learning by temporal shortcuts. We remove the bias against *beginning* and *end* by randomly

Table 3: Test performance (% accuracy) on CLAVI after finetuning

Metric	JustAsk (Yang et al., 2021b)	FrozenBiLM (Yang et al., 2022a)	Singularity-T (Lei et al., 2023)	All-In-One+ (Wang et al., 2023a)
Balanced Acc	72.2 ± 0.2	80.5 ± 0.1	76.8 ± 0.5	73.9 ± 0.1
C _{Acc_V}	50.6 ± 0.3	74.0 ± 0.1	47.2 ± 1.1	49.6 ± 0.5
C _{Acc_T}	50.3 ± 0.1	75.5 ± 0.1	47.0 ± 1.0	49.5 ± 0.3
C _{Acc_V} -control	98.0 ± 0.2	93.2 ± 0.2	92.7 ± 2.0	98.1 ± 0.5
C _{Acc_T} -control	98.2 ± 0.2	93.7 ± 0.2	93.5 ± 1.9	98.2 ± 0.7
C _{Acc_V} -counter	3.6 ± 0.1	54.1 ± 0.2	1.7 ± 0.2	1.2 ± 0.3
C _{Acc_T} -counter	2.4 ± 0.1	57.2 ± 0.2	0.5 ± 0.2	0.8 ± 0.1

extending the boundaries of the action-segments in the video. The detailed curation process and dataset statistics are presented in Appendix A.3.1.

We want to evaluate the sensitivity of the model to the temporal cues in language and video independently. Hence, we define consistent accuracies. **Given a question, if the model predicts the answers correctly for both – the video and its corresponding counterfactual video, it is called video-consistent. Similarly, for a given video, if the model correctly answers a question and its corresponding counterfactual question, it is called text-consistent.** The proportion of video and question consistent predictions are reported as **video-consistent accuracy (C_{Acc_V})** and **text-consistent accuracy (C_{Acc_T})** respectively. We report the consistent accuracies separately for the **control subset** (E, E-NC, and BA-NC question types) and the **counterfactual subset** (BE and BA question types). The control subset can be answered by leveraging shortcuts while answering the counterfactual subset necessitates joint multimodal understanding.

3.2 EXPERIMENT

We finetune and evaluate 4 models: JustAsk (Yang et al., 2021a), FrozenBiLM (Yang et al., 2022b), Singularity-Temporal (Lei et al., 2023) and All-In-One+ (Wang et al., 2023b) on CLAVI using the official finetuning instructions (Appendix A.3.5). We follow the same experimental settings as discussed in Section 2.5. To account for class imbalance in the answers, we use balanced accuracy for validation and testing. The results are summarized in Table 3. All the models have greater than 70% balanced accuracy. At first, it might give an illusion of good multimodal understanding in VideoQA models. However, the consistent accuracy metrics demystify the illusion.

Text and video consistent accuracies are greater than 90% for the control subset for all the models. This is because, unlike the counterfactual subset, the control subset does not require coupled understanding. That is, the model can answer it correctly by simple shortcuts – irrespective of the context of the negative control action in the question and the location of the object and/or the action in the video. However, for achieving high consistent accuracies on the counterfactual subset, the model needs to jointly understand the order of the events and the temporal cues in the question along with the order of the events in the video. We get significantly lower consistent accuracies (less than 4%) for the counterfactual subset, except for FrozenBiLM. Overall, this means that the other **models are able to exploit shortcuts but unable to learn joint multimodal representations.**

How can we be sure that FrozenBiLM is not learning spurious shortcuts on CLAVI? We find that the video-average QUAG-attention on FrozenBiLM cause the C_{Acc_T}-counter and C_{Acc_V}-counter to drop to **23%** and **3.6%** respectively. That is, the performance on the counterfactual subset significantly drops under multimodal impairment. However, C_{Acc_T}-control and C_{Acc_V}-control values increase to **98.6%** and **99.2%** respectively, perhaps because QUAG-attention promotes reliance on shortcuts, and the control subset can be solved easily by shortcuts. These confirm FrozenBiLM’s reliance on multimodal representations for its high performance relative to the other models.

Beyond the consistency accuracy metrics we can use CLAVI for diverse representation analyses. As an example, we present a qualitative representation sensitivity analysis for FrozenBiLM in Appendix A.3.7. We align the attention matrices for counterfactual pairs and find that the representations of correctly answered counterfactual pairs are more distinct than the wrongly answered pairs to validate joint multimodal understanding.

4 RELATED WORK

Dataset Biases: Works in NLP (Papadimitriou et al., 2022; Sinha et al., 2021), vision (Brendel & Bethge, 2019) and vision-language (Yuksekgonul et al., 2023) demonstrate that models can achieve high performance without even understanding the sequence of the embeddings. This is partly because the current benchmarks have unintended biases that could potentially be exploited by models to learn shortcuts; hence accuracy is not always a faithful metric (Pham et al., 2021; Yuksekgonul et al., 2023; Kaffe & Kanan, 2017; Sevilla-Lara et al., 2021). For VideoQA, MovieQA (Tapaswi et al., 2016) and TVQA (Lei et al., 2018) datasets are biased towards plot understanding or dialogue comprehension (Winterbottom et al., 2020). Biases are not always immediately apparent; for example, Social-IQ (Zadeh et al., 2019) contains sentiment-biased annotations (Gat et al., 2021). Moreover, statistical regularities like answer length, answer frequency (Goyal et al., 2017; Agrawal et al., 2016) and co-occurrence (Dancette et al., 2021a; Manjunatha et al., 2019; Subramanian et al., 2019) introduce spurious features. Overall, these biases allow the models learn shortcuts (Geirhos et al., 2020) that circumvent multimodal reasoning (Chao et al., 2018; Ye & Kovashka, 2021). While synthetic VideoQA benchmarks such as VQuAD (Gupta et al., 2022), CLEVRER (Yi et al., 2019) have been carefully curated to mitigate many biases, they are unable to capture the intricate dynamics of the real world. Recently proposed Preception Test (Pătrăucean et al., 2023), while comprehensive, does not contain diagnostic metrics that penalize the effect of shortcut learning. We curate CLAVI by systematically augmenting real-world videos to faithfully represent the complexity of the physical world while controlling the biases to confidently evaluate multimodal temporal understanding.

Shortcut Learning: Tangential to the bias amelioration methods (Cadene et al., 2019; Clark et al., 2019), Lei et al. (2023) and Winterbottom et al. (2020) achieve state-of-the-art performance with simple models by leveraging VideoQA dataset shortcuts in the model. ATP (Buch et al., 2022) demonstrates single frame bias by re-training the models with an informative frame-selection module to achieve competitive performance. Perceptual Score (Gat et al., 2021) quantifies modality bias in terms of relative performance drop under modality-permutation operation. QUAG combines these ideas to evaluate the dependence of models on shortcuts for circumventing multimodal understanding in terms of performance drop under multimodal representation collapse. Unlike others, it assists in identifying sub-optimal representations in a combined model-dataset approach at test time.

Leveraging Counterfactuals: We share our motivation for developing CLAVI with VQA-CP (Agrawal et al., 2018): that iid train-test splits in the presence of strong priors leads to learning via shortcuts. However, rather than reducing the bias by mining new complementary image instances, CLAVI weakens prior of multimodal understanding with synthesized balanced video-question temporal hard-negatives. Concurrent to our work, Momeni et al. (2023) and Wang et al. (2023c) have employed hard-negatives for improving verb-understanding in VideoQA models. Bagad et al. (2023) stitch pairs of unrelated videos to improve the temporal understanding of video-language models. However, unlike CLAVI that uses synthesized negative video instance from the same video, stitched video dataset cannot be a robust diagnostic benchmark because the incoherent contexts can be exploited as a static bias shortcut (Choi et al., 2019).

5 CONCLUSION

In this work, we perform a rigorous analysis of VideoQA models, focusing on multimodal representations. We introduced QUAG, a coupled dataset-model approach, to conduct a systematic analysis of learned multimodal representations. It provides deep insights into *how* the models infer and *why* the models fail. We found that VideoQA models can learn shortcuts on seemingly multimodal datasets without truly learning to align and fuse the information both – within and between the modalities. Using this understanding, we developed QUAG-attention and exposed the sub-optimality of VideoQA models. Hence, we proposed CLAVI, a diagnostic benchmark for analyzing joint multimodal understanding in VideoQA models. With the simple task of temporal ordering we find that most of the current models are unable to jointly infer from text and video modalities. All our proposed approaches – QUAG, QUAG-attention and CLAVI are simple, compute-friendly and generic to be extended to any combination of modalities, datasets and models. Our thorough and systematic dataset-model combined representation analysis provides insights that are shrouded and misled by the standard datasets and evaluation metrics that create the illusion of joint multimodal understanding.

6 ETHICAL STATEMENT

Datasets in machine learning have a history of containing unintentional biases like race, gender, age along with safety and privacy concerns (Birhane & Prabhu, 2021; Peng et al., 2021; Hirota et al., 2022). We curate CLAVI from existing and popular Charades (Sigurdsson et al., 2016) dataset because it is well-studied and collected in controlled settings with consent. Further the owners of Charades ensure the anonymity and privacy of the participants. However, our approach of developing CLAVI is quite generic and can be easily extended to multiple datasets and/or modalities. Also, for automatically generated questions, we ensure to keep the question templates gender neutral.

One of the central goals of our work was to reassert the brittleness in multimodal models by presenting a combined dataset-model centric interpretable representation learning approach through QUAG and CLAVI. We hope our work galvanizes the research community further to not just blindly trust the accuracy score on benchmarks but thoroughly investigate the potential biases that are (1) present in the dataset and (2) are learned by the models.

7 REPRODUCIBILITY STATEMENT

We provide the code to implement QUAG and QUAG-attention in Appendix A.2.2. It can be easily integrated by replacing the self-attention module of the models. Also, QUAG and QUAG-attention are deterministic and non-parametric. Hence, the reproducibility of the results is ensured by reproducibility of the model codebase which we have thoroughly verified.

We provide the scripts to curate CLAVI from Charades along with a representative subset of the dataset in the supplementary material because of licensing restrictions. We have ensured that the dataset curation code is reproducible and it will be released on GitHub under GPL3.0 license. Following the guidelines of Gebru et al. (2021), we provide a comprehensive documentation of CLAVI with the intended use in Appendix A.3.8. For the finetuning experiments on CLAVI, all the experiments were repeated with three random seeds (0, 42 and 71) and the mean results reported with the standard error. We will release the finetuned checkpoints with the code on GitHub post acceptance notification.

REFERENCES

- Estelle Aflalo, Meng Du, Shao-Yen Tseng, Yongfei Liu, Chenfei Wu, Nan Duan, and Vasudev Lal. VI-interpret: An interactive visualization tool for interpreting vision-language transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 21406–21415, 2022.
- Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. Analyzing the behavior of visual question answering models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1955–1960, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1203. URL <https://aclanthology.org/D16-1203>.
- Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Aniruddha Kembhavi. Don’t just assume; look and answer: Overcoming priors for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4971–4980, 2018.
- Piyush Bagad, Makarand Tapaswi, and Cees GM Snoek. Test of time: Instilling video-language models with a sense of time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2503–2516, 2023.
- Abeba Birhane and Vinay Uday Prabhu. Large image datasets: A pyrrhic win for computer vision? In *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1536–1546. IEEE, 2021.
- Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. Token merging: Your ViT but faster. In *International Conference on Learning Representations*, 2023.

- Wieland Brendel and Matthias Bethge. Approximating cnns with bag-of-local-features models works surprisingly well on imagenet. *International Conference on Learning Representations*, 2019. URL <https://openreview.net/pdf?id=SkfMWhAqYQ>.
- 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*, pp. 2917–2927, 2022.
- Remi Cadene, Corentin Dancette, Matthieu Cord, Devi Parikh, et al. Rubi: Reducing unimodal biases for visual question answering. *Advances in neural information processing systems*, 32, 2019.
- Santiago Castro, Ruoyao Wang, Pingxuan Huang, Ian Stewart, Oana Ignat, Nan Liu, Jonathan Stroud, and Rada Mihalcea. FIBER: Fill-in-the-blanks as a challenging video understanding evaluation framework. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2925–2940, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.209. URL <https://aclanthology.org/2022.acl-long.209>.
- Wei-Lun Chao, Hexiang Hu, and Fei Sha. Being negative but constructively: Lessons learnt from creating better visual question answering datasets. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 431–441, 2018.
- Hila Chefer, Shir Gur, and Lior Wolf. Generic attention-model explainability for interpreting bi-modal and encoder-decoder transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 397–406, 2021.
- Jinwoo Choi, Chen Gao, Joseph CE Messou, and Jia-Bin Huang. Why can’t I dance in the mall? Learning to mitigate scene bias in action recognition. *Advances in Neural Information Processing Systems*, 32, 2019.
- Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. Don’t take the easy way out: Ensemble based methods for avoiding known dataset biases. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 4069–4082, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1418. URL <https://aclanthology.org/D19-1418>.
- David F Crouse. On implementing 2d rectangular assignment algorithms. *IEEE Transactions on Aerospace and Electronic Systems*, 52(4):1679–1696, 2016.
- Corentin Dancette, Rémi Cadène, Damien Teney, and Matthieu Cord. Beyond question-based biases: Assessing multimodal shortcut learning in visual question answering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 1574–1583, October 2021a.
- Corentin Dancette, Remi Cadene, Damien Teney, and Matthieu Cord. Beyond question-based biases: Assessing multimodal shortcut learning in visual question answering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1574–1583, 2021b.
- Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. Tall: Temporal activity localization via language query. In *Proceedings of the IEEE international conference on computer vision*, pp. 5267–5275, 2017.
- Itai Gat, Idan Schwartz, and Alex Schwing. Perceptual score: What data modalities does your model perceive? *Advances in Neural Information Processing Systems*, 34:21630–21643, 2021.
- Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé Iii, and Kate Crawford. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92, 2021.

- Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the V in VQA matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- Suriya Gunasekar, Jason D Lee, Daniel Soudry, and Nati Srebro. Implicit bias of gradient descent on linear convolutional networks. *Advances in neural information processing systems*, 31, 2018.
- Vivek Gupta, Badri N Patro, Hemant Parihar, and Vinay P Nambodiri. Vquad: Video question answering diagnostic dataset. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 282–291, 2022.
- Jiachang Hao, Haifeng Sun, Pengfei Ren, Jingyu Wang, Qi Qi, and Jianxin Liao. Can shuffling video benefit temporal bias problem: A novel training framework for temporal grounding. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVI*, pp. 130–147. Springer, 2022.
- Jack Hessel and Lillian Lee. Does my multimodal model learn cross-modal interactions? it’s harder to tell than you might think! In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 861–877, 2020.
- Yusuke Hirota, Yuta Nakashima, and Noa Garcia. Gender and racial bias in visual question answering datasets. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pp. 1280–1292, 2022.
- Soumya Jahagirdar, Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. Watching the news: Towards videoqa models that can read. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 4441–4450, 2023.
- Samy Jelassi, Michael Sander, and Yuanzhi Li. Vision transformers provably learn spatial structure. *Advances in Neural Information Processing Systems*, 35:37822–37836, 2022.
- Baoxiong Jia, Ting Lei, Song-Chun Zhu, and Siyuan Huang. Egotaskqa: Understanding human tasks in egocentric videos. *Advances in Neural Information Processing Systems*, 35:3343–3360, 2022.
- Kushal Kafle and Christopher Kanan. An analysis of visual question answering algorithms. In *Proceedings of the IEEE international conference on computer vision*, pp. 1965–1973, 2017.
- Matthew Kowal, Mennatullah Siam, Md Amirul Islam, Neil DB Bruce, Richard P Wildes, and Konstantinos G Derpanis. A deeper dive into what deep spatiotemporal networks encode: Quantifying static vs. dynamic information. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13999–14009, 2022.
- Jie Lei, Licheng Yu, Mohit Bansal, and Tamara Berg. TVQA: Localized, compositional video question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 1369–1379, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1167. URL <https://aclanthology.org/D18-1167>.
- Jie Lei, Tamara Berg, and Mohit Bansal. Revealing single frame bias for video-and-language learning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 487–507, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.29. URL <https://aclanthology.org/2023.acl-long.29>.
- Paul Pu Liang, Yiwei Lyu, Gunjan Chhablani, Nihal Jain, Zihao Deng, Xingbo Wang, Louis-Philippe Morency, and Ruslan Salakhutdinov. Multiviz: Towards visualizing and understanding multimodal models. In *The Eleventh International Conference on Learning Representations*, 2022a.

- Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. Foundations and recent trends in multimodal machine learning: Principles, challenges, and open questions. *arXiv preprint arXiv:2209.03430*, 2022b.
- Paul Pu Liang, Yun Cheng, Ruslan Salakhutdinov, and Louis-Philippe Morency. Multimodal fusion interactions: A study of human and automatic quantification. *arXiv preprint arXiv:2306.04125*, 2023.
- Xiao Liu, Ankur Sikarwar, Joo Hwee Lim, Gabriel Kreiman, Zenglin Shi, and Mengmi Zhang. Reason from context with self-supervised learning. *arXiv preprint arXiv:2211.12817*, 2022.
- Varun Manjunatha, Nirat Saini, and Larry S Davis. Explicit bias discovery in visual question answering models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9562–9571, 2019.
- Liliane Momeni, Mathilde Caron, Arsha Nagrani, Andrew Zisserman, and Cordelia Schmid. Verbs in action: Improving verb understanding in video-language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 15579–15591, October 2023.
- Nihal Murali, Aahlad Manas Puli, Ke Yu, Rajesh Ranganath, and Kayhan Batmanghelich. Shortcut learning through the lens of early training dynamics. *arXiv preprint arXiv:2302.09344*, 2023.
- Mayu Otani, Yuta Nakahima, Esa Rahtu, and Janne Heikkilä. Uncovering hidden challenges in query-based video moment retrieval. In *The British Machine Vision Conference (BMVC)*, 2020.
- Isabel Papadimitriou, Richard Futrell, and Kyle Mahowald. When classifying grammatical role, bert doesn’t care about word order... except when it matters. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 636–643, 2022.
- Mandela Patrick, Dylan Campbell, Yuki M. Asano, Ishan Misra Florian Metze, Christoph Feichtenhofer, Andrea Vedaldi, and Joao F. Henriques. Keeping your eye on the ball: Trajectory attention in video transformers. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Kenneth L Peng, Arunesh Mathur, and Arvind Narayanan. Mitigating dataset harms requires stewardship: Lessons from 1000 papers. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- Thang Pham, Trung Bui, Long Mai, and Anh Nguyen. Out of order: How important is the sequential order of words in a sentence in natural language understanding tasks? In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 1145–1160, 2021.
- Viorica Pătrăucean, Lucas Smaira, Ankush Gupta, Adrià Recasens Contente, Larisa Markeeva, Dylan Banarse, Skanda Koppula, Joseph Heyward, Mateusz Malinowski, Yi Yang, Carl Dorsch, Tatiana Matejovicova, Yury Sulsky, Antoine Miech, Alex Frechette, Hanna Klimczak, Raphael Koster, Junlin Zhang, Stephanie Winkler, Yusuf Aytar, Simon Osindero, Dima Damen, Andrew Zisserman, and João Carreira. Perception test: A diagnostic benchmark for multimodal video models. In *Advances in Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=HYEGXFnPog>.
- Laura Sevilla-Lara, Shengxin Zha, Zhicheng Yan, Vedanuj Goswami, Matt Feiszli, and Lorenzo Torresani. Only time can tell: Discovering temporal data for temporal modeling. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 535–544, 2021.
- Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pp. 510–526. Springer, 2016.
- Koustuv Sinha, Robin Jia, Dieuwke Hupkes, Joelle Pineau, Adina Williams, and Douwe Kiela. Masked language modeling and the distributional hypothesis: Order word matters pre-training for little. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 2888–2913, 2021.

- Sanjay Subramanian, Sameer Singh, and Matt Gardner. Analyzing compositionality in visual question answering. *ViGIL@ NeurIPS*, 7, 2019.
- Makarand Tapaswi, Yukun Zhu, Rainer Stiefelwagen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. MovieQA: Understanding stories in movies through question-answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4631–4640, 2016.
- Alex Jinpeng Wang, Yixiao Ge, Rui Yan, Ge Yuying, Xudong Lin, Guanyu Cai, Jianping Wu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. All in one: Exploring unified video-language pre-training. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023a.
- Jinpeng Wang, Yixiao Ge, Rui Yan, Yuying Ge, Kevin Qinghong Lin, Satoshi Tsutsui, Xudong Lin, Guanyu Cai, Jianping Wu, Ying Shan, et al. All in one: Exploring unified video-language pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6598–6608, 2023b.
- Xingbo Wang, Jianben He, Zhihua Jin, Muqiao Yang, Yong Wang, and Huamin Qu. M2lens: Visualizing and explaining multimodal models for sentiment analysis. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):802–812, 2021.
- Zhenhailong Wang, Ansel Blume, Sha Li, Genglin Liu, Jaemin Cho, Zineng Tang, Mohit Bansal, and Heng Ji. Paxion: Patching action knowledge in video-language foundation models, 2023c.
- T. Winterbottom, S. Xiao, A. McLean, and N. Al Moubayed. On modality bias in the TVQA dataset. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2020.
- Torsten Wörtwein, Lisa Sheeber, Nicholas Allen, Jeffrey Cohn, and Louis-Philippe Morency. Beyond additive fusion: Learning non-additive multimodal interactions. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pp. 4681–4696, 2022.
- 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.
- 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 (CVPR)*, pp. 9777–9786, June 2021.
- Junbin Xiao, Pan Zhou, Tat-Seng Chua, and Shuicheng Yan. Video graph transformer for video question answering. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVI*, pp. 39–58. Springer, 2022.
- 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.
- Yufei Xu, Qiming Zhang, Jing Zhang, and Dacheng Tao. Vitae: Vision transformer advanced by exploring intrinsic inductive bias. *Advances in Neural Information Processing Systems*, 34:28522–28535, 2021.
- Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Just ask: Learning to answer questions from millions of narrated videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1686–1697, 2021a.
- Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Just ask: Learning to answer questions from millions of narrated videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1686–1697, 2021b.
- Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Zero-shot video question answering via frozen bidirectional language models. In *Advances in Neural Information Processing Systems*, 2022a.

- Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Zero-shot video question answering via frozen bidirectional language models. In *NeurIPS 2022-36th Conference on Neural Information Processing Systems*, 2022b.
- Keren Ye and Adriana Kovashka. A case study of the shortcut effects in visual commonsense reasoning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 3181–3189, 2021.
- Kexin Yi, Chuang Gan, Yunzhu Li, Pushmeet Kohli, Jiajun Wu, Antonio Torralba, and Joshua B Tenenbaum. Clever: Collision events for video representation and reasoning. In *International Conference on Learning Representations*, 2019.
- 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*, volume 33, pp. 9127–9134, 2019.
- Mert Yuksekogonul, Federico Bianchi, Pratyusha Kalluri, Dan Jurafsky, and James Zou. When and why vision-language models behave like bags-of-words, and what to do about it? In *International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=KRLUvxh8uaX>.
- Amir Zadeh, Michael Chan, Paul Pu Liang, Edmund Tong, and Louis-Philippe Morency. Social-iq: A question answering benchmark for artificial social intelligence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8807–8817, 2019.
- Yaoyao Zhong, Wei Ji, Junbin Xiao, Yicong Li, Weihong Deng, and Tat-Seng Chua. Video question answering: Datasets, algorithms and challenges. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 6439–6455, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.emnlp-main.432. URL <https://aclanthology.org/2022.emnlp-main.432>.

A APPENDIX

A.1 ADDITIONAL RELATED WORKS

Multimodal Fusion Interpretability and Visualization: Liang et al. (2022b) and Liang et al. (2023) analyze multimodal fusion interactions along the dimensions of response, information, and mechanics. The closest alignment of QUAG and CLAVI is at the interface of multimodal fusion response and mechanics. Previous works have quantified the presence or absence of specific kinds of modality interactions through the study of datasets (Dancette et al., 2021b), models (Chefer et al., 2021), projections onto simpler models (Hessel & Lee, 2020; Wörtwein et al., 2022) and visualization studies (Liang et al., 2022a; Aflalo et al., 2022; Wang et al., 2021). However, using QUAG, unlike other methods, we perform a combined dataset-model analysis without any additional parameters or finetuning.

A.2 QUAG

A.2.1 ATTENTION MAP VISUALIZATION

We provide a visualization example of the attention values before and after short-circuiting operations in Figure 3.

A.2.2 CODE

Below is the implementation of QUAG as an augmentation of the existing self-attention function. We use row-wise average and replace operation in each if-clause statements, while ignoring the padding, to ablate the effect of the quadrant.

```
1 def self_attention(inputs, mask, dim_model, l_v, l_t, quads):
2     # Inputs:
```




Figure 3: Visualization of the first attention head, as a heatmap, from the second layer of JustAsk model with $l_v = 20$ and $l_t = 20$. Note that here the text embeddings are pre-concatenated to the video embedding in the input. The lengths of the video and text tokens are 9 and 7 respectively. The text and video tokens are individually padded to length 20 each. We visualize (a) the original attention values and (b)-(d) after short-circuiting (SC) operations.

```

3  # inputs: Tensor of shape (batch_size, sequence_length,
   # dim_model)
4  # mask: Tensor of shape (batch_size, sequence_length)
5  # dim_model: Dimension of the model (e.g., 512)
6  # l_v: int    maximum length of video tokens
7  # l_t: int    maximum length of question tokens
8  # quads: list containing elements from {'VV', 'VT', 'TV', 'TT'}
9  query = linear_transform_query(inputs)
10 key = linear_transform_key(inputs)
11 value = linear_transform_value(inputs)
12 attention_scores = compute_attention_scores(query, key, mask)
13 apply_quag(attention_scores, mask, l_v, l_t, quads)
14 attended_output = apply_attention_scores(attention_scores, value)
15 return attended_output
16
17 def compute_attention_scores(query, key, mask):
18     scaled_dot_product = dot_product(query, key) / sqrt(dim_model)
19     attention_scores = softmax(scaled_dot_product + (1 - mask) * -1e9)
20     return attention_scores
21
22 def apply_quag(attention_scores, mask, l_v, l_t, quads):
23     if 'VV' is in quads:
24         replace_with_rowwise_average(attention_scores[:, :l_v, :l_v],
25         mask[:, :l_v, :l_v])
26     if 'VT' is in quads:
27         replace_with_rowwise_average(attention_scores[:, :l_v, -l_t:],
28         mask[:, :l_v, -l_t:])

```



```

27     if 'TV' is in quads:
28         replace_with_rowwise_average(attention_scores[:, -l_t:, :l_v],
29         mask[:, -l_t:, :l_v])
30     if 'TT' is in quads:
31         replace_with_rowwise_average(attention_scores[:, -l_t:, -l_t
32         :], mask[:, -l_t:, -l_t:])
33
34 def replace_with_rowwise_average(scores, mask):
35     rowwise_sum = sum(scores, axis=-1)
36     rowwise_mean = rowwise_sum / sum(mask, axis=-2)
37     expanded_rowwise_mean = expand_dims(rowwise_mean, axis=-1)
38     replace_elements(scores, expanded_rowwise_mean)
39
40 def apply_attention_scores(attention_scores, value):
41     attended_output = dot_product(attention_scores, value)
42     return attended_output

```

Next, we provide the code for QUAG-attention. QUAG-attention modifies the existing self-attention block in the fusion module by replacing the block with the block average. We also demonstrate the normalizing the softmax function so that the each single average sequence is representative of the constituent sequences.

```

1 def quag_attention(inputs, mask, dim_model, l_v, l_t, type):
2     # Inputs:
3     # inputs: Tensor of shape (batch_size, sequence_length,
4     # dim_model)
5     # mask: Tensor of shape (batch_size, sequence_length)
6     # dim_model: Dimension of the model (e.g., 512)
7     # l_v: int    maximum length of video tokens
8     # l_t: int    maximum length of question tokens
9     # type: one of 'text', 'video', 'text-video'
10    query = linear_transform_query(inputs)
11    avg_input = compute_avg_input(inputs, l_v, l_t, type)
12    key = linear_transform_key(avg_input)
13    value = linear_transform_value(avg_input)
14    mask = apply_mask(mask, l_v, l_t, type)
15    scaled_dot_product = compute_scaled_dot_product(query, key,
16    dim_model, mask)
17    attention_scores = softmax(scaled_dot_product)
18    attended_output = apply_attention_scores(attention_scores, value)
19    return attended_output
20
21 def compute_avg_input(inputs, l_v, l_t, type):
22     if type == "video":
23         avg_upper_block = sum(inputs[:, :l_v, :], axis=-2)
24         avg_upper_block = expand_dims(avg_upper_block, axis=1)
25         avg_input = concatenate((avg_upper_block, inputs[:, :-l_t, :])
26         , axis=1)
27     elif type == "text":
28         avg_lower_block = sum(inputs[:, :-l_t, :], axis=-2)
29         avg_lower_block = expand_dims(avg_lower_block, axis=1)
30         avg_input = concatenate((inputs[:, :l_v, :], avg_lower_block),
31         axis=1)
32     elif type == "text-video":
33         avg_upper_block = sum(inputs[:, :l_v, :], axis=-2)
34         avg_upper_block = expand_dims(avg_upper_block, axis=1)
35         avg_lower_block = sum(inputs[:, :-l_t, :], axis=-2)
36         avg_lower_block = expand_dims(avg_lower_block, axis=1)
37         avg_input = concatenate((avg_upper_block, avg_lower_block),
38         axis=1)
39     return avg_input
40
41 def apply_mask(mask, l_v, l_t, type):
42     mask = expand_dims(mask, axis=-1)

```

Table 4: Fine-grained performance of JustAsk on ActivityNet-QA

Config	Motion	Spatial	Temp	Y/N	Color	Obj	Loc	Num	Other
Baseline	30.6	19.9	4.9	64.2	34.7	26.7	35.5	48.9	36.8
Lang-only	1.4	9.1	4.3	51.8	28.7	23.0	16.6	46.9	29.1
Vid-only	20.3	0.9	1.8	0.0	0.0	1.6	1.3	0.0	0.7
SC: unimodal	30.1	19.1	4.9	63.9	33.6	26.4	36.8	48.4	37.0
SC: crossmodal	28.0	18.9	4.8	64.7	34.7	25.8	35.5	48.5	36.4
SC: text	30.4	19.3	5.0	64.1	34.0	26.4	35.5	46.7	37.2
SC: video	28.6	18.8	4.5	64.3	34.6	25.5	35.5	48.4	36.1
QUAG-attention	28.1	18.5	4.9	64.1	33.6	25.2	34.7	48.0	36.6

```

38     mask = tile(mask, [1, 1, sequence_length])
39
40     if "video" in type:
41         video_length = sum(mask[:, :l_v, 0], axis=1)
42         video_length = expand_dims(video_length, axis=-1)
43         scaled_dot_product[:, :, 0] = scaled_dot_product[:, :, 0] *
log(video_length)
44         upper_mask = ones(mask.shape[0], mask.shape[1], 1)
45         mask = concatenate((upper_mask, mask[:, :, l_v:]), axis=-1)
46
47     if "text" in type:
48         text_length = sum(mask[:, :-l_t, 0], axis=1)
49         text_length = expand_dims(text_length, axis=-1)
50         scaled_dot_product[:, :, -1] = scaled_dot_product[:, :, -1] *
log(text_length)
51         lower_mask = ones(mask.shape[0], mask.shape[1], 1)
52         mask = concatenate((mask[:, :, :-l_t], lower_mask), axis=-1)
53
54     return mask
55
56 def compute_scaled_dot_product(query, key, dim_model, mask):
57     scaled_dot_product = dot_product(query, key) / sqrt(dim_model)
58     return scaled_dot_product
59
60 def apply_attention_scores(attention_scores, value):
61     attended_output = dot_product(attention_scores, value)
62     return attended_output

```

A.2.3 EXPERIMENT DETAILS

As mentioned in the main manuscript, we use the official checkpoints and code of JustAsk [website] and FrozenBiLM [website]. For all the experiments with JustAsk, we use the checkpoints of the model pretrained on HowToVQA69M and WebVidVQA3M. For FrozenBiLM, we use the WebVid10M-pretrained checkpoint for all our experiments. Since QUAG operates at inference time, we do not need to perform any training. Since the model owners do not report results on NeXT-QA, we finetune the models with the official recipe to achieve performance similar to that independently reported by others Xiao et al. (2022). While FrozenBiLM can also take subtitles as the input, for fair comparison, we do not pass it in any of the experiments. We provide the hardware details in the main manuscript.

A.2.4 FINEGRAINED ACCURACIES

A.2.5 JUSTASK MODEL

We present the fine-grained performance of JustAsk on the discussed datasets in Tables 4, 5, 6, and 7

Table 5: Fine-grained performance of JustAsk on MSRVTT-QA

Config	What	How	Color	Where	Who	When
Baseline	35.8	83.7	51.7	39.4	51.3	82.3
Lang-only	24.3	83.3	43.4	30.5	37.1	72.3
Vid-only	8.5	0.0	3.5	0.4	3.0	10.1
SC: unimodal	35.6	83.3	51.8	39.8	50.8	82.3
SC: crossmodal	35.35	83.75	51.98	39.8	50.8	81.8
SC: text	35.7	83.2	51.8	39.0	50.8	82.1
SC: video	35.4	83.8	51.8	39.8	50.7	81.6
QUAG-attention	35.1	83.5	51.1	38.6	50.2	82.1

Table 6: Fine-grained performance of JustAsk on NeXT-QA

Config	Causal	Temporal	Descriptive
Baseline	50.8	52.8	65.0
Lang-only	39.5	44.3	47.1
Vid-only	39.2	37.9	44.0
SC: unimodal	50.5	52.5	65.3
SC: crossmodal	50.8	51.8	65.0
SC: text	50.7	52.7	65.0
SC: video	50.7	52.1	65.0
QUAG-attention	50.8	52.0	65.1

Table 7: Fine-grained performance of JustAsk on ATP-Hard subset of NeXT-QA

Config	Causal	Temporal
Baseline	44.4	43.4
Lang-only	41.2	43.1
Vid-only	23.5	22.3
SC: unimodal	43.2	43.3
SC: crossmodal	44.2	44.4
SC: text	43.7	43.4
SC: video	44.3	44.4
QUAG-attention	44.2	43.9

A.2.6 FROZENBiLM MODEL

We present the fine-grained performance of FrozenBiLM on the discussed datasets in Tables 8, 9, 10, and 11

A.2.7 ADDITIONAL RESULTS

We evaluated QUAG on All-in-one model and find that, as the authors claim, the model utilizes both – unimodal and cross-modal modality interactions. The results are summarized in Table A.2.7.

A.2.8 PROGRESSIVE QUAG

We apply QUAG progressively to the first n fusion layers of FrozenBiLM to find their relative importance in multimodality. We apply it in the step size of 4 (24 fusion blocks in total) and report the results for ActivityNet-QA (Figure 4) and MSRVTT-QA (Figure 5) Datasets. We consistently find that the first 8 layers are the most important in unimodal interaction and text interactions (and cross-modal interactions as well for ActivityNet-QA).

Table 8: Fine-grained performance of FrozenBiLM on ActivityNet-QA

Config	Motion	Spatial	Temp	Y/N	Color	Obj	Loc	Num	Other
Baseline	30.1	22.5	6.4	75.6	34.6	27.7	37.1	55.8	41.6
Lang-only	2.6	10.5	4.8	63.3	32.3	23.9	16.6	44.7	31.6
Vid-only	0.0	0.0	0.5	0.0	0.0	0.0	0.0	0.0	0.0
SC: unimodal	0.0	0.1	0.1	8.3	0.0	0.0	0.0	1.3	0.5
SC: crossmodal	1.8	11.1	3.88	64.5	32.7	21.7	16.8	46.0	32.1
SC: text	0.0	0.1	0.1	4.4	0.1	0.3	0.0	1.2	0.3
SC: video	28.8	21.8	6.5	75.1	34.3	29.3	36.0	55.3	41.0
QUAG-attention	28.9	22.3	6.0	74.4	35.0	27.3	37.3	54.1	41.1

Table 9: Fine-grained performance of FrozenBiLM on MSRVT-QA

Config	What	How	Color	Where	Who	When
Baseline	40.5	87.2	57.9	41.5	56.6	81.4
Lang-only	27.3	83.6	50.0	35.8	41.2	77.6
Vid-only	0.0	0.0	0.0	0.0	0.0	0.0
SC: unimodal	0.7	0.0	1.2	0.8	1.8	0.2
SC: crossmodal	27.1	83.4	50.9	32.9	41.1	66.3
SC: text	0.3	0.0	0.8	0.0	2.8	0.0
SC: video	39.8	85.5	58.8	41.9	55.4	80.9
QUAG-attention	39.9	86.2	58.1	42.7	55.2	81.1

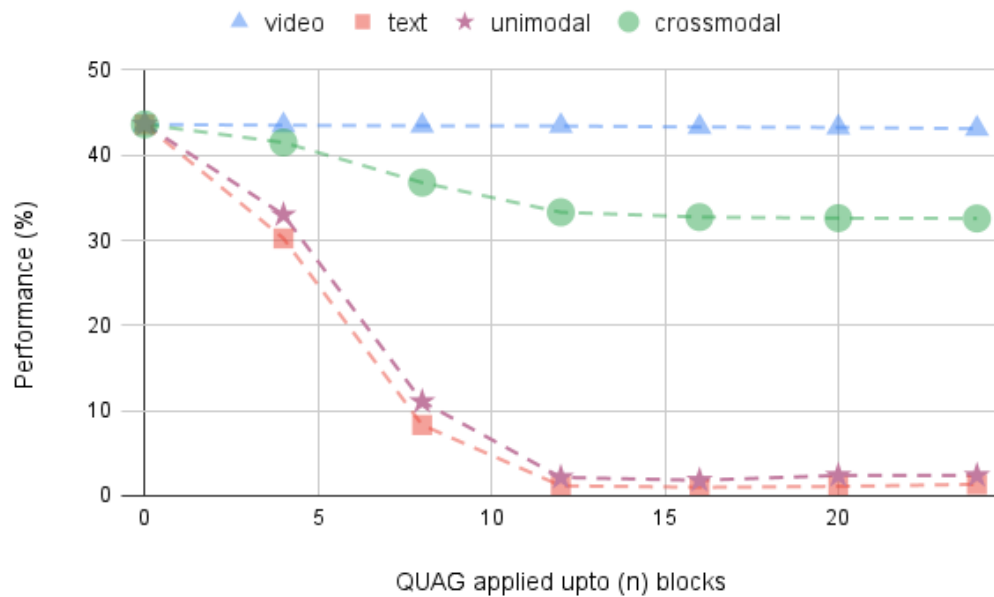


Figure 4: Result of progressive application of QUAG on FrozenBiLM model on ActivityNet-QA dataset

Table 10: Fine-grained performance of FrozenBiLM on NeXT-QA

Config	Causal	Temporal	Descriptive
Baseline	56.0	56.1	54.5
Lang-only	55.9	56.1	54.2
Vid-only	20.7	19.1	20.9
SC: unimodal	19.7	21.1	17.3
SC: crossmodal	56.1	56.5	54.3
SC: text	20.0	21.6	19.9
SC: video	56.1	56.1	54.5
QUAG-attention	55.9	55.8	54.1

Table 11: Fine-grained performance of FrozenBiLM on ATH-Hard subset of NeXT-QA

Config	Causal	Temporal
Baseline	55.2	56.3
Lang-only	55.5	56.2
Vid-only	20.0	20.1
SC: unimodal	20.7	22.5
SC: crossmodal	54.9	56.6
SC: text	20.2	22.3
SC: video	55.3	56.3
QUAG-attention	55.3	56.7

Table 12: Short-circuit (SC) results for All-in-one+ model on ActivityNet-QA (A-QA), and MSRVT-QA (M-QA) datasets.

	All-in-one+	
	ActivityNet-QA	MSRVTT-QA
Acc	41.9	43.1
text_only	23.5	20.8
vid_only	14.2	4.2
SC: unimodal	11.4	3.8
SC: crossmodal	20.6	27.6
SC: video	19.2	12.7
SC: text	5.6	7.3

A.3 CLAVI

A.3.1 DATASET CREATION

We curate CLAVI by leveraging Charades-STA (<https://prior.allenai.org/projects/data/charades/license.txt>) (Gao et al., 2017), containing 9,848 videos of humans performing actions based on a short script written by composing predefined vocabulary that describe multiple daily actions. The videos are annotated with the start and end times of each action. The action category, the start, and the end of each action segment are referred to as the *action tuple*. Each video may contain more than two action tuples. We select pairs of action tuples based on the uniqueness of the action category and complete exclusivity (that is no overlap between the occurrence of the actions). In a given selected pair of action tuples, the two actions along with the inter-action region constitute the video segment. We ensure that the two action categories in the pair are distinct. Additionally, to address temporal boundary ambiguities in the annotations, we filter out segments where either of the selected action classes occurs in close proximity to the segment boundaries




Figure 5: Result of progressive application of QUAG on FrozenBiLM model on NeXT-QA dataset

We also extend the boundaries of the two actions in the pair. We define two boundary extensions – out-extension and in-extension. The out-extension encompasses regions that are not a part of the selected segment but extend outwards in both directions into the original video. Similarly, in-extension extends inwards into the inter-action segment. To avoid temporal position bias (Hao et al., 2022; Otani et al., 2020), the lengths of the extension boundaries are selected randomly. However, since inter-action separation can affect their recognition (Bagad et al., 2023), we constraint the inter-action separation in the original and the corresponding negative video to be the same. That is, the sum of out-extension boundaries is always equal to the sum of in-extension boundaries.

We trim each boundary-extended contiguous segment from the original video to curate a positive video instance. To create the counterfactual video, we swap the boundary-extended action regions as shown in Figure 2. Note that the region between the boundary-extended actions is unaffected. Swapping operation preserves the actions but only alters their chronology, and can be applied independently to question negatives (unlike manipulations like video reversal (Wang et al., 2023c)). This independence provides fine-grained control to create a balanced benchmark for comprehensive analysis.

We create three types of questions using pre-defined templates and action-class annotations:

- 1) **Existence (E) type:** The E-type questions for both the action classes follow the template "Was someone $\langle A \rangle$?", where $\langle A \rangle$ is one of two action classes in video. We use it as a positive control to verify if the model is able to correctly recognize the action classes. We use the exact same question for negative video instance as well, totalling to 4 control (questions, video, answer) instances for a Charades-extracted video segment.
- 2) **Beginning/End (BE) type:** BE type questions the absolute location of the action in the video. The question is of the form, "Was the person $\langle A \rangle$ at the $\{beginning/end\}$?" where $\langle A \rangle$ is one of two action classes in the video, and we select one of *beginning* and *end*. Hence, for a given video and its negative, we have, in total, 8 instances of BE (questions, video, answer) tuples combined. Note that the answer for a given BE question is complemented in the negative video.
- 3) **Before/After (BA) type:** BA type comprises of questions on the relative order of occurrence of actions. The question is of the form "Did $\langle A1 \rangle$ happen $\{after/before\}$ $\langle A2 \rangle$?", where $\langle A1 \rangle$ and $\langle A2 \rangle$ are the selected action classes. We consider all the permutations of action classes. Hence, we have a

total of 8 instances of BA type (questions, video, answer) tuples per extracted video. Similar to BE type, the answer is complemented in the negative video.

Further, we add negative controls for E and BA type questions. A negative control action is an action that does not occur in the video. Since we want to probe only for temporal understanding, we keep the negative control action-class easy to detect by randomly selecting an action-class that does not contain any of the objects or actions in the original video. Hence, answering the negative control does not require understanding temporal cues in language and video and can be answered by object elimination. It serves the dual purpose of sanity check of learning and a baseline for learning by temporal shortcuts. The answer of negative control questions is always false. This adds two E type and sixteen BA type negative control questions for the video and its negative combined. Hence, including the negative control questions, each video in CLAVI is associated with 19 questions: 2 E, 4 BE, 4 BA, 1 E negative control and 8 BA negative controls. The ratio of "yes": "no" answers is 6:13.

A.3.2 COMPARISON WITH EXISTING DATASETS

We provide a comparison of size of CLAVI with established VideoQA datasets in Table 13.

Table 13: Comparison of CLAVI with other other VideoQA datasets sorted in the reverse order of recency.

Dataset	Number of (V,Q,A) samples
MSRVTT-QA (Xu et al., 2017)	243K
ActivityNet-QA (Yu et al., 2019)	58K
Social-IQ QA (Zadeh et al., 2019)	7.5K
NeXT-QA (Xiao et al., 2021)	52K
iVQA (Yang et al., 2021b)	10K
STAR (Wu et al., 2021)	60K
EgoTaskQA (Jia et al., 2022)	40K
FIBER (Castro et al., 2022)	28K
NewsQA (Jahagirdar et al., 2023)	8.6K
CLAVI (Ours)	114K

A.3.3 COMPREHENSIVE LIST OF QUESTIONS

We provide a comprehensive list of the questions for the example presented in Fig 2 of the main paper. We define the actions as: **A**: *turning on light* **B**: *holding clothes* **C**: *washing mirror*, where action A occurs before action B in the original video and action C does not occur anywhere in the original video.

Enlisted below are the questions and its negatives (Q and Q' respectively) for the video (V) (that is event A occurs after event B): Note that the color of the panel is representative of the answer of the question (red: "no", green: "yes").

E-Type:

- Q : Was someone turning on light?
- Q : Was someone holding clothes?

E-Type (negative control):

- Q : Was someone washing mirror?

BE-Type

- Q : Was the person turning on light at the **beginning**?
- Q': Was the person turning on light at the **end**?

Q : Was the person holding clothes at the **end**?

Q': Was the person holding clothes at the **beginning**?

BA-Type

Q : Did turning on light happen **before** holding clothes?

Q': Did turning on light happen **after** holding clothes?

Q : Did holding clothes happen **after** turning on light?

Q': Did holding clothes happen **before** turning on light?

BA-Type (negative-control)

Q': Did washing mirror happen **before** turning on light?

Q': Did washing mirror happen **after** turning on light?

Q': Did turning on light happen **before** washing mirror?

Q': Did turning on light happen **after** washing mirror?

Q': Did washing mirror happen **before** holding clothes?

Q': Did washing mirror happen **after** holding clothes?

Q': Did holding clothes happen **before** washing mirror?

Q': Did holding clothes happen **after** washing mirror?

Enlisted below are the questions and its negatives (Q and Q' respectively) for the negative video instance (V') (that is event B occurs after event A).

E-Type:

Q : Was someone turning on light?

Q : Was someone holding clothes?

E-Type (negative control):

Q : Was someone washing mirror?

BE-Type

Q : Was the person turning on light at the **beginning**?

Q': Was the person turning on light at the **end**?

Q : Was the person holding clothes at the **end**?

Q': Was the person holding clothes at the **beginning**?

BA-Type

Q : Did turning on light happen **before** holding clothes?

Q': Did turning on light happen **after** holding clothes?

Q : Did holding clothes happen **after** turning on light?

Q': Did holding clothes happen **before** turning on light?

BA-Type (negative-control)

Q': Did washing mirror happen **before** turning on light?

Q': Did washing mirror happen **after** turning on light?




Figure 6: Distribution of length of (a) action and (b) video durations

Q': Did turning on light happen before washing mirror?
Q': Did turning on light happen after washing mirror?
Q': Did washing mirror happen before holding clothes?
Q': Did washing mirror happen after holding clothes?
Q': Did holding clothes happen before washing mirror?
Q': Did holding clothes happen after washing mirror?

A.3.4 DATASET METRICS

The duration of individual action in CLAVI lies in the range [4.0 sec, 36.0 sec]; the average length of action is 7.7 ± 3.42 sec. The average video length is 19.95 ± 7.34 secs and the range is [8.67 sec, 65.73 sec]. We plot the distribution of the action and video durations in Fig. 6.

CLAVI consists of **141** unique action classes. Each action class is composed of noun (objects) and verb. There are **37** unique noun classes and **28** unique verb classes. We show the frequency distributions of action, verb and noun classes in Fig 7.

A.3.5 EXPERIMENT DETAILS

As mentioned in the main manuscript, we use the official checkpoints, finetuning code and hyperparameters of JustAsk [website], FrozenBiLM [website], Singularity-Temporal [website], and All-in-one+ [website]. For JustAsk, we use the checkpoint of the model pretrained on HowToVQA69M and WebVidVQA3M. For FrozenBiLM, we use the WebVid10M-pretrained checkpoint. All-in-one+ is pretrained on eight datasets comprising of both images and videos (videos: Webvid, YT-Temporal-180M, HowTo100M and images: CC3M, CC12M, COCO, Visual Genome, SBU Captions). Singularity-Temporal is pretrained on a 17.28M images and video subset (images: COCO, Visual Genome, SBU Captions, CC3M, CC12M and videos: WebVid). We have depicted the finetuning details in Table 14.

A.3.6 FINE-GRAINED ACCURACIES

In Table 15 we provide error bars for the finetuning experiments. The experiments were performed thrice on the same hardware with the same set of hyperparameters.

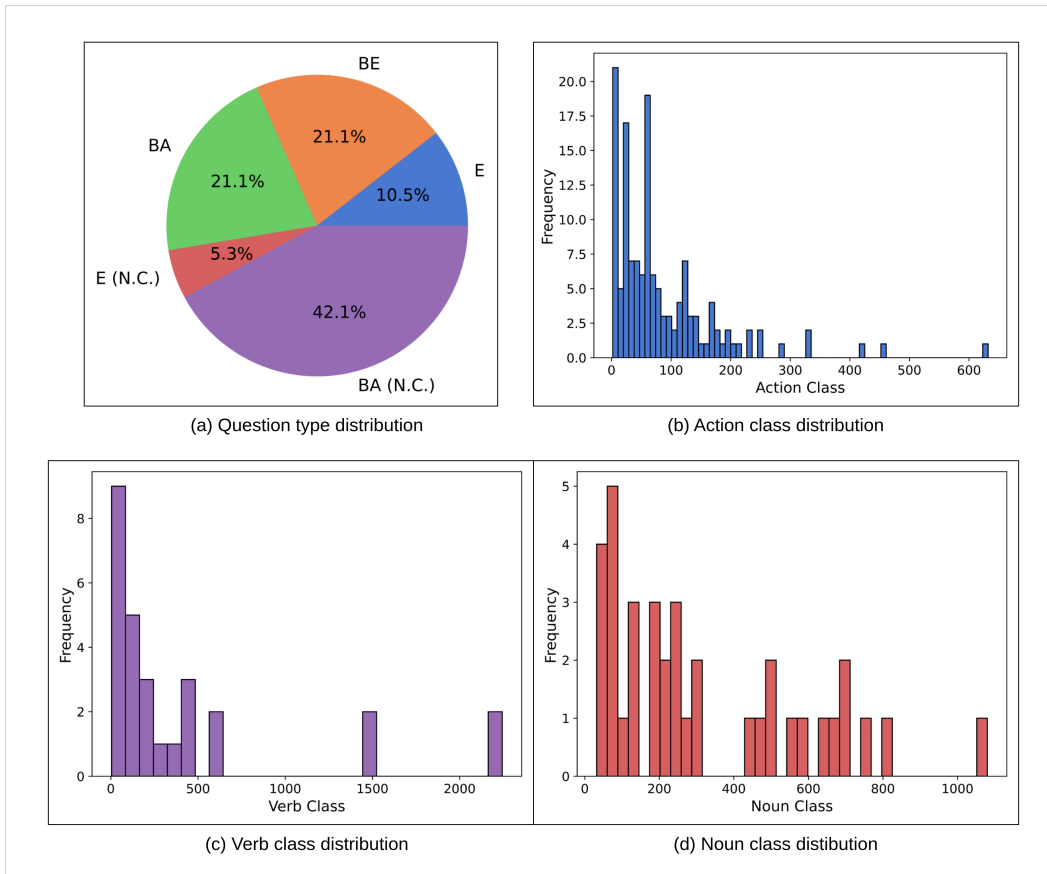


Figure 7: Metrics of the dataset (a) distribution of question types (same for training and testing set), (b) histogram plot of frequencies of action classes (c) histogram plot of frequencies of verb classes (d) histogram plot of frequencies of noun classes.

Table 14: Hyperparameters and checkpoint details of CLAVI finetuning experiment

Model	Checkpoint	Epochs	LR
JustAsk	HowToVQA69M, WebVidVQA3M	20	1.00E-05
FrozenBiLM	WebVid10M	20	5.00E-05
All-In-One+	Webvid, YT-Temporal-180M, HowTo100M, CC3M, CC12M, COCO, Visual Genome, SBU Captions	10	1.00E-04
Singularity-T	COCO, Visual Genome, SBU Captions, CC3M, CC12M, WebVid	20	1.00E-05

Table 15: Fine-grained performance (% of accuracy) on CLAVI for question (Q) and counterfactual question (Q'), video (V) and counterfactual video (V') (Note: N.C. refers to Negative Control)

V/V'	Question	Q/Q'	JustAsk	FrozenBiLM	Singularity-T	All-in-one+
V	E-type	Q	89.55 ± 0.01	87.51 ± 0.00	90.75 ± 0.03	86.08 ± 2.59
	E-type (N.C.)	-	75.28 ± 0.02	88.66 ± 0.00	79.16 ± 0.03	69.34 ± 11.72
	BE-type	Q	69.80 ± 0.07	69.15 ± 0.01	98.23 ± 0.01	99.31 ± 0.84
		Q'	30.58 ± 0.07	73.25 ± 0.01	1.87 ± 0.01	0.73 ± 0.84
	BA-type	Q	27.81 ± 0.02	56.88 ± 0.01	62.55 ± 0.09	25.82 ± 5.49
		Q'	72.31 ± 0.02	86.79 ± 0.01	37.23 ± 0.09	74.31 ± 0.84
BA-type (N.C.)	-	98.23 ± 0.00	96.79 ± 0.00	93.72 ± 0.03	98.44 ± 1.02	
V'	E-type	Q	89.17 ± 0.01	86.96 ± 0.01	90.58 ± 0.02	86.03 ± 2.66
	E-type (N.C.)	Q	76.10 ± 0.03	88.45 ± 0.01	79.04 ± 0.03	69.17 ± 11.26
	BE-type	Q	30.18 ± 0.07	73.61 ± 0.01	1.80 ± 0.01	0.76 ± 1.00
		Q'	69.88 ± 0.07	70.00 ± 0.02	98.28 ± 0.01	99.12 ± 1.02
	BA-type	Q	71.61 ± 0.02	85.43 ± 0.01	38.00 ± 0.08	74.24 ± 5.12
		Q'	28.34 ± 0.02	54.44 ± 0.00	62.15 ± 0.07	25.90 ± 4.93
BA-type (N.C.)	-	98.51 ± 0.00	96.87 ± 0.00	93.51 ± 0.03	98.46 ± 1.04	

Table 16: Statistics of L2 distance values between aligned attention matrices of BA-type CLAVI questions, averaged over all heads and layers for FrozenBiLM. We report the statistics separately for correctly and incorrectly answered consistent counterfactual predictions.

Type	Consistent Prediction	Mean	Variance
Video Counterfactual	Correct	0.70	0.02
	Incorrect	0.50	0.03
Text Counterfactual	Correct	0.55	0.01
	Incorrect	0.38	0.02

A.3.7 REPRESENTATION SENSITIVITY ANALYSIS

CLAVI can be used for diverse analyses to understand and interpret the joint multimodal representations in VideoQA models. We present one such analysis here. We want to find out the difference in representations between correctly and wrongly-answered counterfactual pairs. Ideally, the counterfactual pairs should have distinctly dissimilar representations to be answered correctly.

We use L2 norm as the distance metric. For CLAVI, we construct counterfactuals by augmenting the sequence of the frames (video counterfactuals) or replacing before/after and beginning/end (text counterfactual). Hence, we cannot directly compute the distance between the attention matrices of the counterfactuals because they contain different tokens (text counterfactual) or different order of the same tokens (video counterfactual). We solve this by finding token correspondence between the counterfactual pairs for each layer and head. By treating each attention matrix as a graph, we model the matrix alignment problem to finding the node correspondence between two isomorphic weighted directed complete graphs. Node correspondence between two graphs can be viewed as an instance of a linear sum assignment problem. That is, we want to learn a permutation transformation so that the two attention matrices are similar. We define similarity as negative of L2 distance. We solve this using modified Jonker-Volgenant algorithm as described by Crouse (2016).

We plot the histogram of L2 distance (averaged over heads and layers) for BA-type video and question counterfactuals in Figure 8. As expected, we find that if the answer is correct, then the average L2 distance is generally higher (skewed towards right). The mean and variance values L2 mean distribution of the correctly and incorrectly answered counterfactual pairs is summarized in Table 16. We find that the correct predictions have higher mean and lower variance than the incorrectly inferred counterfactual pairs. These findings validate that the model is indeed learning joint multimodal representations rather than creating its illusion.

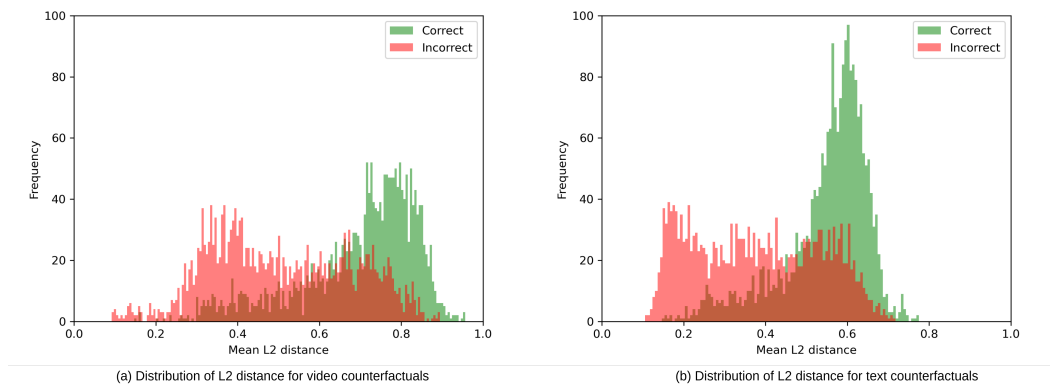


Figure 8: Histogram plots of mean l2 distance between counterfactual BA-type pairs for (a) video and (b) text counterfactuals for FrozenBiLM predictions (note that green is consistently correct; that is both the pairs in the counterfactuals are correctly answered, Similarly, red is consistently incorrect; that is at least one of the instance from the counterfactual pair is incorrectly answered).

A.3.8 DATASHEET

In this section we provide a more detailed documentation of the dataset with the intended uses. We base ourselves on the datasheet proposed by Gebu et al. (2021)

Motivation

- **For what purpose was the dataset created?** CLAVI is curated to diagnose and benchmark the joint multimodal understanding in VideoQA models. It uses temporal counterfactuals in video and question domains to assess the impact of multimodal shortcuts that can create an illusion of joint temporal understanding.
- **Who created the dataset and on behalf of which entity?** ANONYMOUS
- **Who funded the creation of the dataset?** ANONYMOUS

Composition

- **What do the instances that comprise the dataset represent?** Each instance in CLAVI comprises of a video, question, question type, and answer (“yes” or “no”).
- **How many instances are there in total?** CLAVI consists of 6,018 videos composing of 3,830 training and 2,188 testing videos. Each video is associated with 19 question-answer pairs, hence 114,342 data-points (72,770 training and 41,572 testing).
- **Does the dataset contain all possible instances or is it a sample (not necessarily random) of instances from a larger set?** The videos and question answer pairs in CLAVI are generated by manipulating real-world real-world videos. In theory, we can generate more instances with more videos with temporal annotations. We will release the code to generate the true and counterfactual video and question instances.
- **What data does each instance consist of?** Each instance in CLAVI comprises of a video, question, question type, and answer (“yes” or “no”).
- **Is there a label or target associated with each instance?** Yes, each question is associated with a ‘type’ label depending on the type of the question (types described in the main manuscript).
- **Is any information missing from individual instances?** No, all the instances have complete information the corresponding attributes.
- **Are relationships between individual instances made explicit?** Yes, the *video_name* attribute if of the form XXXXXXXX_1 for the original video segment and XXXXXXXX_2

for the counterfactual video segment, where XXXXXXXX is a unique 8-digit video id. The relationship between counterfactual questions is tabulated in the README file of the dataset.

- **Are there recommended data splits (e.g., training, development/validation, testing)?** We provide the split files which are curated from the original split files of Charades.
- **Are there any errors, sources of noise, or redundancies in the dataset?** No. The owners of Charades do not report any known errors. And since our data is generated by machine, we do not expect any errors. For unforeseen errors in temporal annotation boundaries in the original dataset, we eliminate it by selecting the segments where the actions of interest do not occur in the immediate neighbourhood (detailed in the main manuscript).
- **Is the dataset self-contained, or does it link to or otherwise rely on external resources?** This dataset provides video IDs from the Charades dataset under their Non-Commercial license.
- **Does the dataset contain data that might be considered confidential?** No. We curate our dataset from publicly and non-commercially available Charades dataset.
- **Does the dataset identify any sub-populations (e.g., by age, gender)?** No. While it is possible to identify gender from Charades temporal captions, we do not use it in the curation of CLAVI. We only use neutral pronoun *someone*.
- **Is it possible to identify individuals (i.e., one or more natural persons), either directly or indirectly (i.e., in combination with other data) from the dataset?** No. The owners of Charades dataset have anonymized the subject information.
- **Does the dataset contain data that might be considered sensitive in any way?** No. The owners of Charades dataset ensure this and we curate CLAVI from Charades.

Collection Process

- **How was the data associated with each instance acquired?** Each sample of CLAVI associates with a question, answer (yes/no) and video-id from Charades dataset. We generated instances in CLAVI that with corresponding video, question, answer and their respective counterfactuals from Charades.
- **What mechanisms or procedures were used to collect the data (e.g., hardware apparatuses or sensors, manual human curation, software programs, software APIs)?** We design a template-based VideoQA generation process to generate data each instance from Charades.
- **Who was involved in the data collection process (e.g., students, crowdworkers, contractors) and how were they compensated (e.g., how much were crowdworkers paid)?** Not applicable
- **If the dataset is a sample from a larger set, what was the sampling strategy (e.g., deterministic, probabilistic with specific sampling probabilities)?** Yes. We have outlined the process of filtering the data in detail in the appendix.
- **Who was involved in the data collection process (e.g., students, crowdworkers, contractors) and how were they compensated (e.g., how much were crowdworkers paid)?** Not applicable.
- **Over what timeframe was the data collected?** Our dataset is generated from Charades. We generate the dataset from February 2023 to June 2023.
- **Were any ethical review processes conducted (e.g., by an institutional review board)?** Not Applicable.
- **Does the dataset relate to people?** Yes. The Charades dataset contains videos of humans performing actions and we use it to curate CLAVI under their non-commercial license. However, we do not use the information pertaining to humans anyway.
- **Did you collect the data from the individuals in question directly, or obtain it via third parties or other sources (e.g., websites)?** No. Our video data is from Charades under their Non-Commercial license. Charades Homepage: <https://prior.allenai.org/projects/charades>.

- **Were the individuals in question notified about the data collection?** Not applicable. We curate our dataset from Charades and the original owners have ensured this.
- **Did the individuals in question consent to the collection and use of their data?** Not applicable. We curate our dataset from Charades and the original owners have ensured this.
- **If consent was obtained, were the consenting individuals provided with a mechanism to revoke their consent in the future or for certain uses?** Not applicable.

Preprocessing/cleaning/labeling

- **Was any preprocessing/cleaning/labeling of the data done (e.g., discretization or bucketing, tokenization, part-of-speech tagging, SIFT feature extraction, removal of instances, processing of missing values) No.**
- **Was the “raw” data saved in addition to the preprocessed/cleaned/labeled data (e.g., to support unanticipated future uses)?** No. The raw data (Charades) is distributed under Non-Commerical license.
- **Is the software that was used to preprocess/clean/label the data available?** Not applicable.

Uses

- **Has the dataset been used for any tasks already?** Video Question Answering.
- **Is there a repository that links to any or all papers or systems that use the dataset?** No.
- **What (other) tasks could the dataset be used for?** Video-Text and Text-Video retrieval.
- **Is there anything about the composition of the dataset or the way it was collected and preprocessed/cleaned/labeled that might impact future uses?** No.
- **Are there tasks for which the dataset should not be used?** No.
- **Will the dataset be distributed to third parties outside of the entity (e.g., company, institution, organization) on behalf of which the dataset was created?** CLAVI is an academic dataset for public non-commercial use.
- **How will the dataset be distributed (e.g., tarball on website, API, GitHub)?** The dataset files will be released on GitHub.
- **When will the dataset be distributed?** Latest by the official paper acceptance.
- **Will the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)?** Yes. The dataset will be released under GPL3.0 license and terms of usage will be outlined on the dataset hosting website along with the license and the required scripts.
Have any third parties imposed IP-based or other restrictions on the data associated with the instances? No.
Do any export controls or other regulatory restrictions apply to the dataset or to individual instances? No.

Maintenance

- **Who will be supporting/hosting/maintaining the dataset?** ANONYMOUS
How can the owner/curator/manager of the dataset be contacted (e.g., email address)? ANONYMOUS
- **Is there an erratum?** Not yet.
- **Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)?** Updates, if any, will be clearly mentioned on GitHub.
- **If the dataset relates to people, are there applicable limits on the retention of the data associated with the instances (e.g., were the individuals in question told that their data would be retained for a fixed period of time and then deleted)?** No.

- **Will older versions of the dataset continue to be supported/hosted/maintained?** Yes.
- **If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so?** Yes, we will provide the necessary code files with the dataset.

A.4 LIMITATIONS AND FUTURE WORK

Our dataset is intentionally simple, so as to focus the benchmark only on simple temporal sequence understanding, which preempts spatio-temporal referential understanding. We plan to include more complex temporal organizations of action classes like containment and partial-overlap that are defined using prepositions like *during* and *while* in future work. As the current state-of-the-art models catch-up to our benchmark, our future plan is to curate a more complex dataset with more natural questions that include temporal referring expressions with similar balanced doubly-negative strategy.