

MILES: Visual BERT Pre-training with Injected Language Semantics for Video-text Retrieval

Yuying Ge^{1*} Yixiao Ge² Xihui Liu⁴ Jinpeng Wang⁵
Jianping Wu⁶ Ying Shan² Xiaohu Qie³ Ping Luo¹
¹The University of Hong Kong ²ARC Lab, ³Tencent PCG ⁴UC Berkeley
⁵National University of Singapore ⁶Tsinghua University
yuyingge@hku.hk {yixiaoge, yingsshan, tigerqie}@tencent.com
xihui.liu@berkeley.edu jinpengwang@u.nus.edu
jianping@cernet.edu.cn pluo@cs.hku.hk

Abstract. Dominant pre-training work for video-text retrieval mainly adopt the “dual-encoder” architectures to enable efficient retrieval, where two separate encoders are used to contrast global video and text representations, but ignore detailed local semantics. The recent success of image BERT pre-training with masked visual modeling that promotes the learning of local visual context, motivates a possible solution to address the above limitation. In this work, we for the first time investigate masked visual modeling in video-text pre-training with the “dual-encoder” architecture. We perform Masked visual modeling with Injected Language Semantics (MILES) by employing an extra snapshot video encoder as an evolving “tokenizer” to produce reconstruction targets for masked video patch prediction. Given the corrupted video, the video encoder is trained to recover text-aligned features of the masked patches via reasoning with the visible regions along the spatial and temporal dimensions, which enhances the discriminativeness of local visual features and the fine-grained cross-modality alignment. Our method outperforms state-of-the-art methods for text-to-video retrieval on four datasets with both zero-shot and fine-tune evaluation protocols. Our approach also surpasses the baseline models significantly on zero-shot action recognition, which can be cast as video-to-text retrieval.

Keywords: Masked Visual Modeling, Video-text Retrieval

1 Introduction

Pre-training visual-language models to learn transferable representations for downstream video-text retrieval has attracted increasing attention in recent years. The “dual-encoder” architectures [6,31,51,18,53,41,35], where two individual encoders are used to contrast global video and text representations, become the most popular practice to enable efficient retrieval. Despite the high efficiency, recent studies [30] have shed light on the limitations of such dual-encoder representation learning: the coarse-grained alignment constraint on global video-text

* Work done during internship in ARC Lab, Tencent PCG.

features hinders the capture of detailed local semantics and the further improvements in video-text retrieval.




Fig. 1. (a) The diagram of Masked visual modeling with Injected Language Semantics (MILES) in video-text pre-training, which aims to reconstruct the masked video content to the targets that are aligned with the language semantics through reasoning with the spatial and temporal context of the visible video patches. (b) Comparison between recent “dual-encoder” methods for video-text pre-training in text-to-video retrieval on MSR-VTT and MSVD with both the zero-shot (ZS) and fine-tune (FT) evaluation protocols (R@1 as the metric).

Inspired by masked language modeling [12] in natural language processing, a pioneering work, BEIT [7], introduces the pretext task of masked visual modeling (MVM) to promote the learning of local visual context in image pre-training. A proportion of image patches are randomly masked, and the vision Transformer [14] is trained to recover the vision tokens that are obtained from a pre-learned image “tokenizer” [38] as the reconstruction targets. The great success of MVM in image pre-training provides a possible solution to encourage the learning of local visual semantics in video-text retrieval.

We would like to recall the key factor that makes MVM successful is the denoising training objective, where the design of the masked prediction targets turns out to be the most critical. The follow-up works [54,13] of BEIT confirm this by introducing better image tokenizers that are more aware of the high-level perceptions. So *what makes for good targets of masked visual prediction in video-text pre-training?* We argue that, towards the goal of accurate video-text retrieval, besides the spatial and temporal visual context understanding, the local alignment with language semantics serves as another important objective for the prediction of masked video patches as illustrated in Fig. 1 (a).

To build masked video prediction targets with injected language semantics, we use a snapshot video encoder as an evolving “tokenizer” to produce regressed targets for masked video patches. The snapshot video encoder aggregates the knowledge of the in-training video encoder in prior epochs, whose predictions gradually approach the text domain under the global video-text contrastive constraint. By imposing MVM regularizations towards the targets from the snapshot encoder, the in-training video encoder can be iteratively improved to capture more detailed video semantics that are locally aligned with the languages, which, in turn, further enhances the evolving “tokenizer”, *i.e.*, snapshot encoder. Our

method successfully applies the idea of MVM in video-text pre-training without extra pre-training stages for obtaining a proper cross-modality “tokenizer”.

Specifically, based on the “dual-encoder” architecture, we employ an extra snapshot video encoder to provide masked vision modeling regularization only for pre-training if not specified, that is, retaining the high efficiency of an ordinary dual-encoder structure in retrieval. In each iteration, we randomly mask a large proportion of patches in sparsely sampled videos along both the spatial and temporal dimensions to enforce a high-level understanding of local contents and temporal dynamics. The masked videos are fed into the video encoder for performing denoising auto-encoding, while the raw videos are fed into the snapshot video encoder for producing the reconstruction targets. Intuitively, given the highly corrupted video, the video encoder is trained to recover text-aligned features of the masked video content via reasoning among the visible regions along the spatial and temporal dimensions, enhancing not only the discriminativeness of local visual features but also the fine-grained cross-modality alignment.

Our contributions are three-fold. (1) We are the first to explore the potential of BERT-style pre-training in video-text retrieval with dual-encoder models. We study the pretext task of masked visual modeling in video-text pre-training and indicate its advantages in both fine-grained video context understanding and video-text local semantic alignment. (2) We introduce a flexible and effective method with a snapshot video encoder as the evolving “tokenizer” to produce learning targets for the masked video patch prediction. The video encoder gradually improved by the denoising regularizations can be used, in turn, to enhance the “tokenizer”. (3) Extensive empirical results on text-to-video retrieval on four datasets with both zero-shot and fine-tune evaluation protocols fully demonstrate the superiority of our method (Fig. 1 (b)). We further evaluate zero-shot action recognition on two datasets, which can be cast as a video-to-text retrieval task. Our method significantly surpasses its competitive counterparts by a large margin. As an additional benefit, we surprisingly find that our method achieves competitive performance on single-modality action recognition with much fewer video hours for pre-training.

2 Related Work

Pre-training for video-text retrieval. Previous pre-training methods for video-text retrieval can be divided into two categories. Methods in the first category [6,31,51,53,18,41,35,15,29,16,48], adopt the “dual-encoder” architectures, where two individual encoders are used to contrast global video and text representations, and contrastive learning [33,24] is utilized to distinguish paired video-text data with unpaired data. Although these methods are efficient for video-text retrieval, they ignore local semantics and fine-grained alignment between modalities. Methods in the second category [27,46,55,28,30,50] adopt the “joint-encoder” architectures to interact cross-modality local features through concatenating videos and texts as inputs with a binary classifier to predict whether videos and texts are aligned or not. Despite they can build local associ-

ations between videos and texts, they sacrifice the retrieval efficiency since every text-video pair needs to be fed into the encoder during inference. In this work, we adopt the “dual-encoder” architecture for efficient retrieval and use the pretext task of masked visual modeling in video-text pre-training to enhance both fine-grained video context understanding and video-text local semantic alignment.

Image Pre-training with Masked Visual Modeling. Recent works introduce masked visual modeling (MVM) to image self-supervised pre-training, where MVM masks a proportion of the visual patches and optimizes the vision Transformers to reconstruct the missing content. The reconstruction targets are proven to be the most critical. For example, [21] reconstructs the masked image patches in the pixel space, which makes the model focus on short-range dependencies and high-frequency details. [7] predicts the discrete visual tokens from a pre-learned image “tokenizer” [38], which requires one more pre-training stage on extra data (250M images [38]). Our method uses a snapshot encoder as the evolving “tokenizer” without additional pre-training stages. A concurrent work, iBOT [54], uses a similar online tokenizer to guide MVM for image pre-training. However, we have different purposes. While iBOT encourages MVM to focus more on the high-level visual semantics rather than trivial low-level reconstruction, our method uses a self-training pipeline to progressively inject the text-aware semantics into the MVM targets, aligning the text and visual domains in both global and local representations.

Masked Region/Frame Modeling in Video-text Pre-training. Similar techniques, masked region modeling (MRM) and masked frame modeling (MFM), were introduced in the “joint-encoder” methods for video-text pre-training. For example, [55] uses MRM, which masks the object regions of video frames with a pre-trained detection model [39] and predicts a distribution over fixed vocabulary for the masked-out frame region. [30,28,50] adopts MFM, which masks video frames and recovers the masked frames to the features encoded from an off-the-shelf video feature extraction network [49]. Both MRM and MFM rely on pre-trained models with extra data to obtain visual-only reconstruction targets, while our work evolves a snapshot video encoder to provide video-text aligned reconstruction targets without additional training stage on extra data.

3 Method

3.1 Revisiting Dual-encoder for Video-text Pre-training

As shown in Fig. 2, we adopt the “dual-encoder” structure to maintain high efficiency for video-text retrieval, which contains a text encoder to encode text representations from natural languages, and a video encoder to produce video representations from raw video frames. We first feed a video to the video encoder and the corresponding text description (*e.g.*, “A woman in a jumpsuit is walking along the edge of a pool”) to the text encoder to embed their respective representations, which are projected to a common feature space as v^{cls} and t^{cls} via two separate linear layers. We calculate the similarity between the video and the text by performing the dot product between two projected embeddings v^{cls} and




Fig. 2. The video-text pre-training pipeline of **M**asked visual modeling with **I**njected **L**anguage **S**emantics (MILES). Based on the “dual-encoder” architecture, *i.e.* a text encoder and a video encoder, we use a snapshot video encoder to provide reconstruction targets with injected language semantics. We first mask a proportion of a video along the spatial and temporal dimension. The masked video is fed into the video encoder to predict features of the masked video patches as v^{mask} , while the raw video is fed into the snapshot encoder to produce the reconstruction targets as s^{mask} . The snapshot encoder is progressively updated from the in-training video encoder under the constraint of contrasting global video representations v^{cls} and text representations t^{cls} .

t^{cls} . We adopt a contrastive objective [33,24] to maximize the similarity between positive video-text pairs in the batch and minimize the similarity between all other negative combinations in the batch. The separate dual encoders ensure the high efficiency for retrieval since only the dot product between video and text representations is calculated during inference.

3.2 Masked Visual Modeling in Video-text Pre-training

Overview. As illustrated in Fig. 2, we impose the regularization with injected language semantics for masked visual modeling (MVM) through using a snapshot video encoder to produce reconstruction targets for masked patches in the video, besides dual encoders for efficient retrieval. Specifically, we train the model using the pretext task of MVM with the following steps.

In the first step, a sparsely sampled video is divided and projected into a sequence of tokens following [6]. A proportion of video tokens are masked out along the spatial and temporal dimensions by being replaced with a [MASK] token, which is a learnable embedding to indicate masked patches. Positional embeddings are further added to the token sequence after [MASK] token replacement as the input token sequence. **In the second step**, the input token sequence is fed into the video encoder for performing denoising auto-encoding. The video encoder predicts the features of the masked video tokens as v^{mask} through resorting to visible video patches of spatial and temporal neighbors. **In the third step**, the raw video is fed into the snapshot video encoder to provide regularization for MVM. Features of those tokens that correspond to masked video patches for the video encoder are obtained as the reconstruction targets s^{mask} . **In the fourth step**, token-level supervision is imposed on the video en-

coder through minimizing the ℓ_2 distance between its predicted features v^{mask} and the reconstruction targets s^{mask} embedded from the snapshot encoder.

We progressively update the snapshot video encoder from the video encoder under the constraint of contrasting global video representations v^{cls} from the video encoder and text representations t^{cls} from the text encoder, which will be explained in the following section. Since the snapshot video encoder aggregates the knowledge of the in-training video encoder, its embedded visual features gradually align with language semantics. Imposing the regularization of MVM towards the output of the snapshot encoder iteratively improves video encoder to capture detailed local visual semantics that are aligned with texts, which in turn enhances the snapshot encoder to provide more effective reconstruction targets. Taking the corrupted video as input, the video encoder is trained to recover text-aligned local features of videos through reasoning with the visible patches in the video along the spatial and temporal dimension, which enhances both the discriminativeness of local features and the fine-grained video-text alignment.

Evolving Snapshot Video Encoder. We use a snapshot video encoder to embed text-aligned video features as reconstruction targets for MVM, which does not require pre-training with extra data. Inspired by previous work [9], we freeze the snapshot encoder over an epoch, and update its parameters in the k -th epoch denoted as $\{\theta_s\}_k$ with Exponentially Moving Averaged (EMA) [22] mechanism as $\{\theta_s\}_k = \lambda\{\theta_s\}_{k-1} + (1 - \lambda)\{\theta_v\}_{k-1}$, where $\{\theta_v\}_{k-1}$ denotes the parameters of the video encoder at the end of the $(k-1)$ -th epoch. The updating mechanism makes the snapshot encoder evolve more smoothly and thus provide more consistent reconstruction targets for MVM. We also explore other update rules for the snapshot encoder in Sec. 4.5, which show inferior performances.

Masking Strategy. Videos usually exhibit similar visual patterns in adjacent patches within a frame or patches of neighboring frames (spatio-temporal neighbors), which makes the masked video patches easy to recover through interpolating between the spatio-temporal neighbors. To make masked visual modeling more challenging and improve model’s spatial and temporal understanding, we adopt the “tube” masking strategy, which masks blocks of video patches along the spatial and temporal dimension instead of independently masking random patches for each frame. Specifically, we first sample a 2D mask through block-wise masking following [7], and then extend the 2D mask to 3D mask through repeating it in the temporal dimension, such that the spatially masked patches are the same for each frame. Our masking strategy refrains the video encoder from reconstructing the masked video content by extrapolation from adjacent visual patterns, and instead requires actual visual reasoning among visible patches along the spatial and temporal dimension. Other masking strategies are discussed in Sec. 4.5, which achieve worse results.

3.3 Pre-training Objectives

We combine two objectives to optimize the entire model in an end-to-end manner including a contrastive objective with Noise-Contrastive Estimation (NCE) [33,24]

and a regressive objective with ℓ_2 distance, formulated as below,

$$\mathcal{L} = \mathcal{L}_{\text{vanilla}} + \mathcal{L}_{\text{mvm}} \quad (1)$$

$$\mathcal{L}_{\text{vanilla}} = \sum_{i=1}^B \text{NCE}(v_i^{\text{cls}}, t_i^{\text{cls}}) + \sum_{i=1}^B \text{NCE}(t_i^{\text{cls}}, v_i^{\text{cls}}) \quad (2)$$

$$\text{NCE}(x_i, y_i) = -\log \frac{\exp(x_i^T y_i / \tau)}{\sum_{j=1}^B \exp(x_i^T y_j / \tau)} \quad (3)$$

$$\mathcal{L}_{\text{mvm}} = \sum_{i=1}^B \|s_i^{\text{mask}} - v_i^{\text{mask}}\|_2 \quad (4)$$

where B is the batch size and τ is the temperature hyper-parameter.

3.4 Model Architecture

Video Encoder. Video encoder takes the video as input and produces final video representations and predicts features of the masked video content. A video $V \in R^{M \times 3 \times H \times W}$ sampling M frames is first divided into $M \times N$ patches, and are further fed into a linear projection head to get a sequence of $M \times N$ tokens. We follow BERT [12] to add a learnable [CLS] token to the the beginning of the token sequence for global video representations. A proportion of video tokens are replaced with a [MASK] token, which is a learnable embedding. The token sequence after [MASK] token replacement are further added with learnable spatial and temporal positional embeddings. All tokens within the same frame are given the same temporal positional embedding, and all tokens in the same spatial location of different frames are given the same spatial positional embedding, so that the video encoder learns to ascertain the position of video patches. Following [6], the video encoder consists of a stack of space-time self-attention blocks, where each block sequentially performs temporal self-attention and then spatial self-attention on the output of previous block.

Text Encoder. Text encoder takes the nature language as input and outputs final text representations from the [CLS] token, which is concatenated to the beginning of the input text. We adopt a multi-layer bidirectional transformer structure [42] for the text encoder.

Snapshot Encoder. Snapshot encoder takes the original video as input and produces the reconstruction targets for MVM only in pre-training and is neglected for retrieval. It has exactly the same architecture as the video encoder.

4 Experiments

4.1 Pre-training Datasets

We follow [6] to jointly pre-train our model on an image dataset Google Conceptual Captions (CC3M) [43] and a video dataset WebVid-2M [6]. CC3M contains

3.3M image-text pairs with captions harvested from the web. WebVid-2M contains 2.5M video-text pairs with manually generated captions. We do not use the large-scale video-text dataset HowTo100M [32] with 136M video-text pairs. As [6] points out, the captions in HowTo100M are extracted from ASR transcription of continuous narration with incomplete sentences, thus are noisy.

4.2 Downstream Tasks

Text-to-Video Retrieval. (a). **MSR-VTT** [52] consists of 10K YouTube videos with 200K descriptions, which is divided into 9K and 1K videos for training and testing. (b). **MSVD** [11] contains 1,970 videos from YouTube with 80K descriptions. 1200, 100 and 670 videos are split out for training, validation and testing respectively. (c). **LSMDC** [40] consists of 118,081 video clips from 202 movies, where the validation set and the test set contain 7,408 and 1,000 videos. (d). **DiDeMo** [5] contains 10K Flickr videos with 40K sentences. 1K videos are split out for testing. We follow [6] to concatenate all sentences of a video as a single description. We perform evaluation with both the zero-shot and fine-tune setups, and adopt Recall and Median Rank as the evaluation metric.

Action Recognition. (a). **HMDB51** [26], which consists of 6,766 videos with 51 categories. (b). **UCF101** [44], which consists 13,320 videos with 101 action classes. Both datasets have three standard training/test splits. We explore three experimental settings for evaluation, including **linear**, where we fix the parameters of the video encoder and only optimize a linear classifier, **fully fine-tuning**, where we optimize the parameters of the video encoder and the linear classifier together, and **zero-shot**, where we perform video-to-text retrieval through describing a video with the name of its action class following [37]. Averaged results over three training/test splits are reported.

4.3 Implementation Details

For fair comparison, we follow the recent work [6] for implementation. We first resize a video to 224×224 , and then divide a video into M equal segments. We randomly sample a single frame within each segment for training and uniformly sample a single frame within each segment for testing. We adopt the same model architecture as [6] with a video encoder and a text encoder. The video encoder consists of 12 space-time self-attention blocks [8] with patch size $P = 16$, and sequence dimension $D = 768$. We initialize the video encoder with ViT [14] weights trained on ImageNet-21k following [6,34]. The text encoder is instantiated as DistilBERT [42] pre-trained on English Wikipedia and Toronto Book Corpus. The snapshot video encoder has the same architecture as the video encoder, and is updated over an epoch with the weights of the video encoder using $\{\theta_s\}_k = \lambda\{\theta_s\}_{k-1} + (1 - \lambda)\{\theta_v\}_{k-1}$, where θ is set as 0.996 in our experiments. We set the dimension of the common feature space as 256 and the temperature hyper-parameter of the contrastive objective as 0.05. For visual augmentation, we randomly crop and horizontally flip during training, and center crop the

maximal square crop during testing. We adopt a temporal curriculum learning following [6], where we first pre-train our model on the image dataset CC3M and video dataset WebVid-2M sampling 1 frame, and then on the video dataset WebVid-2M sampling 4 frames. Pre-training with 1 frame takes 16 epochs with the batch size of 2048 and the learning rate of 1×10^{-4} , where the first epoch trains the model with only the contrastive objective as a warm-up. Pre-training with 4 frames takes 4 epochs with the batch size of 1024 and the learning rate of 3×10^{-5} . We adopt random mask sampling with the masking ratio of 75% for 1-frame pre-training and block-wise mask sampling with the masking ratio of 75% for 4-frame pre-training, where the spatial mask is repeated in the temporal dimension. For evaluating downstream tasks, we uniformly sample 4 frames for text-to-video retrieval and 16 frames for action recognition following the setting of [6,31]. We follow [6] to expand the temporal embeddings through filling zeros to enable the training of longer frames.

Table 1. Text-to-video retrieval results on MSR-VTT test set, where **higher** R@k and **lower** MedR (Median Rank) are better. Zero-shot evaluation results are shown on the top while fine-tuning on the bottom. “Video Input” lists the model to extract 3D features, where “Raw Videos” means training on raw video frame pixels without using pre-extracted features. “Pairs” lists the number of video-text pairs for pre-training.

Method	Year	Video Input	Pre-train Dataset	Pairs	R@1	R@5	R@10	MedR
ActBERT [55]	2020	ResNet-3D	HowTo100M	120M	8.6	23.4	33.1	36.0
MMV [2]	2020	Raw Videos	HowTo100M, AudioSet	138M	9.3	23.0	31.1	38.0
MIL-NCE [31]	2020	Raw Videos	HowTo100M	120M	9.9	24.0	32.4	29.6
VATT [1]	2021	Raw Videos	HowTo100M, AudioSet	138M	-	-	29.7	49.0
NoiseEst [4]	2021	ResNeXt-101	HowTo100M	110M	8.0	21.3	29.3	33.0
TACo [53]	2021	I3D, S3D	HowTo100M	120M	9.8	25.0	33.4	29.0
VideoCLIP [51]	2021	S3D	HowTo100M	110M	10.4	22.2	30.0	-
MCN [10]	2021	ResNeXt-101	HowTo100M	120M	10.5	25.2	33.8	-
SupportSet [35]	2021	R(2+1)D-34	HowTo100M	120M	12.7	27.5	36.2	24.0
Frozen [6]	2021	Raw Videos	CC3M, WebVid-2M	5.5M	18.7	39.5	51.6	10.0
AVLnet [41]	2021	ResNeXt-101	HowTo100M	120M	19.6	40.8	50.7	9.0
Ours	2022	Raw Videos	CC3M, WebVid-2M	5.5M	26.1	47.2	56.9	7.0
ActBERT [55]	2020	ResNet-3D	HowTo100M	120M	16.3	42.8	56.9	10.0
UniVL [30]	2020	S3D	HowTo100M	110M	21.2	49.6	63.1	6.0
MMT [15]	2020	S3D	HowTo100M	120M	26.6	57.1	69.6	4.0
HERO [28]	2021	SlowFast	TV and HowTo100M	120M	16.8	43.4	57.7	-
NoiseEst [4]	2021	ResNeXt-101	HowTo100M	110M	17.4	41.6	53.6	8.0
ClipBert [27]	2021	Raw Videos	COCO, VisGenome	5.6M	22.0	46.8	59.9	6.0
AVLnet [41]	2021	ResNeXt-101	HowTo100M	120M	27.1	55.6	66.6	4.0
VLM [50]	2021	S3D	HowTo100M	110M	28.1	55.5	67.4	4.0
TACo [53]	2021	I3D, S3D	HowTo100M	120M	28.4	57.8	71.2	4.0
SupportSet [35]	2021	R(2+1)D-34	HowTo100M	120M	30.1	58.5	69.3	3.0
VideoCLIP [51]	2021	S3D	HowTo100M	110M	30.9	55.4	66.8	-
Frozen [6]	2021	Raw Videos	CC3M, WebVid-2M	5.5M	31.0	59.5	70.5	3.0
Ours	2022	Raw Videos	CC3M, WebVid-2M	5.5M	37.7	63.6	73.8	3.0

4.4 Main Results

Text-to-Video Retrieval Results on MST-VTT [52] can be seen in Table. 1. We have the following observations. First of all, our method outperforms all recent work by a large margin. The significant improvement of the performance under the zero-shot evaluation protocol indicates that our pre-trained model is more generalizable that can be used for out-of-domain text-to-video retrieval. Fine-tuning our pre-trained model on the training set of MSR-VTT also achieves better performance. Second, the majority of previous work pre-extract 3D features from “expert” models as the input of the video encoder (*e.g.* SupportSet [35] uses expert features from a 34-layer, R(2+1)-D model pre-trained on IG65M [17]). By contrast, our model takes raw video frame pixels as inputs without using any pre-extracted features and surpasses its counterparts. Third, previous work mainly pre-train their models on the large-scale HowTo100M [32], which is 20× larger in terms of video-text pairs than CC3M [43] and WebVid-2M [6], thus higher computation cost is required. Our pre-trained model achieves higher performance with lower computation cost. Finally, some work [27,28,30,50,55] adopts a joint encoder to concatenate videos and texts as inputs, thus every text-video pair needs to be fed into the encoder during inference, resulting in low efficiency for retrieval. By comparison, our model adopts the efficient “dual-encoder” architecture with only a video encoder and a text encoder for inference.

We further show text-to-video retrieval results on MSVD [11] in Table. 2, LSMDC in Table. 3, and DiDeMo [5] in Table. 4. Under both the zero-shot (top) and the fine-tune (bottom) evaluation protocols, our model achieves the best performance on three datasets, demonstrating the effectiveness of our method in utilizing the pretext task of masked visual modeling in video-text pre-training to learn fine-grained cross-modality alignment between videos and texts.

Table 2. Text-to-video retrieval results on MSVD test set. **Table 3.** Text-to-video retrieval results on LSMDC test set.

Method	R@1	R@5	R@10	MedR
NoiseEst [4]	13.7	35.7	47.7	12.0
SupportSet [35]	21.4	46.2	57.7	6.0
Frozen [6]	38.7	70.1	80.1	2.0
Ours	44.4	76.2	87.0	2.0

Method	R@1	R@5	R@10	MedR
AVLnet [41]	1.4	5.9	9.4	273.5
NoiseEst [4]	4.2	11.6	17.1	119.0
Frozen [6]	9.3	22.0	30.1	51.0
Ours	11.1	24.7	30.6	50.7

Method	R@1	R@5	R@10	MedR
NoiseEst [4]	20.3	49.0	63.3	6.0
SupportSet [35]	28.4	60.0	72.9	4.0
Frozen [6]	45.6	79.8	88.2	2.0
Ours	53.9	83.5	90.2	1.0

Method	R@1	R@5	R@10	MedR
NoiseEst [4]	6.4	19.8	28.4	39.0
MMT [15]	12.9	29.9	40.1	19.3
Frozen [6]	15.0	30.8	39.8	20.0
Ours	17.8	35.6	44.1	15.5

Action Recognition We evaluate zero-shot action recognition on HMDB51 [26] and UCF101 [44], which is cast as video-to-text retrieval by using the name of a video’s action class as its description following [37]. Table. 5 lists the results, where our model outperforms the competitive baselines by a large margin. Our model improves the averaged top-1 accuracy over three splits by 16.9% and 10.5% on HMDB51, 24.9% and 6.8% on UCF101 compared with ClipBert [27] and Frozen [6]. Our method learns powerful cross-modality representations that enable effective zero-shot action recognition.




Fig. 3. The visualizations of the local cross-modality alignment between the text token and visual tokens. Compared with the model that is not trained with MVM, our pre-trained model aligns words with corresponding visual regions accurately, indicating that MVM can effectively train the model to capture local video-text alignment.

Table 4. Text-to-video retrieval results on DiDeMo test set.

Method	R@1	R@5	R@10	MedR
VideoCLIP [51]	16.6	46.9	-	-
Frozen [6]	21.1	46.0	56.2	7.0
Ours	27.2	50.3	63.6	5.0
HERO [28]	2.1	-	11.4	-
CE [29]	16.1	41.1	82.7	8.3
ClipBert [27]	20.4	48.0	60.8	6.0
Frozen [6]	31.0	59.8	72.4	3.0
Ours	36.6	63.9	74.0	3.0

We further explore the single-modality video representations, where the features from the video encoder are extracted and are fed into a linear classifier. In Table. 6, our model achieves higher accuracy than some competitive methods, which pre-train their models on datasets with considerably longer video duration. For example, XDC [3] pre-trains on IG-Kinetics that is $14\times$ longer than our pre-training dataset, but gets worse results. Although MMV [2] surpasses our method when it pre-trains the model on HowTo100M and AudioSet ($11\times$ longer than ours) and utilizes extra modalities such as audio and text, its performance is far worse than ours with only text and video or audio and video as inputs.

Table 5. Zero-shot action recognition results on HMDB51 and UCF101, with top-1 accuracy as the evaluation metric. “S” denotes different testing splits and “Mean” reports the averaged results over three splits.

Method	HMDB51				UCF101			
	S1	S2	S3	Mean	S1	S2	S3	Mean
ClipBert [27]	20.0	22.0	22.3	21.4	27.5	27.0	28.8	27.8
Frozen [6]	27.5	28.3	27.7	27.8	45.4	44.7	47.7	45.9
Ours	38.4	38.6	37.8	38.3	51.8	53.4	52.8	52.7

Table 6. Action recognition results on HMDB51 and UCF101 under linear evaluation (Linear) and fully fine-tuning evaluation (Full), with top-1 accuracy as the evaluation metric. “Modality” denotes the extra modality used for pre-training besides videos, *i.e.*, audio (A), text (T), optical flow (OF), motion vector (MV). “Length” denotes the video length for pre-training in K hours.

Method	Modality Length (K)		HMDB		UCF101	
			Linear	Full	Linear	Full
CCL [25]	-	1.8	29.5	37.8	54.0	69.4
CBT [45]	-	1.8	29.5	44.5	54.0	79.5
MemDPC [19]	OF	1.8	30.5	54.5	54.1	86.1
CoCLR [20]	OF	1.8	52.4	62.9	77.8	90.6
MVCGC [23]	MV	1.8	53.0	63.4	78.0	90.8
XDC.R [3]	A	188.3	49.9	61.2	80.7	88.8
XDC.K [3]	A	188.3	56.0	63.1	85.3	91.5
MIL-NCE [31]	T	134.5	54.8	59.2	83.4	89.1
Frozen [6]	T	13.0	61.3	66.3	87.8	89.8
VATT [1]	A, T	139.8	63.3	-	89.2	-
ELO [36]	A, OF	115.0	64.5	67.4	-	93.8
MMV [2]	A	134.5	53.6	-	77.1	-
MMV [2]	T	134.5	55.1	-	86.8	-
MMV [2]	A, T	139.8	67.1	75.0	91.8	95.2
Ours	T	13.0	65.4	70.5	89.5	92.2

Visualization We visualize the cross-modality alignment between text tokens and visual tokens by calculating the similarity map between features embedded from the text encoder and video encoder. As shown in Fig. 3, compared with the model that is not pre-trained with MVM, our pre-trained model aligns words with corresponding visual regions accurately. For example, visual features of the football region above the man’s head in the second row exhibits large similarity with the features of the word “football”. Performing MVM with the video-text aligned features as the reconstruction targets effectively trains the model to enhance local video-text alignment with the “dual-encoder” architecture.

4.5 Ablation Study

Reconstruction targets. We explore different reconstruction targets for MVM in video-text pre-training including raw frame pixels as in [21], discrete visual tokens from a learned image “tokenizer” [38] as in [7], and the text-aligned features in this work. As shown in Table. 7, using masked visual modeling with different targets improves performance than the baseline model in the first row with only the contrastive objective. Imposing the regularization for MVM with text-aligned features achieves better results than with other visual-only reconstruction targets, which indicates that the local video-text alignment serves as an important objective for the prediction of masked video patches. Besides the “dual-encoder”, we further adopt a joint encoder to interact videos with texts for performing denoising auto-encoding, which improves the performance over the “dual-encoder” structure with pixels and discrete tokens as the reconstruction targets. Using the joint encoder with the aligned features as the targets brings poorer results, since the learning of cross-modality alignment is mainly

achieved by the joint encoder, which does not benefit dual encoders for retrieval. When the reconstruction targets contain video-text aligned semantics, it is more effective to impose MVM regularization on the output local video tokens.

Table 7. Ablation studies on different reconstruction targets, where “Joint” denotes the use of a joint encoder to concatenate videos and texts. Results of zero-shot text-to-video retrieval on MSR-VTT are reported.

Targets	Joint	R@1↑	R@5↑	R@10↑	MnR↓
-	×	21.9	43.5	53.5	52.8
Pixels	×	22.9	44.2	53.7	53.3
Pixels	✓	24.4	45.2	55.3	53.1
Discrete tokens	×	24.1	46.6	56.0	51.9
Discrete tokens	✓	25.0	47.1	56.6	51.8
Aligned features	✓	25.2	47.1	55.6	48.3
Aligned features	×	26.1	47.2	56.9	46.9

Updating mechanism of the snapshot encoder. We compare different strategies to update the snapshot video encode from the video encoder. In our method, we use the video encoder of the previous epoch with momentum update as the snapshot encoder and freeze the the snapshot encoder over an epoch. As shown in Table. 8, compared with our updating mechanism in the last row, using the video encoder of the current iteration, or of the previous iteration with or without momentum update all achieve less competitive performance. Since these mechanisms update the snapshot encoder in each iteration, the reconstruction targets are less consistent to impose effective regularization for MVM. Using the video encoder of previous epoch without momentum update also drops performances, since it masks the evolution of the snapshot encoder more sharp.

Table 8. Ablation studies on the updating mechanism of the snapshot video encoder. “Mom” indicates whether the momentum update is used. Results of zero-shot text-to-video retrieval on MSR-VTT are reported.

Mechanism	Mom	R@1↑	R@5↑	R@10↑	MnR↓
Current iter	-	25.0	47.1	55.7	48.8
Previous iter	×	24.7	46.4	54.6	48.9
Previous iter	✓	24.5	46.2	55.3	48.8
Previous epoch	×	23.8	47.0	56.3	48.2
Previous epoch	✓	26.1	47.2	56.9	46.9

Masking strategy. We explore different strategies of masking the videos along the spatial and temporal dimension to perform MVM for multi-frame pre-training. We first propose a frame-wise masking, where a proportion of full frames are masked (*e.g.* 25% masking ratio means masking one random frame among four sampled frames), so that the model can only resort to neighboring visible frames

to recover the masked frames. As shown in Table. 9, frame-wise masking strategy degrades performance, which indicates that reasoning with visible patches along the spatial dimension is also essential. When “tube” masking [47] is not adopted and each frame is spatially masked individually, the model achieves the worst results because the model could temporally “interpolate” between frames with visible patches for reconstruction. We further compare different spatial masking strategies with “tube” temporal masking. Compared with block-wise masking, which tends to mask large blocks, masking with random sampling achieves less satisfactory performance because it makes the model easier to reconstruct the masked patches by resorting to visible patches within the frame. Finally, we study the effect of the masking ratio and find that our method with block-wise masking works reasonably well at a ratio of 75%, but degrades at 65% when the reconstruction task becomes easier and also degenerates at a ratio 85% when the reconstruction task is too difficult.

Table 9. Ablation studies on the masking strategy. “Tube” denotes whether the spatially masked patches are the same for each frame. Results of zero-shot text-to-video retrieval on MSR-VTT are reported.

Masking	Ratio	Tube	R@1↑	R@5↑	R@10↑	MnR↓
Frame	25%	-	25.3	46.2	55.8	47.9
Frame	50%	-	25.5	46.6	55.9	47.8
Random	65%	✓	24.9	46.6	56.1	47.4
Random	75%	✓	25.0	46.8	56.4	47.2
Random	85%	✓	25.1	47.0	56.4	47.2
Block	75%	×	24.9	45.9	55.4	47.3
Block	65%	✓	25.7	46.9	56.2	47.0
Block	85%	✓	25.8	46.6	55.5	47.0
Block	75%	✓	26.1	47.2	56.9	46.9

5 Conclusion

In this work, we explore masked visual modeling in video-text pre-training with the “dual-encoder” architecture for efficient video-text retrieval. We introduce an effective method with a snapshot video encoder to produce reconstruction targets with injected language semantics for the masked video patch prediction, which does not require extra pre-training stages. Training the video encode to recover the text-aligned features of masked video patches through reasoning with the visible regions along the spatial and temporal dimension strengthens both the awareness of local visual features and the fine-grained cross-modality alignment. Extensive evaluations on the text-to-video retrieval and action recognition clearly show the great advantage of our method.

Acknowledgment Ping Luo is supported by the General Research Fund of HK No.27208720, No.17212120, and No.17200622.

References

1. Hassan Akbari, Linagzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong. Vatt: Transformers for multimodal self-supervised learning from raw video, audio and text. *arXiv preprint arXiv:2104.11178*, 2021. 9, 12
2. Jean-Baptiste Alayrac, Adria Recasens, Rosalia Schneider, Relja Arandjelovic, Jason Ramapuram, Jeffrey De Fauw, Lucas Smaira, Sander Dieleman, and Andrew Zisserman. Self-supervised multimodal versatile networks. *NeurIPS*, 2(6):7, 2020. 9, 11, 12
3. Humam Alwassel, Dhruv Mahajan, Bruno Korbar, Lorenzo Torresani, Bernard Ghanem, and Du Tran. Self-supervised learning by cross-modal audio-video clustering. In *NeurIPS 2020*, 2020. 11, 12
4. Elad Amrani, Rami Ben-Ari, Daniel Rotman, and Alex Bronstein. Noise estimation using density estimation for self-supervised multimodal learning. In *AAAI*, volume 35, pages 6644–6652, 2021. 9, 10
5. Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *ICCV*, pages 5803–5812, 2017. 8, 10
6. Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval, 2021. 1, 3, 5, 7, 8, 9, 10, 11, 12
7. Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *ICLR*, 2022. 2, 4, 6, 12
8. Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding. *arXiv preprint arXiv:2102.05095*, 2021. 8
9. Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, pages 9650–9660, 2021. 6
10. Brian Chen, Andrew Rouditchenko, Kevin Duarte, Hilde Kuehne, Samuel Thomas, Angie Boggust, Rameswar Panda, Brian Kingsbury, Rogerio Feris, David Harwath, et al. Multimodal clustering networks for self-supervised learning from unlabeled videos. *arXiv preprint arXiv:2104.12671*, 2021. 9
11. David Chen and William B Dolan. Collecting highly parallel data for paraphrase evaluation. In *ACL*, pages 190–200, 2011. 8, 10
12. Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *NAACL*, 2019. 2, 7
13. Xiaoyi Dong, Jianmin Bao, Ting Zhang, Dongdong Chen, Weiming Zhang, Lu Yuan, Dong Chen, Fang Wen, and Nenghai Yu. Peco: Perceptual codebook for bert pre-training of vision transformers. *arXiv preprint arXiv:2111.12710*, 2021. 2
14. Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xi-aohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2020. 2, 8
15. Valentin Gabeur, Chen Sun, Karteek Alahari, and Cordelia Schmid. Multi-modal transformer for video retrieval. In *ECCV*, pages 214–229, 2020. 3, 9, 10
16. Yuying Ge, Yixiao Ge, Xihui Liu, Dian Li, Ying Shan, Xiaohu Qie, and Ping Luo. Bridging video-text retrieval with multiple choice questions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16167–16176, 2022. 3
17. Deepti Ghadiyaram, Du Tran, and Dhruv Mahajan. Large-scale weakly-supervised pre-training for video action recognition. In *CVPR*, pages 12046–12055, 2019. 10

18. Simon Ging, Mohammadreza Zolfaghari, H Pirsiavash, and Thomas Brox. Coot: Cooperative hierarchical transformer for video-text representation learning. In *NeurIPS*, 2020. [1](#), [3](#)
19. Tengda Han, Weidi Xie, and Andrew Zisserman. Memory-augmented dense predictive coding for video representation learning. In *ECCV*, pages 312–329, 2020. [12](#)
20. Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning. *NeurIPS*, 33:5679–5690, 2020. [12](#)
21. Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. *arXiv preprint arXiv:2111.06377*, 2021. [4](#), [12](#)
22. Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. [6](#)
23. Yuqi Huo, Mingyu Ding, Haoyu Lu, Nanyi Fei, Zhiwu Lu, Ji-Rong Wen, and Ping Luo. Compressed video contrastive learning. *NeurIPS*, 34, 2021. [12](#)
24. Rafal Jozefowicz, Oriol Vinyals, Mike Schuster, Noam Shazeer, and Yonghui Wu. Exploring the limits of language modeling. *arXiv preprint arXiv:1602.02410*, 2016. [3](#), [5](#), [6](#)
25. Quan Kong, Wenpeng Wei, Ziwei Deng, Tomoaki Yoshinaga, and Tomokazu Murakami. Cycle-contrast for self-supervised video representation learning. 2020. [12](#)
26. Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *ICCV*, pages 2556–2563. IEEE, 2011. [8](#), [10](#)
27. Jie Lei, Linjie Li, Luowei Zhou, Zhe Gan, Tamara L Berg, Mohit Bansal, and Jingjing Liu. Less is more: Clipbert for video-and-language learning via sparse sampling. In *CVPR*, pages 7331–7341, 2021. [3](#), [9](#), [10](#), [11](#)
28. Linjie Li, Yen-Chun Chen, Yu Cheng, Zhe Gan, Licheng Yu, and Jingjing Liu. Hero: Hierarchical encoder for video+ language omni-representation pre-training. In *EMNLP*, pages 2046–2065, 2020. [3](#), [4](#), [9](#), [10](#), [11](#)
29. Yang Liu, Samuel Albanie, Arsha Nagrani, and Andrew Zisserman. Use what you have: Video retrieval using representations from collaborative experts. *arXiv preprint arXiv:1907.13487*, 2019. [3](#), [11](#)
30. Huaishao Luo, Lei Ji, Botian Shi, Haoyang Huang, Nan Duan, Tianrui Li, Jason Li, Taroon Bharti, and Ming Zhou. Univl: A unified video and language pre-training model for multimodal understanding and generation. *arXiv preprint arXiv:2002.06353*, 2020. [1](#), [3](#), [4](#), [9](#), [10](#)
31. Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-end learning of visual representations from uncurated instructional videos. In *CVPR*, pages 9879–9889, 2020. [1](#), [3](#), [9](#), [12](#)
32. Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *ICCV*, pages 2630–2640, 2019. [8](#), [10](#)
33. Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. [3](#), [5](#), [6](#)
34. Junting Pan, Ziyi Lin, Xiatian Zhu, Jing Shao, and Hongsheng Li. Parameter-efficient image-to-video transfer learning. *arXiv preprint arXiv:2206.13559*, 2022. [8](#)
35. Mandela Patrick, Po-Yao Huang, Yuki Asano, Florian Metze, Alexander G Hauptmann, Joao F Henriques, and Andrea Vedaldi. Support-set bottlenecks for video-text representation learning. In *ICLR*, 2020. [1](#), [3](#), [9](#), [10](#)

36. AJ Piergiovanni, Anelia Angelova, and Michael S Ryoo. Evolving losses for unsupervised video representation learning. In *CVPR*, pages 133–142, 2020. [12](#)
37. Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021. [8](#), [10](#)
38. Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, 2021. [2](#), [4](#), [12](#)
39. Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *NeurIPS*, 28, 2015. [4](#)
40. Anna Rohrbach, Marcus Rohrbach, Niket Tandon, and Bernt Schiele. A dataset for movie description. In *CVPR*, pages 3202–3212, 2015. [8](#)
41. Andrew Rouditchenko, Angie Boggust, David Harwath, Brian Chen, Dhiraj Joshi, Samuel Thomas, Kartik Audhkhasi, Hilde Kuehne, Rameswar Panda, Rogerio Feris, et al. Avlnet: Learning audio-visual language representations from instructional videos. *arXiv preprint arXiv:2006.09199*, 2020. [1](#), [3](#), [9](#), [10](#)
42. Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019. [7](#), [8](#)
43. Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*, pages 2556–2565, 2018. [7](#), [10](#)
44. Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. [8](#), [10](#)
45. Chen Sun, Fabien Baradel, Kevin Murphy, and Cordelia Schmid. Learning video representations using contrastive bidirectional transformer. *arXiv preprint arXiv:1906.05743*, 2019. [12](#)
46. Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *ICCV*, pages 7464–7473, 2019. [3](#)
47. Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *arXiv preprint arXiv:2203.12602*, 2022. [14](#)
48. Jinpeng Wang, Yixiao Ge, Guanyu Cai, Rui Yan, Xudong Lin, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Object-aware video-language pre-training for retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3313–3322, 2022. [3](#)
49. Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*, pages 305–321, 2018. [4](#)
50. Hu Xu, Gargi Ghosh, Po-Yao Huang, Prahal Arora, Masoumeh Aminzadeh, Christoph Feichtenhofer, Florian Metze, and Luke Zettlemoyer. Vlm: Task-agnostic video-language model pre-training for video understanding. *arXiv preprint arXiv:2105.09996*, 2021. [3](#), [4](#), [9](#), [10](#)
51. Hu Xu, Gargi Ghosh, Po-Yao Huang, Dmytro Okhonko, Armen Aghajanyan, Florian Metze, Luke Zettlemoyer, and Christoph Feichtenhofer. Videoclip: Contrastive pre-training for zero-shot video-text understanding. *arXiv preprint arXiv:2109.14084*, 2021. [1](#), [3](#), [9](#), [11](#)
52. Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*, pages 5288–5296, 2016. [8](#), [10](#)

53. Jianwei Yang, Yonatan Bisk, and Jianfeng Gao. Taco: Token-aware cascade contrastive learning for video-text alignment. In *ICCV*, pages 11562–11572, 2021. 1, 3, 9
54. Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *ICLR*, 2022. 2, 4
55. Linchao Zhu and Yi Yang. Actbert: Learning global-local video-text representations. In *CVPR*, pages 8746–8755, 2020. 3, 4, 9, 10