

Rethinking Video-Text Understanding: Retrieval from Counterfactually Augmented Data

Wufei Ma^{1,2*}, Kai Li¹, Zhongshi Jiang¹, Moustafa Meshry¹, Qihao Liu²
Huiyu Wang³, Christian Häne¹, and Alan Yuille²

¹ Meta Reality Labs

² Johns Hopkins University

³ Meta AI

Abstract. Recent video-text foundation models have demonstrated strong performance on a wide variety of downstream video understanding tasks. Can these video-text models genuinely understand the contents of natural videos? Standard video-text evaluations could be misleading as many questions can be inferred merely from the objects and contexts in a single frame or biases inherent in the datasets. In this paper, we aim to better assess the capabilities of current video-text models and understand their limitations. We propose a novel evaluation task for video-text understanding, namely *retrieval from counterfactually augmented data* (RCAD), and a new *Feint6K* dataset. To succeed on our new evaluation task, models must derive a comprehensive understanding of the video from cross-frame reasoning. Analyses show that previous video-text foundation models can be easily fooled by counterfactually augmented data and are far behind human-level performance. In order to narrow the gap between video-text models and human performance on RCAD, we identify a key limitation of current contrastive approaches on video-text data and introduce *LLM-teacher*, a more effective approach to learn action semantics by leveraging knowledge obtained from a pre-trained large language model. Experiments and analyses show that our approach successfully learn more discriminative action embeddings and improves results on Feint6K when applied to multiple video-text models. Our Feint6K dataset and project page is available [here](#).

Keywords: Video-Text Understanding · Retrieval

1 Introduction

Video-text foundation models have gained increasing attention due to their simple formulation and strong transferability [12, 16, 23]. By pretraining on web scale video-text datasets, these models demonstrate strong performance across a wide range of downstream tasks, such as video-text retrieval [22, 28] and video question answering [25].

* All data collection and experiments were conducted at JHU.

As video-text foundation models evolve and achieve increasingly better performance on various benchmarks, we raise the following questions: Can these video-text model truly grasp the semantics of natural videos? Are these models genuinely reaching a level of understanding comparable to humans? Despite the remarkable achievements made in previous works, our study suggests that current video-text models still fall far behind human-level understanding.

We argue that existing prominent results on standard video-text tasks can be misleading as models largely exploit the shortcuts and biases inherent in the dataset. Many of the questions can be answered by objects or context extracted from a single frame without capturing cross-frame relations. For the examples in Fig. 2a, the video-text alignment can be easily inferred from shortcuts such as “cymbals” or “football”. Moreover, the models may utilize biases in the datasets, such as the spurious correlation between “outdoor” and “football”. Current evaluations of video-text understanding are compromised by shortcuts and biases, which obscure us from analyzing the limitations of current models. As we proceed from image-text understanding to video-text understanding, we should focus on more challenging semantics in the video domain that require cross-frame reasoning to solve, such as the interaction between person and object or the change of appearances over a sequence of frames (see Fig. 1a).

In this paper we propose a novel evaluation paradigm of video-text understanding, *i.e.*, *retrieval from counterfactually augmented data* (RCAD). As demonstrated in Fig. 2b, the goal is to retrieve the only *positive caption* with matched semantics among “hard” negative captions. *Negative captions* are counterfactually augmented so that video-text models must derive a holistic understanding of the video sequence for both objects and actions, in order to retrieve the correct caption. As a comparison, negative examples from standard video-to-text retrieval [22, 24, 28] are captions of different videos in the same dataset, often containing different object entities that are easy to distinguish.

We follow the previous human-in-the-loop system [9] and develop a benchmark dataset for our new RCAD task, *i.e.*, the *Feint6K* dataset. We test a wide range of public video-text methods with various pretraining strategies on this new benchmark. Two failure examples are visualized in Fig. 1a and quantitative results are summarized in Fig. 1b. Note how state-of-the-art video-text model, InternVideo [23], features a 87.9% rank-1 accuracy on standard video-to-text retrieval but drops by 32.8% (pretraining only) and 29.7% (finetuned) when tested on our benchmark with counterfactually augmented data. By establishing a human-level baseline on our benchmark, the InternVideo model surprisingly falls far behind human performance by 38.6%. Our findings sharply contrast to the prominent performance obtained on previous video-text benchmarks and question the common belief that latest video-text models can develop a fairly effective representation for texts and videos in existing datasets.

In light of our evaluation results on Feint6K, we identify a significant limitation of the widely adopted contrastive representation learning on video-text data, which is the issue of shortcut learning (see Sec. 4.1). To address this, we propose *LLM-teacher* that enables a more effective learning of action se-

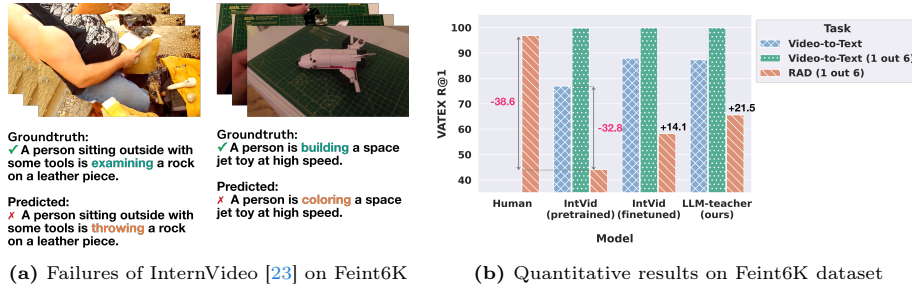


Fig. 1: (a): Although with large-scale pretraining on web-scale data, current video-text model like [23] can be easily fooled by counterfactually augmented data. (b): The performance of InternVideo on *retrieval from counterfactually augmented data* (RCAD) drops by over 30% when compared to the standard video-to-text retrieval and by 38.6% when compared to human-level performance. We also evaluate models on standard video-text retrieval from only 6 sampled candidates and show that our RCAD task is indeed much more challenging. Our LLM-teacher successfully improves the performance on RCAD by enforcing a more effective learning of action semantics.

mantics by introducing extra knowledge from pretrained large language models (LLMs). Specifically we generate synthetic captions by modifying the contents of an existing caption and LLM serves as the teacher to determine if a synthetic caption matches the semantics in the original caption using binary pseudo-labels or continuous logits. This approach enables a more effective learning of action embeddings from the available video-text data. Experimental results show that LLM-teacher learns more discriminative action embeddings from our advanced contrastive objectives and improves the performance of retrieval from counterfactually augmented data when applied to multiple video-text models.

In summary, our key contributions are as follows. (1) We propose to evaluate video-text models on questions that require cross-frame reasoning and develop a new task, retrieval from counterfactually augmented data, and a new dataset, Feint6K (Sec. 3). (2) Extensive experimental results on our new evaluation paradigm suggest that existing video-text models demonstrate very limited understanding of the action semantics in a video, which is contrary to the prominent performance achieved on standard video-text retrieval benchmarks. (3) From our results on Feint6K dataset, we identify a key shortcoming of contrastive learning approaches on video-text data. We present LLM-teacher to enforce a more effective learning of action embeddings by injecting knowledge from pretrained LLMs (Sec. 4). Our approach effectively improves the results on Feint6K when applied to multiple text-video models.

2 Related Works

Video-text pretraining. With the availability of web-scale video-text paired datasets, such as WebVid2M [1] and HowTo100M [17], recent developments of

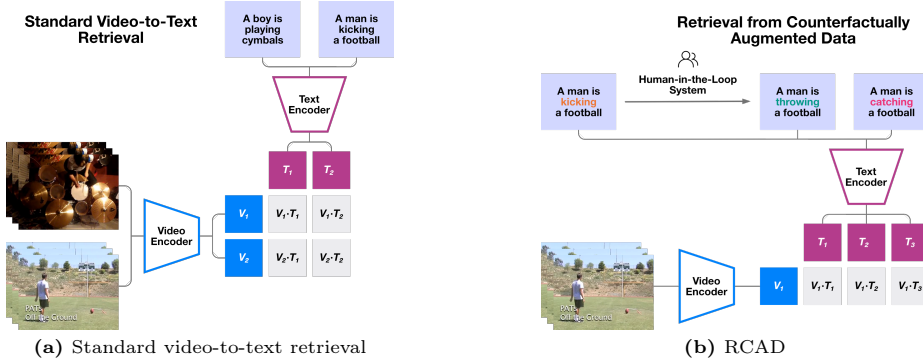


Fig. 2: Different evaluations of video-text understanding. (a): In standard video-to-text retrieval, negative captions are sampled from different videos in the same dataset. Image-text models can achieve good performance by exploiting shortcuts (e.g., “football” and “cymbals”) and biases (e.g., spurious correlation between “outdoor” and “football”). (b): In our proposed RCAD paradigm, we adopt a human-in-the-loop system (see Sec. 3.2) to obtain “hard” negative captions with unchanged object entities but modified actions. Models must develop a holistic understanding of the semantics from the sequence of frames to retrieve the matched caption.

video-text pretraining models achieved improved results on a wide range of video understanding tasks. Visual and textual embeddings jointly learned from the large-scale data demonstrate strong transferability and largely reduce the efforts in downstream tasks. Mainstream pretraining objectives can be categorized into discriminative and generative. Discriminative approaches extend the objective of CLIP [18] to the video-text domain and learn multi-modal representations by contrasting between matched and unmatched pairs [15, 16, 23, 26] or simply predicting if a video-text pair is matched [16]. Generative methods follow the masked modeling idea in BERT [5] and adopted masked language modeling (MLM) and masked video modeling (MVM) for video-text pretraining [16, 21, 23]. These generative approaches are often considered superior for action recognition, while discriminative approaches learn better semantics from texts [21].

Evaluation of video-text models. In order to analyze the effectiveness of video-text pretraining, previous works [15, 16, 21, 23, 26, 29] focused on the following tasks: (1) *video-text retrieval (zero-shot or finetuned)* [22, 28]: retrieving the matched video (or text) given a text (or video) as query where the negative candidates are unmatched pairs from the same dataset, (2) *video question answering (finetuned)* [25]: predicting the answer with a classification head, (3) *video classification (zero-shot or finetuned)* [6, 10, 20]: classifying videos using label names as text prompts, and (4) *video captioning (zero-shot or finetuned)* [11, 28]: summarizing the contents of a video.

We argue that results on these video-text tasks could be deceitful as most tasks heavily rely on the alignment of object entities in the video and text

(which image-text models are also capable of), and “action understanding” can often be developed through shortcuts and biases. For instance, a commonly-used evaluation for “action recognition” is video classification on ActivityNet [6], where a model classifies a given video into classes such as “playing badminton”, “kayaking”, or “volleyball”. However, video-text models would classify the videos by exploiting objects and contexts as shortcuts (*e.g.*, “badminton” and “kayak”), as well as other biases (*e.g.*, spurious correlation between “kayak” and “on the water”) without genuinely understanding the semantics of the action represented by interactions between the person and the objects over time. For instance, classifying the action as “kayaking” is essentially a “kayak” detection problem in ActivityNet. Therefore we propose *retrieval from counterfactually augmented data*, a new evaluation paradigm where we aim to eliminate the shortcuts from the questions so models must establish a comprehensive understanding of the semantics from cross-frame reasoning in order to predict the correct answer.

3 Retrieval from Counterfactually Augmented Data

Previous evaluation tasks of video-text models focused on the alignment of feature embeddings of the video-text pairs (*e.g.*, video-text retrieval [22, 28] and video classification [10, 20]) or reconstruction of the semantics with text (*e.g.*, video question answering [25] and video captioning [11, 28]). However, these evaluation tasks are largely limited by the paired data available in existing video-text datasets. As demonstrated by in Fig. 2a, we can extract most of the semantics of a video, such as the objects, people, and contexts in the video, by looking at only one frame from the video. The action in the video can be further inferred by exploiting biases in the datasets. Therefore, models pretrained on image-text data often perform surprisingly well on existing video-text tasks.

To effectively evaluate how video-text models can understand the contents and semantics of videos beyond images, we must develop video understanding tasks that are free from shortcuts and require cross-frame reasoning to solve. To this end, we propose a new evaluation task for video-text models, *i.e.*, retrieval from counterfactually augmented data (Sec. 3.1), and a new dataset Feint6K (Sec. 3.2). In contrast to previous datasets with matched video-text pairs scrapped from the web [24, 28] or annotated by human [22], we adopt a human-in-the-loop system and annotate counterfactually manipulated texts as negative pairs with the original video. This allows us to evaluate video-text models with novel tasks and better understand their limitations. Lastly we measure human performance on our dataset as a direct comparison with the state-of-the-art video-text models (Sec. 3.3).

3.1 Task Formulation

Retrieval from counterfactually augmented data is a variant of the standard video-to-text retrieval. As shown in Fig. 2b, given a video sequence and a list of candidate captions, the model is asked to retrieve the caption that best matches

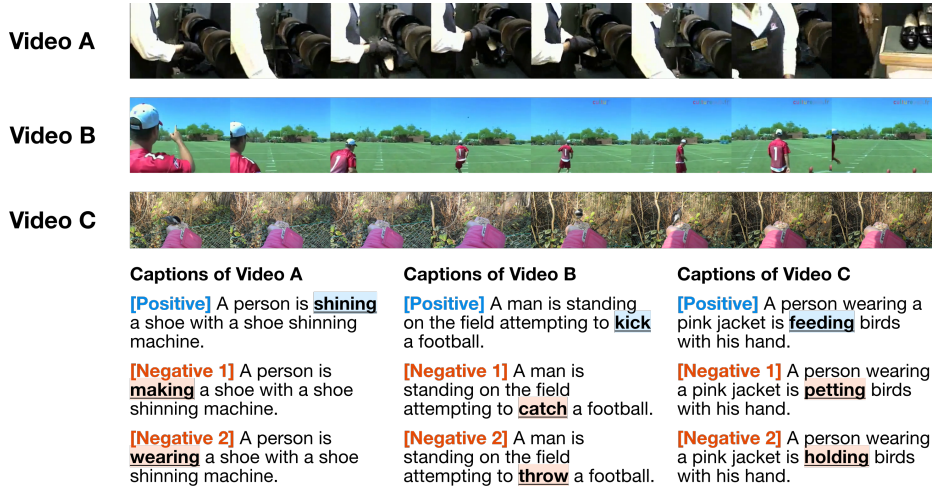


Fig. 3: Examples of *RCAD* on our Feint6K dataset.

the semantics in the video. However, unlike standard video-to-text retrieval where negative captions come from other video-text pairs in the same dataset, negative captions in our task are modified from the positive captions, with the same text structure and object entities but different actions (see Fig. 3).

Similar to video-text retrieval, our proposed task provides the benefit of zero-shot evaluation. This allows us to quantitatively analyze the effectiveness of video-text pretraining without downstream finetuning.

What distinguishes our task from all previous video-text tasks is its focus on questions that require cross-frame reasoning. Consider “Video B” in Fig. 3 as an example. All captions contain the same object entities present in the video, *i.e.*, “a man”, “the field”, and “a football”, but differ in the action specified. We cannot determine the most accurate caption because all suggested actions, “kick”, “catch”, and “throw”, are plausible given the video’s context. To retrieve the corresponding caption, it is crucial to grasp the action’s semantics by inspecting the interactions between the man and the football across a series of frames. This underscores the necessity of understanding the video semantics over time in order to succeed in our task, rather than relying shortcuts or biases.

3.2 Feint6K: Data Collection

We utilize a human-in-the-loop system [9] to counterfactually manipulate the positive captions in existing video-text datasets [22, 28]. The generated counterfactually augmented data, when paired with the original videos, forms negative pairs to be used in our proposed task. We recruit 40 annotators to manually manipulate the actions in existing captions. Specifically, the new actions introduced must be plausible within the context in the caption but are not occurring in the corresponding video. In summary, a total of 6,243 videos from the validation




Fig. 4: Overview of our data collection pipeline for Feint6K dataset featuring a human-in-the-loop system.

set of the MSR-VTT dataset [28] and the test set of the VATEX dataset [22]. For detailed statistics of our Feint6K dataset, please refer to our supplementary materials.

The overview of our human-in-the-loop system is demonstrated in Fig. 4. We start by sampling a small subset of video-text pairs and manually annotate the counterfactually augmented captions for the purpose of training and demonstration. Specific guidance and feedback are given to the annotators on the practice questions during the training stage. Moreover, we adopt a reviewing process during the annotation stage to further ensure the quality of our annotated data. Each annotation is reviewed before acceptance, where rejected annotations are marked and sent back for refinement. We refer the readers to the supplementary materials regarding the annotation interface and annotator guidance.

3.3 Human Performance

We employ the same group of annotators to establish a human-level baseline on our Feint6K dataset. Analyzing the human performance serves two goals: (i) verifying the legitimacy of our Feint6K dataset – if each question is answerable and has one unique answer, and (ii) comparing the state-of-the-art video-text models with human-level performance.

Each annotator is first presented with the full video and then asked to select the one out of six captions that best matches the semantics in the video. To avoid leakage, we make sure the annotators assigned to the testing questions are not the same annotators annotating this video. We also randomly shuffle the order of the candidate captions to remove any biases. In Sec. 5.2 we present the quantitative comparisons between human-level performance and current state-of-the-art video-text models on our Feint6K dataset.

4 LLM-teacher

Contrastive learning are widely adopted in previous video-text representation learning works [16, 23, 26]. The contrastive models learn to distinguish between

similar and dissimilar pairs of video-text data by contrasting their feature representations. Despite the simple formulation and great generalization ability, our results on Feint6K show that current video-text models built heavily on contrastive objectives have very limited understanding of action semantics in natural videos and can be easily fooled by counterfactually augmented data. Therefore we aim to develop a more effective contrastive approach and learn more discriminative action representations from existing video-text data.

Sec. 4.1 discusses the limitations of current contrastive approaches due to numerous shortcuts in video-text dataset. Inspired by this finding, we present LLM-teacher in Sec. 4.2, a simple but effective approach to learn more discriminative action embeddings by incorporating knowledge from pretrained large language models.

4.1 Shortcuts in Video-Text Data


The CLIP model [18] trained on image-text data with a contrastive objective demonstrated strong zero-shot capabilities in many vision tasks [13, 30] and was widely adopted as the vision encoder in many large vision-language models [14]. However, directly generalizing this idea to video-text data may face unprecedented challenges due to shortcuts in the video-text data.

As we proceed from image-text to video-text representation learning, we aim to learn a powerful embedding for not only the object entities, but also the actions specified by texts or sequence of frames. Naturally we hope to achieve this by contrasting videos with different actions. However, as we are optimizing the contrastive objectives in a mini-batch of videos, object entities become the shortcuts that contrastive models exploit to saturate the contrastive objective, leading to an ineffective action embeddings. Consider the example in Fig. 2a, the contrastive loss would be small as long as the model learns discriminative embeddings for shortcut objects such as “cymbal” and “football”, even without any understanding of the actions. Since many video-text models start with pretrained CLIP encoders [23, 26], they often possess strong embeddings for shortcut objects from the beginning of the video-text pretraining, which hinders the model from further learning effective action representations.

In Fig. 5 we present quantitative results of changes in cosine similarities when the video is unchanged but the object or action in the caption is manipulated. We find that state-of-the-art video-text models learn less discriminative embeddings for actions as compared to objects, which also explains the small gap between image-text and video-text models on standard video-text tasks.

4.2 LLM-teacher

We introduce LLM-teacher, an LLM-powered approach to learn better action representations with contrastive objectives. Specifically, LLM serves as the “teacher” and provides synthesized captions as extra knowledge for the video-text model to learn from. These new captions serve as negative captions of the original video,



(a) Comparison between Δs w.r.t. object or action changes. (b) Comparison between Δs w.r.t. action using InternVideo and our LLM-teacher.

Fig. 5: Change of cosine similarity w.r.t. objects or actions. (a): Comparison between the change in cosine similarity when the action or object is swapped. Results show that current video-text models learn a more effective embedding for objects than for actions. **(b):** Comparison between the change in cosine similarity using InternVideo or our LLM-teacher. This demonstrate that LLM-teacher learns a more discriminative embedding for actions by enabling a more effective contrastive learning using knowledge from LLMs. Refer to Sec. 5.2 for more details.

removing possible shortcuts and biases and enabling a more effective learning of action semantics.

We start with positive captions from common video-text datasets and run an abstract meaning representation (AMR) parser, resulting in a list of action or object tokens. To generate the hard negative captions, two methods are considered to find novel actions or objects given the original caption and the substitution token w .

Method I: Mask filling. Masked language modeling (MLM) [5] is a widely adopted self-supervised task for LLMs. In the pretraining stage, a certain percentage of word tokens are substituted with a special token [mask] and the LLM is optimized to recover the masked tokens. To find negative captions with novel actions given the original caption c_0 and substitution token a , we first substitute the token a with [mask] and use a LLM pretrained with MLM to predict a list of k possible tokens $\{\hat{w}_i\}_{i=1}^k$ such that $\hat{w}_i \neq w$ and \hat{w}_i is an action token.

Method II: LLM-powered chatbot. Although mask filling can find appropriate words to substitute w , it is limited by one single token. In many cases, we want to update the prepositions following the change of verbs. For example, when “install” is substituted by “uninstall”, we should also update the preposition from “install ... to” to “uninstall ... from”. Here we consider a more flexible approach that builds on LLM-powered chatbots. Besides a prompt describing the text substitution, we leverage the in-context learning ability of LLMs [2, 7] and provide multiple in-context examples to obtain the desirable negative captions.

LLM-teacher. With the procedure above, we obtain k synthesized negative captions $\{n_i\}_{i=1}^k$ besides the original positive caption p for each video v in existing datasets [22, 28]. For each of the original and generated captions, LLM serves as a “teacher” and provide binary pseudo-labels – if the caption have matched (“1”) or unmatched (“0”) semantics with the video. Now we may optimize the video-text model with the normalized temperature-scaled cross entropy loss

$$l = -\log \frac{\exp(\text{sim}(f_v, f_p)/\tau)}{\exp(\text{sim}(f_v, f_p)/\tau) + \sum_{i=1}^k \exp(\text{sim}(f_v, f_{n_i})/\tau)} \quad (1)$$

where f_v, f_p, f_{n_i} are the visual and textual embeddings.

In practice, we find that certain negative captions have unmatched but similar semantics to the video. It is undesirable to regard them as a strictly negative pair in the contrastive loss. Therefore, we extend the binary pseudo-labels to soft logits by computing caption similarities with a pretrained LLM and then optimize the video-text model to match model outputs with the soft logits from the LLM teacher model [8].

$$\begin{aligned} l &= \mathcal{L}_{\text{KL}}(z_{\text{video-text}}, z_{\text{LLM}}) \\ z_{\text{video-text}, t} &= \frac{\exp(\text{sim}(f_v, f_t)/\tau)}{\exp(\text{sim}(f_v, f_p)/\tau) + \sum_{i=1}^k \exp(\text{sim}(f_v, f_{n_i})/\tau)} \\ z_{\text{LLM}, t} &= \frac{\exp(\text{sim}(e_p, e_t)/\tau)}{\exp(\text{sim}(e_p, e_p)/\tau) + \sum_{i=1}^k \exp(\text{sim}(e_p, e_{n_i})/\tau)} \end{aligned} \quad (2)$$

where $t \in \{n_i\}_{i=1}^k \cup \{p\}$, f ’s are textual and visual embeddings of the video-text model, and e ’s are textual embeddings of a pretrained LLM.

5 Experiments

5.1 Experimental Setup

Datasets. We use MSR-VTT dataset [28] and VATEX dataset [22] for standard video-to-text retrieval and our Feint6K dataset for retrieval from counterfactually augmented data. Note that Feint6K dataset contains the same video sequences for evaluation as in MSR-VTT dataset and VATEX dataset.

Evaluation metrics. For standard video-to-text retrieval we report rank-1 accuracy (R@1) as this is the most challenging metric and distinguish the performance of previous video-text models. For retrieval from counterfactually-augmented data we use rank-1 accuracy (R@1), rank-2 accuracy (R@2), and mean rank (MeanR).

Video-text models. We consider multiple state-of-the-art public models in this work. CLIP4Clip [15] extends CLIP model [18] to the video domain using a

Model	MSR-VTT	Feint6K (MSR-VTT)			VATEX	Feint6K (VATEX)		
	R@1	R@1	R@2	MeanR	R@1	R@1	R@2	MeanR
Human		95.2				96.8		
Random	<1e-3	16.7			<1e-3	16.7		
Zero-Shot								
CLIP [18]	26.3	37.3	55.3	2.6	38.8	34.8	54.3	2.7
VideoCLIP [27]	14.5	35.1	71.3	2.6	13.7	33.0	70.7	2.7
InternVideo [23]	37.5	45.8	63.6	2.3	76.9	44.1	63.9	2.3
LanguageBind [31]	42.8	41.3	77.0	2.4		42.3	77.0	2.4
Finetuned								
LanguageBind [31]								
CLIP4Clip [15]	43.1	50.8	72.4	2.0				
VindLU [3]	46.6	53.4	70.9	2.0				
SimVTP [16]	50.2	35.7	70.8	2.6	76.6	33.6	68.4	2.6
w/ LLM-teacher-lbl	49.0 $\downarrow 1.2$	40.0 $\uparrow 4.3$	73.1 $\uparrow 2.3$	2.3 $\downarrow 0.3$	74.8 $\downarrow 1.8$	37.3 $\uparrow 3.7$	72.2 $\uparrow 3.8$	2.4 $\downarrow 0.2$
w/ LLM-teacher-lgt	49.5 $\downarrow 0.7$	43.5 $\uparrow 7.8$	75.0 $\uparrow 4.2$	2.2 $\downarrow 0.4$	75.3 $\downarrow 1.3$	40.1 $\uparrow 6.5$	73.5 $\uparrow 5.1$	2.3 $\downarrow 0.3$
InternVideo [23]	49.1	58.6	80.2	1.8	87.9	58.2	76.9	1.9
w/ LLM-teacher-lbl	48.2 $\downarrow 0.9$	64.2 $\uparrow 5.6$	82.5 $\uparrow 2.3$	1.7 $\downarrow 0.1$	85.2 $\downarrow 2.7$	63.8 $\uparrow 5.6$	80.5 $\uparrow 3.6$	1.7 $\downarrow 0.2$
w/ LLM-teacher-lgt	48.9 $\downarrow 0.2$	65.8 $\uparrow 7.2$	83.8 $\uparrow 3.6$	1.7 $\downarrow 0.1$	87.3 $\downarrow 0.6$	65.6 $\uparrow 7.4$	81.7 $\uparrow 4.8$	1.7 $\downarrow 0.2$

Table 1: Performance of standard video-to-text retrieval on MSR-VTT dataset [28] and VATEX dataset [22], as compared to performance of retrieval from counterfactually-augmented data on our Feint6K dataset. In the “zero-shot” setting, models are evaluated directly after the pretraining stage, while in the “finetuned” setting, models are finetuned on the training set of MSR-VTT or VATEX. With LLM-teacher, we enable a more efficient learning of action semantics and effectively improves R@1 accuracy on Feint6K. Here “LLM-teacher-lbl” stands for our approach with binary pseudo-labels and “LLM-teacher-lgt” uses soft logits.

mean pooling mechanism for zero-shot video-text retrieval and a contrastive loss for finetuning. SimVTP [16] is pretrained on the WebVid2M dataset [1] with a combination of contrastive learning and masked modeling. InternVideo [23] is a video-text foundation model pretrained on a combination of 7 large-scale video-text datasets, with a supervised video post-pretraining for better video recognition. LanguageBind [31] is a language-based multi-modal pretraining model with remarkable performance on a wide range of benchmarks.

To demonstrate the efficacy of our LLM-teacher, we adopt this method on two state-of-the-art video-text models, SimVTP and InternVideo. In the default setting we generate 10 action-based captions with LLM for each training video. For text generation we use a pretrained XLM-RoBERTa [4] for mask filling. To measure text similarities and compute soft logits, we utilize a pretrained Sentence-BERT model [19]. In Sec. 5.3 we ablate on the choices of parameters and caption generation.

5.2 Main Results

Performance of previous state-of-the-art. In Tab. 1 we report the quantitative results of standard video-text retrieval on MSR-VTT and VATEX, and of

retrieval from counterfactually augmented data on our Feint6K. We also report a “Random” performance where a model predicts random guesses. We make the following observations: (i) Results show that previous state-of-the-art video-text models demonstrate limited understanding of the action semantics in a video with less than 60% R@1 accuracy on RCAD, given a 16.7% R@1 accuracy when taking random guesses. (ii) Although it is hard to compare the quantitative results between two tasks, we note that our RCAD is much more challenging given the performance gap between “Random” and previous state-of-the-arts. This supports our previous arguments that our RCAD removes shortcuts from standard video-text datasets and focuses on harder questions in the video domain that require cross-frame reasoning.

Performance of LLM-teacher. We apply our LLM-teacher to two pretrained video-text models, SimVTP [16] and InternVideo [23]. Specifically, we consider the two objectives in Eq. (1) and Eq. (2), labeled as “LLM-teacher-lbl” and “LLM-teacher-lgt” respectively. Results in Tab. 1 show that by exploiting knowledge from pretrained LLMs, LLM-teacher enables a more efficient leaning of action semantics and achieved improved results on various metrics of RCAD, increasing the R@1 accuracy on RCAD by 7.2% and 7.4%. We also note that with binary pseudo-labels, results demonstrate a trade-off between the performance on standard video-to-text retrieval and RCAD; and in comparison, LLM-teacher with soft logits achieves the highest performance on RCAD (65.8% and 65.6%) with negligible drops (0.2% and 0.6%) on standard retrieval. This is because soft logits account for the text similarities between different actions and avoid overfitting on binary pseudo-labels.

Comparison with human-level performance. The human-level performance are reported in Tab. 1. We see that human annotators achieved an almost perfect performance on Feint6K dataset, with a 95.2% accuracy on MSR-VTT videos and a 96.8% accuracy on VATEX videos. These results show that: (i) most questions for retrieval from counterfactually-augmented data are answerable with an unique answer, and (ii) this task is fairly simple for human. As a comparison, all previous video-text models with large-scale pretraining fall behind by a wide margin. Specifically, the state-of-the-art video-text model InternVideo [23] features a heavy two-stage pretraining on a collection of 7 large-scale video-text datasets but falls behind by 36.6% and 38.6% even with downstream finetuning.

Analyses of failure cases. In Fig. 1b we present two failure cases by the InternVideo [23] model on our Feint6K dataset. Both examples are trivial for human, yet challenging from video-text models due to incapability of cross-frame reasoning. For the first example, a model must analyze the interactions between the person and the rock over a sequence of frames to predict the correct answer. And in the second example, the model must observe the change of appearances over time. As we can see, with the counterfactually augmented data in our Feint6K,

models cannot “guess” the correct answer by exploiting shortcuts extracted from a single frame. The findings on Feint6K highlight significant weaknesses in current video-text models, offering valuable insights for future researches in this area. Additional qualitative examples of video-text models on Feint6K dataset are provided in the supplementary materials for readers’ reference.

Sensitivity to the change of object or action. With the counterfactually augmented data collected, we investigate how the cosine similarities between the video and text change when the object or action in the text are swapped with an object or action not present in the image. Given the original video-text pair (v, t) and a counterfactually-augmented text \hat{t} , the *change in cosine similarity* is given by $\Delta s = s(v, \hat{t}) - s(v, t)$. As \hat{t} contains objects or actions not present in the video, Δs should be negative for an ideal video-text model. Moreover, larger $|\Delta s|$ implies the model being more sensitive to the changes. We compute Δs using the InternVideo with pretraining only (“IntVid-pt”), InternVideo with downstream finetuning (“IntVid-ft”), and our LLM-teacher (“LLM-teacher”).

Results in Fig. 5a show that the Δs are almost always negative when the objects are swapped and the absolute changes are larger. In comparison, Δs are sometimes positive when actions are modified and the absolute changes are smaller. This demonstrates that InternVideo learns a more effective embedding for objects than for actions. This further highlights the necessity of our new evaluation paradigm for exploring the limitations of video-text models beyond existing tasks. Moreover, in Fig. 5b we note that with LLM-teacher the action embeddings are more effective than the ones learned by InternVideo, as shown by the more reasonable changes in cosine similarity when actions are swapped.

Please refer to the supplementary materials for a detailed discussion regarding the influence of textual encoders.

5.3 Ablation Study

We conduct ablation studies on the VATEX dataset [22] and our Feint6K dataset. We follow the same settings as above and report rank-1 (R@1) accuracy for standard video-to-text retrieval and rank-1 accuracy (R@1), rank-2 accuracy (R@2), and mean rank (MeanR) for counterfactual augmented data.

Choice of captions in LLM-teacher. Tab. 2 shows the comparison between different numbers and types of LLM captions generated. In the “default” setting we use 10 action-based LLM captions and compare to settings with 5 action-based LLM captions or 5 action-based and 5 object-based LLM captions. Results show with 5 action-based LLM captions the R@1 accuracy on Feint6K drops by 0.9% while the R@1 accuracy on VATEX increases by 0.3%. We also experiment on object-based LLM captions and find them not beneficial for both standard video-to-text retrieval or our RCAD. This is consistent with our assumption that pretrained video-text models already learn a discriminative embedding for objects and a more efficient training objective for action is needed.

Model	VATEX	Feint6K (VATEX)			
	R@1	R@1	R@2	MeanR	
Default (10 action captions; XLM-RoBERTa)	87.3	65.6	81.7	1.7	
5 action captions	87.6 $\uparrow 0.3$	64.7 $\downarrow 0.9$	81.0 $\downarrow 0.7$	1.7	0.0
5 object + 5 object captions	87.5 $\uparrow 0.2$	64.2 $\downarrow 1.4$	80.6 $\downarrow 1.1$	1.7	0.0
LLM Chatbot	87.0 $\downarrow 0.3$	65.9 $\uparrow 0.3$	81.8 $\uparrow 0.1$	1.7	0.0

Table 2: Ablation studies on the number of LLM captions and caption generation methods. In the default setting we use a total of 10 action-based LLM captions and utilize a pretrained XLM-RoBERTa [4] for mask filling.

Choice of caption generation. Besides a pretrained XLM-RoBERTa model for mask filling, we also experiment on generating captions with LLM-powered agents that are finetuned on chat datasets for dialogue applications. Specifically we use the “chat” model finetuned for dialogue applications. Empirically we find that captions generated by LLM chatbots achieves a higher overall quality, exploring a more diverse caption space and better flexibility. However we don’t observe significant improvements from our ablation study experiments – the R@1 accuracy on VATEX drops by 0.3% and the R@1 accuracy on Feint6K increases by 0.3%. We choose XLM-RoBERTa in our main experiments as the model runs faster and scale up easily to bigger settings. We provide qualitative comparisons between the two types of caption generation in our supplementary materials.

6 Conclusions

In this work we propose a new evaluation task, retrieval from counterfactually augmented data, and a benchmark dataset, Feint6K. The idea is to remove shortcuts in the video-text questions with a human-in-the-loop system and produce more challenging questions where the video-text model must derive a comprehensive understanding of the video with cross-frame reasoning. Quantitative and qualitative evaluation results on our Feint6K dataset show that despite the task is trivial for human, previous state-of-the-art video-text models can be easily fooled by the counterfactually augmented data. This implies that the prominent results on previous video-text benchmarks could be misleading – models may largely exploit shortcuts without genuinely understand the video contents. Moreover, we identify a key limitation of current contrastive learning on video-text data being the shortcut learning and propose LLM-teacher that enables a more effective learning of action semantics by utilizing knowledge from pre-trained LLMs. Experimental results and analyses show that our method can learn a more discriminative representation for actions.

Supplementary materials. We present the following: (1) limitations of our work, (2) more details about our Feint6K dataset, (3) ethics statements, and (4) extra quantitative and qualitative results.

Acknowledgements. We would like to thank Daniel Khashabi, Yiyan Li, and the anonymous reviewers for their helpful comments and suggestions. Wufei Ma is supported by ONR with N00014-23-1-2641.

References

1. Bain, M., Nagrani, A., Varol, G., Zisserman, A.: Frozen in time: A joint video and image encoder for end-to-end retrieval. In: IEEE International Conference on Computer Vision (2021) 3, 11
2. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. *Advances in neural information processing systems* **33**, 1877–1901 (2020) 9
3. Cheng, F., Wang, X., Lei, J., Crandall, D., Bansal, M., Bertasius, G.: Vindlu: A recipe for effective video-and-language pretraining. *arXiv preprint arXiv:2212.05051* (2022) 11
4. Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., Stoyanov, V.: Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116* (2019) 11, 14
5. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018) 4, 9
6. Fabian Caba Heilbron, Victor Escorcia, B.G., Niebles, J.C.: Activitynet: A large-scale video benchmark for human activity understanding. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 961–970 (2015) 4, 5
7. Gupta, T., Kembhavi, A.: Visual programming: Compositional visual reasoning without training. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 14953–14962 (2023) 9
8. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015) 10
9. Kaushik, D., Hovy, E., Lipton, Z.C.: Learning the difference that makes a difference with counterfactually-augmented data. *arXiv preprint arXiv:1909.12434* (2019) 2, 6
10. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950* (2017) 4, 5
11. Krishna, R., Hata, K., Ren, F., Fei-Fei, L., Carlos Niebles, J.: Dense-captioning events in videos. In: *Proceedings of the IEEE international conference on computer vision*. pp. 706–715 (2017) 4, 5
12. Li, K., Wang, Y., Li, Y., Wang, Y., He, Y., Wang, L., Qiao, Y.: Unmasked teacher: Towards training-efficient video foundation models. *arXiv preprint arXiv:2303.16058* (2023) 1
13. Lin, J., Gong, S.: Gridclip: One-stage object detection by grid-level clip representation learning. *arXiv preprint arXiv:2303.09252* (2023) 8
14. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. In: *NeurIPS* (2023) 8
15. Luo, H., Ji, L., Zhong, M., Chen, Y., Lei, W., Duan, N., Li, T.: Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neuro-computing* **508**, 293–304 (2022) 4, 10, 11

16. Ma, Y., Yang, T., Shan, Y., Li, X.: Simvtp: Simple video text pre-training with masked autoencoders. arXiv preprint arXiv:2212.03490 (2022) [1](#), [4](#), [7](#), [11](#), [12](#)
17. Miech, A., Zhukov, D., Alayrac, J.B., Tapaswi, M., Laptev, I., Sivic, J.: HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips. In: ICCV (2019) [3](#)
18. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021) [4](#), [8](#), [10](#), [11](#)
19. Reimers, N., Gurevych, I.: Sentence-bert: Sentence embeddings using siamese bert-networks. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing. Association for Computational Linguistics (11 2019), <https://arxiv.org/abs/1908.10084> [11](#)
20. Soomro, K., Zamir, A.R., Shah, M.: Ucf101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402 (2012) [4](#), [5](#)
21. Tong, Z., Song, Y., Wang, J., Wang, L.: Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. Advances in neural information processing systems **35**, 10078–10093 (2022) [4](#)
22. Wang, X., Wu, J., Chen, J., Li, L., Wang, Y.F., Wang, W.Y.: Vatex: A large-scale, high-quality multilingual dataset for video-and-language research. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4581–4591 (2019) [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [10](#), [11](#), [13](#)
23. Wang, Y., Li, K., Li, Y., He, Y., Huang, B., Zhao, Z., Zhang, H., Xu, J., Liu, Y., Wang, Z., et al.: Internvideo: General video foundation models via generative and discriminative learning. arXiv preprint arXiv:2212.03191 (2022) [1](#), [2](#), [3](#), [4](#), [7](#), [8](#), [11](#), [12](#)
24. Wu, Z., Yao, T., Fu, Y., Jiang, Y.G.: Deep learning for video classification and captioning. In: Frontiers of multimedia research, pp. 3–29 (2017) [2](#), [5](#)
25. Xu, D., Zhao, Z., Xiao, J., Wu, F., Zhang, H., He, X., Zhuang, Y.: Video question answering via gradually refined attention over appearance and motion. In: Proceedings of the 25th ACM international conference on Multimedia. pp. 1645–1653 (2017) [1](#), [4](#), [5](#)
26. Xu, H., Ghosh, G., Huang, P.Y., Okhonko, D., Aghajanyan, A., Metze, F., Zettlemoyer, L., Feichtenhofer, C.: Videoclip: Contrastive pre-training for zero-shot video-text understanding. arXiv preprint arXiv:2109.14084 (2021) [4](#), [7](#), [8](#)
27. Xu, H., Ghosh, G., Huang, P.Y., Okhonko, D., Aghajanyan, A., Metze, F., Zettlemoyer, L., Feichtenhofer, C.: VideoCLIP: Contrastive pre-training for zero-shot video-text understanding. In: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP). Association for Computational Linguistics, Online (Nov 2021) [11](#)
28. Xu, J., Mei, T., Yao, T., Rui, Y.: Msr-vtt: A large video description dataset for bridging video and language. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5288–5296 (2016) [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [10](#), [11](#)
29. Yan, S., Zhu, T., Wang, Z., Cao, Y., Zhang, M., Ghosh, S., Wu, Y., Yu, J.: Video-text modeling with zero-shot transfer from contrastive captioners. arXiv preprint arXiv:2212.04979 (2022) [4](#)
30. Zhou, Z., Lei, Y., Zhang, B., Liu, L., Liu, Y.: Zegclip: Towards adapting clip for zero-shot semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11175–11185 (2023) [8](#)

31. Zhu, B., Lin, B., Ning, M., Yan, Y., Cui, J., HongFa, W., Pang, Y., Jiang, W., Zhang, J., Li, Z., Zhang, C.W., Li, Z., Liu, W., Yuan, L.: Languagebind: Extending video-language pretraining to n-modality by language-based semantic alignment. In: The Twelfth International Conference on Learning Representations (2024), <https://openreview.net/forum?id=QmZKc7UZCy> 11