

An Empirical Study of End-to-End Video-Language Transformers with Masked Visual Modeling

Tsu-Jui Fu^{†*}, Linjie Li^{‡*}, Zhe Gan[‡], Kevin Lin[‡], William Yang Wang[†], Lijuan Wang[‡], Zicheng Liu[‡]

[†]UC Santa Barbara [‡]Microsoft

{tsu-juiifu, william}@cs.ucsb.edu

{lindsey.li, zhe.gan, keli, lijuanw, zliu}@microsoft.com

Abstract

Masked visual modeling (MVM) has been recently proven effective for visual pre-training. While similar reconstructive objectives on video inputs (e.g., masked frame modeling) have been explored in video-language (VidL) pre-training, previous studies fail to find a truly effective MVM strategy that can largely benefit the downstream performance. In this work, we systematically examine the potential of MVM in the context of VidL learning. Specifically, we base our study on a fully end-to-end Video-Language Transformer (VIOLET) [17], where the supervision from MVM training can be backpropagated to the video pixel space. In total, eight different reconstructive targets of MVM are explored, from low-level pixel values and oriented gradients to high-level depth maps, optical flow, discrete visual tokens, and latent visual features. We conduct comprehensive experiments and provide insights into the factors leading to effective MVM training, resulting in an enhanced model VIOLETv2. Empirically, we show VIOLETv2 pre-trained with MVM objective achieves notable improvements on 13 VidL benchmarks, ranging from video question answering, video captioning, to text-to-video retrieval.¹

1. Introduction

Video, containing multiple modalities in nature, has been used as an epitome to test how AI systems perceive. Video-language (VidL) research aims at extending this ability to convey perception via language. Popular VidL tasks were introduced, such as text-to-video retrieval [31, 62, 84], video question answering [27, 83], and video captioning [7, 84]. Recent progresses in VidL learning mostly focus on VidL pre-training [54, 66, 92] with video-text matching [41, 88] and masked language modeling [12]. There have also been attempts on similar masked modeling on vision inputs.

For example, masked frame modeling [41] aims to recover masked frame representations. However, the pre-extracted video features cannot be refined during pre-training, which may limit its effectiveness. More recently, VIOLET [17] designs an end-to-end video-language transformer and proposes to reconstruct discrete visual tokens for masked frame patches. Though showing some promises in recovering visual semantics, the performance improvements on downstream VidL tasks are still marginal.

Meanwhile, self-supervised visual pre-training has been proven highly effective by reconstructing the masked image patches through raw pixel values [23, 82], discrete visual tokens [4, 90], or visual-semantic features [79, 80]. However, they all only focus on the visual modality. It is unclear which variant of masked visual modeling (MVM) objectives can help VidL learning, especially given that the paired language inputs can already provide high-level semantics.

Motivated by this, we conduct a comprehensive study of MVM for VidL learning. As illustrated in Figure 1, we base our study on the fully end-to-end Video-Language Transformer (VIOLET) [17], and study a broad spectrum of MVM targets, including RGB pixel values (Pixel), histogram of oriented gradients (HOG), depth maps (Depth), optical flow (Flow), discrete visual tokens (VQ), spatial-focused image features (SIF), temporal-aware video features (TVF), and multimodal features (MMF). During pre-training, we mask out some proportions of the video input along both spatial and temporal dimensions, and the model learns to recover the MVM targets for these masked patches. Equipped with another two standard pre-training tasks (*i.e.*, video-text matching and masked language modeling), we empirically verify the effectiveness of different MVM variants on downstream VidL tasks.

Our study reveals that: (*i*) spatial-focused image features (SIF) is the most effective MVM target on video-text inputs; and (*ii*) the effects of different MVM targets on downstream VidL tasks are not shared between video-text and image-text inputs. For example, SIF extracted from the same model brings a large drop on downstream VidL

¹Code has been released at https://github.com/tsujuiifu/pytorch_empirical-mvm.

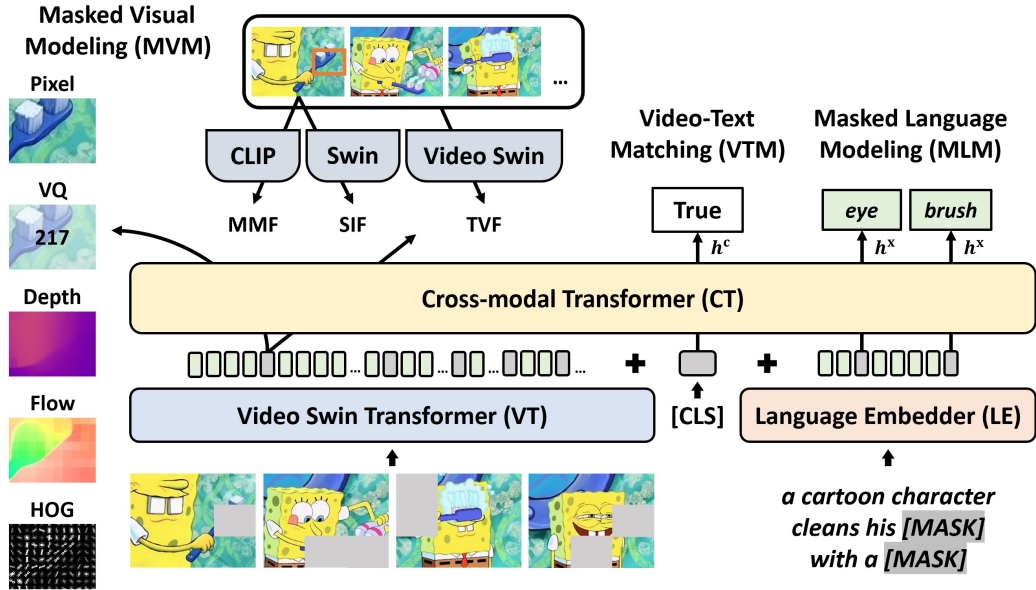


Figure 1. We systematically explore *eight* masked visual modeling (MVM) targets for end-to-end video-language (VidL) pre-training, including RGB pixel values (Pixel), histogram of oriented gradients (HOG), depth maps (Depth), optical flow (Flow), discrete visual tokens (VQ), spatial-focused image features (SIF), temporal-aware video features (TVF), and multimodal features from CLIP (MMF). Besides MVM, we pre-train VIOLET model [17] along with video-text matching (VTM) and masked language modeling (MLM).

performance when pre-trained with image-text pairs. In addition, we conduct comprehensive analyses of the masking strategy and ratio, combination of different MVM targets, to shed light on effective MVM training for VidL learning. We name the enhanced version of the original VIOLET [17] with the best MVM strategy as VIOLETv2.

Our contributions can be summarized as follows. We present an empirical study of masked visual modeling for video-language pre-training, with comprehensive analyses to reveal the ingredients for effective MVM training. VIOLETv2 with the best MVM recipe achieves strong performance on 13 VidL datasets. Concretely, compared to models pre-trained on the same 5M corpus, VIOLETv2 brings mean improvements of +5.4% accuracy on video question answering, +6.6% recall on text-to-video retrieval, and +11.4 CIDEr on video captioning. Direct comparison to VIOLET [17] also shows notable advantages of our model, even when pre-trained with much less data.

2. Related Work

Video-Language Understanding. Joint video-language (VidL) understanding [18, 19, 28, 34, 43, 47, 57] aims at interpreting the physical world via both vision and text perception. Researchers have explored such capability on VidL tasks including text-to-video retrieval [31, 39, 41, 62, 84], video question answering [27, 37, 38, 83], moment retrieval [21, 25, 31, 39], and video captioning [62, 78, 84, 91]. Prior arts before the large-scale pre-training era [15, 20, 34, 35, 38, 89] leverage offline extracted video features [2, 6, 11, 16, 24, 29, 32, 76, 81]. Later on, VidL pre-trained mod-

els [41, 54, 66, 92] built on the above pre-extracted features have shown promising results. To enhance the performance, there have been parallel interests in bringing in more modalities from raw video inputs [18, 46, 63] and end-to-end training [3, 36, 53, 88], aiming to elevate video representations.

Masked Visual Modeling (MVM). Aligned with the success of transformer-based [74] language pre-training [33, 48], image-text pre-training [9, 67] and video-text pre-training [30, 85, 86] have shown promising results on diverse vision-language (VL) tasks. Popular VL pre-training tasks include visual-text matching (VTM) and masked language modeling (MLM), which are directly adapted from language pre-training [12]. Similar masked modeling on visual inputs [9, 14, 41] has also been introduced to VL pre-training, but are not as useful. Among the literature of vision pre-training itself, MAE [23, 70] and SimMIM [82] reconstruct the pixels of the masked image patches to enhance visual representation. BEiT [4], iBOT [90], VIM-PAC [68], and BEVT [77] adopt a BERT-like pre-training strategy to recover the missing visual tokens. On the other hand, MaskFeat [79] and MVP [80] consider latent features for MVM, including hand-crafted HOG features and image features extracted from pre-trained CLIP models [58]. Unlike previous studies exploring MVM on uni-modal data, in this study, we conduct a comprehensive investigation on how different MVM targets can help VidL learning.

The most relevant study to ours is VIOLET [17], which proposes to augment VidL pre-training with masked visual token modeling, while only showing marginal improvements on downstream performance. In contrast, our comprehensive investigation covers diverse MVM targets and

studies different combinations of masking strategies, which encompasses the design of MVM as well as shows large performance improvements on downstream VidL tasks.

3. Method

We first describe the base model VIOLET in Section 3.1, and then introduce the problem formulation of our investigation in Section 3.2. Section 3.3 discusses eight different target features for masked visual modeling (MVM).

3.1. End-to-End Video-Language Transformer

We conduct our empirical study using an end-to-end Video-Language Transformer (VIOLET) [17], with 3 components: Video Swin Transformer (VT), Language Embedder (LE), and Cross-modal Transformer (CT). VIOLET takes video \mathcal{V} and sentence \mathcal{X} as inputs. Sparse-sampled frames $\{f_1, f_2, \dots\}$ from \mathcal{V} are first segmented into a set of video patches, and then processed by VT to compute video features $v = \{v_1, v_2, \dots\}$. LE extracts the word embeddings $w = \{w_1, w_2, \dots\}$ for each word token $\{x_1, x_2, \dots\}$ in \mathcal{X} . Then, CT performs cross-modal fusion on top of v and w to produce joint VidL representations $h = [h^v, h^c, h^x]$, where h^v, h^c, h^x denote the hidden representations of video patches, the special [CLS] token, and other word tokens.

3.2. Problem Setting

Given a large-scale video-language (VidL) dataset D , we aim to pre-train a VidL transformer to learn effective video-text representations. The learned representations can be transferred to downstream tasks for performance improvement. Different from existing works that focus on MVM for pure vision problems [4, 23, 90], we study MVM as a VidL pre-training task. Given a video-text pair $(\mathcal{V}, \mathcal{X})$ where \mathcal{V} is a sequence of video frames and \mathcal{X} is a sequence of word tokens. As shown in Figure 1, we randomly mask out some portions of the input frames \mathcal{V} , and learn to predict the target features corresponding to the masked patches. To output a correct prediction, the model will have to resort to other relevant video frames \mathcal{V} and/or text tokens \mathcal{X} . This facilitates cross-modality learning for better VidL understanding.

In addition, we employ the commonly used VidL pre-training objectives, including video-text matching (VTM) and masked language modeling (MLM), where VTM aims to predict whether an input video-text pair is matched or not, while MLM aims to predict the masked word tokens from the surrounding context.² Our overall pre-training objective can be written as:

$$\mathcal{L} = \mathcal{L}_{\text{MVM}} + \mathcal{L}_{\text{VTM}} + \mathcal{L}_{\text{MLM}}, \quad (1)$$

where $\mathcal{L}_{\text{MVM}}, \mathcal{L}_{\text{VTM}}, \mathcal{L}_{\text{MLM}}$ are the MVM, VTM and MLM objectives, respectively.

²Refer to the Appendix for detailed formulation of VTM and MLM.

3.3. Target Features

Masked visual modeling (MVM) is a generic masked feature prediction task, where we mask out some of the visual input patches, and then predict the target features corresponding to the masked ones. Thus, a core design of MVM is the target features, which enables VIOLET learning a desired aspect of visual modeling. While MVM has been explored in pure vision tasks [4, 23, 79], it remains an open question whether MVM can facilitate the interactions between video and language modalities. In this study, we investigate *what design of MVM is effective in the context of video-language pre-training?*

Following [79, 82], we employ a simple linear layer or 2-layer MLP as the prediction head for MVM, to project the hidden video representations (h^v , of hidden size 768) from CT to the same dimension as the MVM targets. The default MVM loss is the l_1 loss, unless specified otherwise. Next, we introduce the considered target features in details.

RGB Pixel Values (Pixel). We treat the normalized RGB pixel values as a candidate target feature. During MVM, VIOLET learns to reconstruct the pixel values of the masked patches. The linear MVM head projects h^v into the same dimension as the raw video frame patch ($H \times W \times 3$).

Histogram of Oriented Gradients (HOG). HOG [10] is a pioneer feature descriptor that describes the gradients of orientations of the image. While HOG has been proven effective for visual pre-training [79], it is unknown whether it can benefit VidL pre-training. We extract HOG features in a dense grid level, and use such feature descriptors as the prediction targets of MVM. The HOG feature map is of the same size as the input video frame, but with channel size 1. The linear MVM prediction head projects h^v to the same dimension as HOG for the video frame patch ($H \times W \times 1$).

Depth Maps (Depth). Since depth maps usually contains finer-grained details of the object shapes and general scene layout of the foreground objects, it is worth exploring whether depth maps can be used to improve the scene/object understanding capability of a VidL pre-trained model. To obtain such MVM target, we employ a pre-trained dense prediction transformer (DPT) [60] to perform monocular depth estimation given an input video frame. The linear prediction head used for Depth is the same as the one for HOG, as both targets are of channel size 1.

Optical Flow (Flow). Optical flow is commonly used in motion analysis and video understanding. Here, we analyze whether apparent velocity of objects can benefit VidL pre-training. We employ a pre-trained recurrent all-pairs field transforms (RAFT) [69] to compute optical flow given the consecutive video frames. We directly use the estimated optical flow values as the prediction target, and supervise the MVM training with l_1 loss. To obtain the MVM predictions, we concatenate the hidden video representations

Pre-training Tasks	MVM Target	TGIF-Frame		DiDeMo-Retrieval		
		Acc.	R1	R5	R10	AveR
VTM+MLM	None	68.1	28.7	57.0	69.7	51.8
+MVM	RGB Pixel Values	68.3 (+0.2)	29.2 (+0.5)	58.6 (+1.6)	70.1 (+0.4)	52.6 (+0.8)
	Histogram of Oriented Gradients [10]	67.3 (-0.8)	26.6 (-2.1)	54.9 (-2.1)	68.1 (-1.6)	49.8 (-2.0)
	Depth Maps (DPT-L [60])	68.0 (-0.1)	27.3 (-1.4)	55.0 (-2.0)	68.3 (-1.4)	50.2 (-1.6)
	Optical Flow (RAFT-L [69])	67.6 (-0.5)	30.3 (+1.6)	58.0 (+1.0)	70.3 (+0.6)	52.9 (+1.1)
	Spatial-focused Image Features (Swin-B [49])	68.8 (+0.7)	35.4 (+6.7)	62.4 (+5.4)	74.9 (+5.2)	57.6 (+5.8)
	Temporal-aware Video Features (VidSwin-L [50])	68.0 (-0.1)	32.8 (+4.1)	60.5 (+3.5)	73.0 (+3.3)	55.4 (+3.6)
	Discrete Visual Tokens (DALL-E [59])	68.4 (+0.3)	28.1 (-0.6)	56.6 (-0.4)	69.4 (-0.3)	51.3 (-0.5)
Multimodal Features (CLIP-ViT-B [58])	67.7 (-0.4)	29.8 (+1.1)	57.8 (+0.8)	68.5 (-1.2)	52.1 (+0.3)	

Table 1. Comparing target features for MVM applied to video-text data. All variants are pre-trained on WebVid [3] for 5 epochs. Masking is performed randomly (RM) with ratio of 15%. The final pre-training setting is highlighted in gray.

computed by CT on consecutive frames, and employ a linear layer to project the concatenated video representations (of hidden size 768×2) to the same dimension as the estimated optical flow target for a given patch ($H \times W \times 2$).

Spatial-focused Image Features (SIF). We investigate whether image features can be useful for improving VidL pre-training. We employ a well-known vision transformer (such as Swin Transformer [49]) to extract the grid features given an input image. We then normalize the extracted grid features and consider them as ground-truth MVM targets. Likewise, we adopt a 2-layer MLP to project h_v to the same dimension as the image feature target.

Temporal-aware Video Features (TVF). We also study the impact of video features to VidL pre-training. We employ pre-trained video transformer (such as Video Swin Transformer [50]) to compute temporal-aware features for this analysis. Given a set of video frames, we use the transformer to extract video features in the form of space-time cubes, and then apply l_1 regression between normalized video features and MVM predictions from a 2-layer MLP head of the masked video patches.

Discrete Visual Tokens (VQ). In addition to continuous MVM targets, we also consider the discrete variational autoencoder (dVAE) [59, 73] to quantize video inputs. dVAE is learned to tokenize images into discrete visual tokens q from a finite dictionary, and then reconstruct the original visual scene based on q , where q should have a one-to-one correspondence with the input image patches spatially. We first adopt dVAE to tokenize the t^{th} video frame f_t into q_t : $q_t = \text{dVAE}(f_t)$, and then a 2-layer MLP is used to project h_v into the finite VQ vocabularies. As VQ token is discrete, we can model MVM with VQ as a classification problem, and adopt the cross-entropy loss to optimize the MVM training, following [4, 17].

Multimodal Features (MMF). We further study if the features learned via multimodal pre-training can benefit VidL pre-training. We utilize the vision branch of the ViT-Base backbone [13] in CLIP [58] to extract such multimodal features, and use the normalized features as the prediction targets in MVM pre-training. Again, we apply l_1 regression between the MVM predictions made via a 2-layer MLP

head and the MMF targets for the masked patches.

In the following sections, we conduct comprehensive investigation over MVM targets described above, and perform detailed analysis on MVM strategies. To avoid confusion, we denote the strongest model with the most effective MVM training as VIOLETv2.

4. Study: Target Features for MVM

Settings. We conduct pre-training on WebVid-2.5M [3] for 5 epochs, and report accuracy on TGIF-Frame [27] for video question answering and R1/R5/R10/AveR on DiDeMo [26] for text-to-video retrieval.³ We initialize our Video Swin Transformer (VT) with VideoSwin-Base [50], pre-trained on Kinetics-600 [29]. Language Embedder (LE) and Cross-modal Transformer (CT) are initialized from pre-trained BERT-Base [12]. During pre-training, we sparsely sample 4 video frames and randomly crop them into 224×224 to split into patches with $H = W = 32$. For all downstream tasks, we adopt the same video frame size and patch size but 5 sparse-sampled frames. We keep the training recipe (*e.g.*, optimizer settings, masking ratio, training schedule, *etc.*) consistent across all targets, which we find generally good in practice.⁴ For MVM targets that involve a teacher model, we use official models released by the authors. We compare models pre-trained with 8 different MVM variants to the baseline pre-trained with only VTM and MLM. Our goal is to find the best MVM target features that can provide the largest performance improvement over this baseline. Results are summarized in Table 1. We first categorize the MVM targets into 4 groups, and discuss their performance in details.

One-stage Visual Targets. We include *Pixel* and *HOG*, as they do not require training a deep neural network in advance to extract these features. Compared to the baseline without MVM objective, regressing the explicit RGB colors contributes to a relatively small gain of +0.2% on TGIF-Frame and +0.8% on AveR for DiDeMo Retrieval. In

³We base our ablation experiments on these two representative datasets for fast iteration, our main results are reported on 13 benchmarks in Section 6. Details about downstream adaptation are included in the Appendix.

⁴Refer to the Appendix for more on training details.

MVM Targets	TGIF-Frame	DiDeMo-Retrieval			
	Acc.	R1	R5	R10	AveR
Pixel	68.3	29.2	58.6	70.1	52.6
Flow	67.6	30.3	58.0	70.3	52.9
SIF	68.8	35.4	62.4	74.9	57.6
SIF + Pixel	68.8	31.8	60.4	73.0	55.1
SIF + Flow	68.7	34.4	61.5	72.8	56.3

Table 2. Combining MVM targets. All variants are pre-trained on WebVid [3] for 5 epochs, using RM with 15% as the masking strategy. We highlight the final setting in gray.

contrast, HOG renders degradation on downstream video-language (VidL) performance (-0.8% on TGIF-Frame and -2.0% on DiDeMo-Retrieval). We hypothesize that this is due to the missing color information in HOG features, which is critical in VidL understanding.

Supervised Pseudo-label Targets. We include *Depth Maps (Depth)* and *Optical Flow (Flow)*. Intuitively, Depth and Flow can be considered as continuous pseudo “labels”, which are made by models trained to perform depth and optical flow estimation [60, 69]. Depth does not improve over baseline with VTM+MLM. The nature of depth maps are to separate the foreground from the background, thus may guide the model to ignore information from the background, even when they are relevant for solving downstream VidL tasks (-0.1% on TGIF-Frame, -1.6% on DiDeMo Retrieval). Flow only focuses on the moving part between frames, while ignores the spatial details of static components, thus fail on more spatially-focused TGIF-Frame task (-0.5%). We also find that the optical flow estimation model easily fails with sparse sampling strategy, which is widely adopted in VidL pre-training.⁵

Supervised Visual Feature Targets. We include continuous features extracted from the last layers of image classification model [49] (*i.e.*, *Spatial-focused Image Features (SIF)*) and action recognition model [50] (*i.e.*, *Temporal-aware Video Features (TVF)*). We consider regressing supervised features from Swin-B or VidSwin-L⁶ as a type of knowledge distillation from unimodal models to our model. SIF achieves significant improvement over baseline (+0.7% on TGIF-Frame and +5.8% on AveR for DiDeMo-Retrieval). In contrast, TVF fails to improve TGIF-Frame accuracy (-0.1%), though it brings notable improvement on retrieval performance (+3.6% on AveR). By distilling the knowledge from Swin-B, we enforce the model to focus more on spatial details of each frame, which we hypothesize is the main reason behind the large performance improvement. As previous study [5] pointed out, existing VidL benchmarks largely test on spatial understanding about the key frame of the video, with only a fractional of examples actually testing on temporal reasoning over multiple frames.

Self-supervised Multimodal Feature Targets. We use

⁵Please find visualization examples in the Appendix.

⁶VidSwin-L is trained on Kinetics-400 [29] with 83.1% accuracy.

Image Features Model	Train Data	IN-1K	TGIF-Frame	DiDeMo-Retrieval			
		ACC@1	Acc.	R1	R5	R10	AveR
ResNet-50 [24]	IN-1K	76.1	67.3	29.1	58.1	69.3	52.2
Swin-T [49]	IN-1K	81.2	68.9	33.8	63.6	74.2	57.2
Swin-B	IN-1K	83.5	68.3	34.9	63.4	73.9	57.4
Swin-B	IN-22K	85.2	68.8	35.4	62.4	74.9	57.6
Swin-L	IN-22K	86.3	68.2	33.2	62.4	72.6	56.1

Table 3. Comparing different image feature targets for MVM. All variants are pre-trained on WebVid [3] with VTM+MLM+MVM (SIF) for 5 epochs, using RM with 15% as the masking strategy. The final pre-training setting is highlighted in gray.

Discrete Visual Tokens (VQ) from DALL-E [59] and continuous *Multimodal Features (MMF)* extracted from CLIP [58]. Both are pre-trained on large-scale image-text datasets, usually much more expensive than all other targets. Both targets improve the performance by a slight margin on only one task. VQ that can capture patch-level semantics, benefits TGIF-Frame (+0.3%) which mostly focuses on scene understanding. While MMF from CLIP, contrastively pre-trained to measure the high-level similarity between the entire image and text sentence, is helpful for DiDeMo-Retrieval (+0.3% on AveR).

Summary. We hypothesize that many factors could lead to the low performance of an MVM target, such as its own characteristics (*e.g.*, local vs. global features); the target model; the loss design; or the mismatch between pre-train objectives and downstream focus. We try our best to compare them rigorously with controlled experiments to find the best setting. Based on our experiments, regressing RGB values (Pixel) and distilling features from Swin-B [49] (SIF) are the only two that produce consistent gains over the baseline on both downstream tasks. MVM with SIF achieves the best performance, with a gain of +0.7% on TGIF-Frame and +5.8% on AveR for DiDeMo-Retrieval over the baseline. Therefore, we use SIF as the default target for MVM in the following sections, unless specified otherwise.

5. Analyses of MVM

Combining MVM Targets. As different MVM targets focus on different aspects of visual modeling, a naive way to enable models with different visual capabilities is to combine them together. Specifically, the model pre-training can be supervised by more than one MVM loss, which is simply added together to be backpropagated. In Table 2, we find there is no merit in combining different MVM targets, leading to worse downstream performance than using SIF alone. When combining the best two targets found in Table 1: Pixel+SIF, it performs better than Pixel only, but does not improve over using SIF alone. We hypothesize that the explicit details of pixel values may conflict with the high-level visual semantics summarized in the grid features from the image classifier. We further try to combine SIF with Flow in the hope of enforcing both temporal and spatial reasoning over video inputs. In addition, Flow is a better candidate than other targets, as it demonstrates some advantages

Masking Strategy	Time Cost hours	TGIF-Frame	DiDeMo-Retrieval			
		Acc.	R1	R5	R10	AveR
RM	8.0	68.8	35.4	62.4	74.9	57.6
BM	8.0	69.0	35.9	63.3	74.6	57.9
AM	34.5	68.4	31.5	59.9	72.0	54.7
RM+BM	8.0	68.7	36.4	64.2	74.4	58.3
RM+AM	20.5	68.8	33.7	63.2	73.5	56.8
BM+AM	20.5	68.9	35.6	61.9	74.4	57.3
RM+BM+AM	17.0	68.6	34.7	62.0	74.8	57.2

Table 4. Impact of **masking strategy of MVM**. All variants are pre-trained on WebVid [3] with VTM+MLM+MVM (SIF) for 5 epochs. The masking ratio is set as 15% for all masking strategies. The final pre-training setting is highlighted in gray.

on retrieval performance in Table 1, and it is a different type of target from SIF, compared to temporal-aware video features. The results are consistent, with improvements over optical flow only; while the performance drops, compared to SIF alone. Though our results are not encouraging, we believe how to effectively combine different MVM targets is an interesting direction for future study.

MVM Target Extractors vs. Downstream Performance.

In Table 3, we explore different image classification models as the MVM target extractor for SIF, and investigate whether stronger image classification model enables better VidL performance. We compare ResNet-50 [24], Swin-Tiny/Base/Large [49], trained on ImageNet-1K (IN1K) or ImageNet-22K (IN-22K) [11], and summarize the observations below:

- ResNet-50 performs lower than Swin variants. Two potential reasons are (i) ResNet-50 architecture is very different from VidSwin (*i.e.*, with different inductive bias); and (ii) the much lower ImageNet performance (~ 76 vs. > 81) suggest the ResNet-50 features are not as strong.
- When the target model shares *similar inductive biases* to the video encoder (*i.e.*, Swin-T/B/L), the downstream performance is *not directly proportional* to ImageNet accuracy, and is overall better than that of Res50. This suggests that the architecture design of both target model and video encoder should be similar.
- A key difference between different Swin targets is the feature dimension (768/1024/1568 for Swin-T/B/L), while the video tokens from CT are of size 768. Although we project them into the same dimension as the targets, the mismatch may lead to slightly lower performance (with Swin-L especially).

In short, we believe a SIF target model should share similar inductive biases as the video encoder.

Masking Strategy. We investigate the effect of different masking strategies in Table 4, including random masking (RM), blockwise masking (BM), attended masking (AM), and their combinations.

- **Random Masking (RM).** Following the conventional practice in MLM, we randomly select a certain percentage p_m of video frame patches from the whole video inputs to be masked. In Table 5, we explore different mask-

p_m	TGIF-Frame	DiDeMo-Retrieval			
	Acc.	R1	R5	R10	AveR
15%	68.8	35.4	62.4	74.9	57.6
30%	68.8	36.2	64.0	74.5	58.2
45%	68.9	35.6	61.9	74.4	57.3
60%	68.1	34.1	63.9	74.6	57.5
75%	68.3	35.4	62.4	74.2	57.3

Table 5. Impact of **masking ratio of MVM**. All variants are pre-trained on WebVid [3] with VTM+MLM+MVM (SIF) for 5 epochs, using RM as the masking strategy. The final pre-training setting is highlighted in gray.

ing ratios (p_m), and empirically find $p_m = 30\%$ gives the best downstream performance.

- **Blockwise Masking (BM).** To make MVM relying less on similar neighbor patches, we adopt blockwise masking [4, 68] that masks blocks of video patches along spatial-temporal dimension rather than independently masking randomly sampled patches for each frame. Specifically, we randomly sample an (H', W', T') as a masking block, where all $H' \times W'$ visual patches in the following T' consecutive frames will be masked; we repeat this process until $> p_m$ of video patches are masked to perform MVM pre-training.
- **Attended Masking (AM).** Attended masking tries to put more weights on the more important elements based on the attention weights computed by Cross-modal Transformer (CT). A similar idea has been explored in [88] for MLM. Here, we extend AM to both visual and textual modalities. We first keep the video-text inputs intact, feed them into CT to compute the attention weights, to decide which portions in video and text are more important. We then select the top p_m of most-attended patches/tokens to be masked in video-text inputs for MVM and MLM.

To combine different masking strategies, we randomly apply one masking method for each video-text pair in a batch. Results in Table 4 suggest that TGIF-Frame can slightly benefit from BM, and combining BM with RM leads to the best retrieval performance on DiDeMo. As video usually presents analogous visual patterns in spatial-temporal neighbors (*i.e.*, nearby patches within current frame or neighboring frames), when masking patches independently (*i.e.*, RM), these neighbors can make the masked patches easy to recover, and may lead to spurious success in MVM evaluation. By masking a block (*i.e.*, BM) instead of individual patches, the model cannot merely rely on similar neighboring visual cues but requires actual visual reasoning to recover a group of missing patterns. Combining BM with RM leads to more diverse dropout patterns in video inputs, which is in analogy to data augmentation.

In addition, AM and combinations with AM are not effective for both downstream tasks. It is also worth noting that AM greatly increase the training time (4 times more than RM/BM), due to the additional forward pass needed

Pre-training Tasks	MVM Target	TGIF-Frame	DiDeMo-Retrieval			
		Acc.	R1	R5	R10	AveR
ITM+MLM	None	69.8	36.4	64.3	74.7	58.4
+MVM	RGB Pixel Values	69.7 (-0.1)	35.8 (-0.6)	64.4 (+0.1)	74.9 (+0.2)	58.4
	Histogram of Oriented Gradients [10]	69.8	34.9 (-1.5)	64.4 (+0.1)	75.1 (+0.4)	58.1 (-0.3)
	Depth Maps (DPT-L [60])	69.6 (-0.2)	32.3 (-4.1)	63.8 (-0.5)	74.2 (-0.5)	56.9 (-1.5)
	Spatial-focused Image Features (Swin-B [49])	69.7 (-0.1)	31.6 (-4.8)	60.5 (-3.8)	72.5 (-2.2)	54.9 (-3.5)
	Discrete Visual Tokens (DALL-E [59])	69.8	34.4 (-2.0)	62.6 (-1.7)	75.1 (+0.4)	57.4 (-1.0)
	Multimodal Features (CLIP-ViT-B [58])	69.8	33.6 (-2.8)	62.9 (-1.4)	75.6 (+0.9)	57.4 (-1.0)

Table 6. Comparing target features for MVM applied to image-text data. All variants are pre-trained on CC3M [65] for 5 epochs. Masking is performed randomly (RM) with ratio of 15%.

Pre-training Tasks	MVM Target		TGIF-Frame	DiDeMo-Retrieval			
	WebVid2.5M	CC3M	Acc.	R1	R5	R10	AveR
VTM+MLM	None	None	69.7	36.7	66.5	76.6	59.9
+MVM	Spatial-focused Image Features (Swin-B [49])	None	71.1	38.8	69.6	80.0	62.8
	Spatial-focused Image Features (Swin-B)	Pixel	71.3	39.7	69.3	78.4	62.5

Table 7. Combining MVM target features for both video-text and image-text data. All variants are pre-trained on WebVid2.5M [3] +CC3M [65] for 5 epochs. The final pre-training setting is highlighted in gray .

to compute the attention weights. In our implementation, we optimize the three losses altogether in the same forward-backward pass. Hence, the performance drop with AM may be due to the important elements (*e.g.*, visual patches containing the main object or content words) are more likely to be masked together and leaving the less relevant elements (*e.g.*, scene background or stop words) intact, which will especially make the learning of video-text matching harder.

Applying MVM to Image-Text Data. As image can be considered as a special case of video with temporal size 1, video-language (VidL) pre-training can take advantages of image-text data, which has been proven successful in [3,36]. The current trend in VidL pre-training is to leverage both video-text data and image-text data. Therefore, we repeat the experiments in Section 4 and examine which MVM targets work the best on downstream VidL tasks, when pre-trained on image-text data only. We remove optical flow and temporal-aware video features from this study, as the inputs are static images. In Table 6, we pre-train our model on CC3M [65] for 5 epochs and report results on TGIF-Frame and DiDeMo-Retrieval. The performance trend with different MVM targets are not consistent with that observed on video-text data. Pixel is able to largely preserve the baseline (VTM+MLM) performance, while other MVM targets lead to different degrees of performance drop, especially on retrieval. Without visual implications from neighbor frames as video, MVM is more challenging to learn on image data. On the other hand, MVM over an image may easily fit in static visual representation, which could hurt video temporal reasoning and not benefit downstream VidL learning.

Combining Video-Text Data with Image-Text Data. We further follow recent VidL literature [3,40] to use both video-text data and image-text data for pre-training, and investigate different ways to combine MVM targets on image and video data in Table 7. Note that we adopt the best training strategy found in the above investigations, that is, using spatial-focused image feature (SIF) as MVM target for

video inputs, and using blockwise masking (BM) + random masking (RM) with masking ratio of 30% as the masking strategy. As the best MVM target (Pixel) on image data does not show improvement over the baseline without MVM objective in Table 6, we explore with/without MVM objective on images in this combined pre-training. For the baseline with VTM+MLM only, we simply remove the MVM objective on both image and video data, while keeping the rest training settings. Under the strict fair comparison, we observe adding MVM objectives contributes to $>+0.4\%$ gains on TGIF-Frame and $>+2.6\%$ increase on AveR for DiDeMo-Retrieval. Comparing with or without MVM objective on images, adding MVM on image-text brings minor performance difference (+0.2% on TGIF-Frame and degrades by -0.3% on AveR for DiDeMo-Retrieval) over MVM on video-text only. Therefore, in our final setting, we only apply MVM objective on video data.

6. Main Results

To this end, we combine the most effective MVM strategies to pre-train VIOLETv2 and evaluate on 13 video-language (VidL) tasks. Table 8 shows the comparison to prior arts on **video question answering (QA)** and **video captioning**. We observe that VIOLETv2 is effective in learning transferable knowledge for the downstream tasks. For example, considering pre-training data at a similar scale (*i.e.*, $\leq 5M$, the top rows of Table 8), VIOLETv2 achieves better results than prior arts, including ALPRO [40], ClipBERT [36], and SwinBERT [44], across all considered video QA and video captioning benchmarks. Specifically, when pre-training with the exact same data (*i.e.*, WebVid2.5M [3] + CC3M [65]), VIOLETv2 surpasses ALPRO by 2.4% accuracy on MSRVT-QA and 8.4% accuracy on MSVD-QA, respectively. We also compare with other models pre-trained on significantly larger scale of video-text pairs. As shown in the bottom rows of Table 8, although

Method	# Pretrain videos/images	TGIF [27]			MSRVTT [84]		LSMDC [71]		MSVD [7]	Captioning	
		Act.	Trans.	Frame	MC [87]	QA [83]	MC	FiB	QA [83]	MSRVTT	MSVD
ClipBERT [36]	0.2M	82.8	87.8	60.3	88.2	37.4	-	-	-	-	-
ALPRO [40]	5M	-	-	-	-	42.1	-	-	46.3	-	-
SwinBERT [44]	-	-	-	-	-	-	-	-	-	53.8	120.6
<i>Models pre-trained on more data</i>											
JustAsk [85]	69M	-	-	-	-	41.5	-	-	46.3	-	-
MERLOT [88]	180M	94.0	96.2	69.5	90.9	43.1	81.7	52.9	-	-	-
All-in-one [75]	283M	95.5	94.7	66.3	92.3	46.8	84.4	-	48.3	-	-
MV-GPT [64]	53M	-	-	-	-	41.7	-	-	-	60.0	-
VIOLET [17]	186M	92.5	95.7	68.9	91.9	43.9	82.8	53.7	47.9	-	-
VIOLETv2	5M ⁷	94.8	99.0	72.8	97.6	44.5	84.4	56.9	54.7	58.0	139.2

Table 8. Comparison with SOTA on **video question answering** (accuracy) and **video captioning** (CIDEr). VIOLETv2 is pre-trained on WebVid2.5M [3]+CC3M [65] with VTM+MLM+MVM (SIF on videos) for 10 epochs. We gray out methods that use significantly more pre-training data.

Method	# Pretrain videos/images	MSRVTT [84]			DiDeMo [26]			LSMDC [62]		
		R1	R5	R10	R1	R5	R10	R1	R5	R10
ClipBERT [36]	0.2M	22.0	46.8	59.9	20.4	48.0	60.8	-	-	-
Frozen [3]	5M	31.0	59.5	70.5	31.0	59.8	72.4	15.0	30.8	39.8
ALPRO [40]	5M	33.9	60.7	73.2	35.9	67.5	78.8	-	-	-
BridgeFormer [22]	5M	37.6	64.8	75.1	37.0	62.2	73.9	17.9	35.4	44.5
<i>Models pre-trained on more data</i>										
HERO [41]	136M	16.8	43.4	57.7	-	-	-	-	-	-
All-in-one [75]	138M	37.9	68.1	77.1	32.7	61.4	73.5	-	-	-
Clip4Clip [52]	400M	42.1	71.9	81.4	43.4	70.2	80.6	21.6	41.8	49.8
VIOLET [17]	186M	34.5	63.0	73.4	32.6	62.8	74.7	16.1	36.6	41.2
VIOLETv2	5M	37.2	64.8	75.8	47.9	76.5	84.1	24.0	43.5	54.1

Table 9. Comparison with SOTA on **text-to-video retrieval** tasks (R1/5/10). We gray out methods that use significantly more pre-training data.

we use less pre-training data than others, VIOLETv2 still achieves comparable or better performance.

We observe similar findings on video captioning. On MSRVTT captioning, VIOLETv2 is only 2 points behind MV-GPT [64] pre-trained with 53M video-text pairs, which is 10 times larger than ours (5M). In addition, MV-GPT leverages ASR transcripts to enhance the captioning performance, while our captioning model takes only video frames as inputs and outputs the video caption.⁸ We believe augmenting VIOLETv2 with additional modalities, such as audio or ASR transcripts, can further improve captioning performance, which we leave as future work.

Table 9 presents the comparison on **text-to-video retrieval**. When pre-training with the same datasets (*i.e.*, WebVid2.5M [3] + CC3M [65]), VIOLETv2 shows across-the-board improvements with all metrics considered on DiDeMo and LSMDC. It is worth noting that our method performs comparably to BridgeFormer [22] on MSRVTT-Retrieval. BridgeFormer adopts a noun/verb masking strategy during pre-training, which is specially aligned to the simple sentences in MSRVTT. However, it cannot show similar effects on DiDeMo and LSMDC due to more complex texts with multiple nouns/verbs. In contrast, the studied MVM can achieve a comprehensive enhancement in VidL learning and lead to notable improvements (+10.9% R1 on DiDeMo and +6.1% R1 on LSMDC).

⁷The SIF target model Swin-B is trained on IN-22K.

⁸Details about downstream finetuning on captioning are in Appendix.

Direct Comparison to VIOLET [17]. Across Table 8 and 9, it is worth noting that VIOLETv2 outperforms VIOLET with notable margins, even when VIOLET is pre-trained with significantly more data (about 37 times more). Specifically, VIOLETv2 yields an average gain of +3.4% across 8 video QA datasets, and an absolute gain of +8.6% on R1 across all three retrieval benchmarks. These results suggest the importance of an appropriate MVM setting, which is the core belief in our study.

7. Conclusion

We initiate the first empirical study on adopting masked visual modeling (MVM) for video-language (VidL) learning. We explore diverse MVM objectives upon end-to-end ViDeO-LanguagE Transformer (VIOLETv2), including low-level pixel space, high-level visual semantics, and extracted latent features. Our results show that VIOLETv2 pre-trained on 5M video/image-text data with MVM objective achieves strong performance on 3 popular VidL tasks across 13 VidL benchmarks. Our comprehensive analyses on different combinations of MVM targets, various SIF target extractors, and varying masking strategies/ratios shed light on effective MVM design. We believe our study can guide future research on large-scale VidL pre-training and wish to study how MVM can generalize to larger-scale data. In addition, we vision that with the emergence of video/VidL foundation models in future works, better choices of MVM targets can be explored.

A. Additional Results

Fair Comparison to VIOLET [17]. VIOLET proposes to augment VTM+MLM with masked visual token modeling for VidL pre-training, while only showing marginal improvements on downstream performance. In contrast, we conduct comprehensive investigations across different MVM targets and masking strategies to demonstrate that effective MVM training can largely improve downstream performance. Note that our study already encompasses the design of MVM in [17], that is MVM with VQ target and BM+AM as the masking strategy. To make a fair comparison between [17] and our best setting (MVM with SIF target and BM+AM as the masking strategy), we reproduce [17] under the same setting and report downstream performance in Table 10. Results show that our setting obtains a significant improvement, with +2.3% on TGIF-Frame and +13.3% on AveR for DiDeMo-Retrieval, respectively. These results suggest the importance of an appropriate MVM setting, which is the core belief in our study.

Method	MVM Target	TGIF-Frame		DiDeMo-Retrieval			
		Acc.	R1	R5	R10	AveR	
VIOLET [17]	VQ	70.5	32.9	63.0	74.5	56.8	
VIOLETv2	SIF	72.8	47.9	76.5	84.1	69.5	

Table 10. Fair comparison to VIOLET [17]. Both models are pre-trained on WebVid [3]+CC [65].

MVM vs. Temporal self-supervised objectives. In addition to the reconstructive MVM task, other self-supervised video modeling tasks can be explored, for example, Frame Order Modeling (FOM) [41, 88], which reconstructs the temporal orders of shuffled frame inputs. In Table 11, we compare MVM (SIF) with FOM in [88], when pre-trained on WebVid [3]. MVM (SIF) still leads to better performance, with a gain of +0.7% on TGIF-Frame and +4.8% on AveR for DiDeMo-Retrieval, respectively.

VTM+MLM+	TGIF-Frame		DiDeMo-Retrieval			
	Acc.	R1	R5	R10	AveR	
MVM (SIF)	68.8	35.1	63.3	73.1	57.2	
FOM	68.1	27.6	59.0	70.7	52.4	

Table 11. MVM vs. FOM. Both models are pre-trained on WebVid [3].

Initialization and Learning of Video Backbone. We investigate the effect of different initialized video backbones with or without MVM on VL inputs in Table 12. At first, although the used video transformer (VT) is randomly initialized, the MVM training still enhances the visual representation and benefits the downstream video-language (VidL) tasks. Furthermore, MVM can also boost better initialized VT from VidSwin-B and lead to a comprehensive increase. Specifically, the improvement gap is more significant than random initialization, where we can learn better from MVM and enlarge its effectiveness during pre-training. We additionally compare two self-supervised initializations with MVM on video-only inputs, one with TVF as the MVM tar-

get and the other with SIF. Though VidL pre-training with MVM from supervised VidSwin-B initialization leads to the best downstream performance, we observe consistent performance improvement from MVM on VL inputs regardless of the initialization setting.

Weight Init.	MVM on VL	TGIF-Frame		DiDeMo-Retrieval			
		Acc.	R1	R5	R10	AveR	
Random	×	55.9	5.6	19.9	29.8	18.5	
	✓	56.5	7.4	22.9	33.8	21.4	
V-only MVM (TVF)	×	59.9	15.3	38.4	54.7	36.1	
	✓	60.2	17.4	43.2	56.0	38.9	
V-only MVM (SIF)	×	61.0	16.9	42.4	54.9	38.1	
	✓	61.5	18.6	44.0	58.1	40.2	
VidSwin-B	×	68.1	28.7	57.0	69.7	51.8	
	✓	68.8	35.1	63.3	73.1	57.2	

Table 12. Impact of weight initialization and learning of video backbone. All variants are pre-trained on video-text from WebVid [3] for 5 epochs. The MVM target is spatial-focused image features (SIF) from Swin-B [49]), if not specified otherwise. For V-only MVM (TVF/SVF), we first self-supervisedly pre-train the video backbone with MVM on video-only inputs from WebVid for 5 epochs. The final pre-training setting is highlighted in gray.

Type of MVM Loss. We compare the type of loss function for the MVM training by using least absolute deviations (l_1) or least square errors (l_2) in Table 13. It is well known that the l_1 loss can be resistant to outlier data. We show that MVM through l_1 is also more robust and leads to better performance on both video question answering and text-to-video retrieval than the l_2 loss.

MVM Loss	TGIF-Frame		DiDeMo-Retrieval			
	Acc.	R1	R5	R10	AveR	
l_1	68.8	35.4	62.4	74.9	57.6	
l_2	68.8	33.0	60.1	71.9	55.0	

Table 13. Impact of MVM loss type. All variants are pre-trained on WebVid [3] with VTM+MLM+MVM (SIF) for 5 epochs, using RM as the masking strategy with ratio of 15%. The final pre-training setting is highlighted in gray.

MVM Prediction Head. We investigate the prediction head for MVM in Table 14. As a result, a single linear layer is not enough to model the complicated distilling MVM features. (e.g., 31.3 vs. 35.4 R1 on DiDeMo-Retrieval) Therefore, we follow VTM and MLM to use 2-layer MLP as the prediction head for MVM.

MVM Head	TGIF-Frame		DiDeMo-Retrieval			
	Acc.	R1	R5	R10	AveR	
1 Linear Layer	68.8	31.3	60.1	72.8	54.7	
2-layer MLP	68.8	35.4	62.4	74.9	57.6	

Table 14. Impact of MVM prediction head. All variants are pre-trained on WebVid [3] with VTM+MLM+MVM (SIF) for 5 epochs, using RM as the masking strategy with ratio of 15%. The final pre-training setting is highlighted in gray.

TVF Target Extractors vs. Downstream Performance. We compare distilling video features from VidSwin-B vs. VidSwin-L (the default setting in the main text) in Table 15.

Here, for experiments with VidSwin-B, the same VidSwin-B weight is used to initialize the video backbone and to extract the MVM target. Hence, the MVM objective can be easily minimized by simply ignoring the text inputs, which conflict with the other objectives. This variant is in principle similar to masked frame modeling in HERO [41], the key difference lies in whether the video backbone is refined during pre-training. In addition, we investigate whether the sparse sampling of video frames when extracting TVF target is the key reason behind the lower performance of TVF, compared to SIF. Hence, we compare the default sparse sampling of 5 frames, against a dense-version of TVF target (feeding 16 frames into VidSwin-L). While the dense input is slightly beneficial, SIF still performs better, with absolute advantages of +0.4% on TGIF-Frame and +1.8% on AveR for DiDeMo-Retrieval.

MVM Target	TGIF-Frame		DiDeMo-Retrieval		
	Acc.	R1	R5	R10	AveR
TVF (VidSwin-L [50])	68.0	32.8	60.5	73.0	55.4
TVF (VidSwin-B)	67.5	25.8	55.0	68.0	49.6
TVF-dense (VidSwin-L)	68.4	34.3	60.8	72.4	55.8

Table 15. Temporal-aware video feature (TVF) target models vs. downstream performance. All variants are pre-trained on WebVid [3] with VTM+MLM+MVM (TVF) for 5 epochs, using RM as the masking strategy with ratio of 15%.

Additional Exploration in Combining MVM Targets.

We explore the additional combination of distilling MVM targets in Table 16. MVM with SIF has an obvious advantage over TVF only on both video question answering and text-to-video retrieval. While, considering SIF+TVF seems not to bring a robust improvement, especially decreasing text-to-video retrieval. The previous study [5] shows that current VidL benchmarks primarily focus on spatial understanding of the key frame from videos. Furthermore, combining TVF with SIF results in excessive training overhead. Accordingly, we choose SIF as our final pre-training setting.

MVM Targets	TGIF-Frame		DiDeMo-Retrieval		
	Acc.	R1	R5	R10	AveR
SIF	68.8	35.4	62.4	74.9	57.6
TVF	68.0	32.8	60.5	73.0	55.4
SIF + TVF	69.2	33.8	63.0	74.4	57.1

Table 16. Combining target features for MVM. All variants are pre-trained on WebVid [3] for 5 epochs. The final pre-training setting is highlighted in gray.

Additional Exploration of SIF Target. We explore a more advanced SIF target, DeiT [72], in Table 17. These results show that Swin-B still has an advantage (a noticeable higher 34.9 R1 on retrieval), consistent with our previous observations in the main text. That is, SIF should share similar inductive biases to the video encoder (i.e., Swin-T/B/L).

Investigation of Training Recipe with CLIP Target. We presented initial results for varying training settings using CLIP/Swin-B targets and compare them with the default

Image Feat. Model	Train Data	IN-1K	TGIF-Frame	DiDeMo-Retrieval		
		ACC@1	Acc.	R1	R5	R10
ResNet-50	IN-1K	76.1	67.3	29.1	58.1	69.3
DeiT [72]	IN-1K	83.4	68.4	31.4	59.4	72.2
Swin-B	IN-1K	83.5	68.3	34.9	63.4	73.9

Table 17. Comparing Swin-B vs. another SIF model (DEiT) All variants are pre-trained on WebVid with VTM+MLM+MVM (SIF) for 5 epochs, using RM with 15% as the masking strategy.

setting in Table 1. Swin-B had a significant advantage over CLIP. As we adjust the training recipe with CLIP target in Table 18, a better training recipe reduces the performance gap between the SIF target and the CLIP target. The results in turn suggest the importance of an effective MVM strategy (e.g., masking ratio). Though impossible to iterate over all settings, Swin-B remains competitive under the same training recipe, especially with limited training data (IN-22K vs. 400M). Note that we use CLIP image features as the MVM target, while other related works [45, 55] use them as model inputs. One potential enhancement is to leverage the multimodal information from both image and text encoders in CLIP (similar to the use of both in [55]), which is an interesting direction to explore in future studies. However, our setup is still valid, as we aim to train a fusion-encoder architecture rather than a dual-encoder architecture as CLIP.

MVM target	Settings	TGIF-Frame	DiDeMo-Retrieval		
		Acc.	R1	R5	R10
CLIP	Default	67.7	29.8	57.8	68.5
Swin-B		68.8	35.1	63.3	73.1
	$lr \times 2$	70.5	32.9	61.6	73.5
	masking ratio = 0.3	68.0	31.8	59.6	71.3
CLIP	loss type = l_2	68.3	30.1	59.1	71.0
	linear MVM head	68.2	30.5	58.3	69.2
Swin-B	$lr \times 2$	70.6	33.3	63.7	75.2
	masking ratio = 0.3	68.8	36.2	64.0	74.5

Table 18. Investigation of Training Recipe with CLIP Target. All variants are pre-trained on WebVid with VTM+MLM+MVM for 5 epochs. The default setting follows Table 1 in the main text, that is RM with 15% masking ratio, l_1 loss and 2-layer MLP head for MVM prediction.

Extended Results for Table 1. The additional results on MSVD-QA and MSRVT- Retrieval in Table 19 show that SIF is still the most effective, consistent with Table 1. However, the effects of different MVM targets seem to be better on MSVD-QA and MSRVT- Retrieval (on average 15s long) than those on TGIF-Frame (on average 3s long) and DiDeMo-Retrieval (on average 30s long). We hypothesize that different downstream video lengths may contribute to different performance gains/losses when evaluating the effectiveness of an MVM target, which we leave as future directions for investigation.

Qualitative Results. Figure 2 shows good and bad examples of optical flow predictions made by RAFT-L [69] with sparsely sampled frames. As shown in 2b, the top example shows zoom-in shots, and the bottom shows moving shots. All content in the current frame is moving, which is the

Pre-training Tasks	MVM Target	MSVD-QA	MSRVTT-Retrieval			
		Acc.	R1	R5	R10	AveR
VTM+MLM	None	49.2	26.0	56.6	69.4	50.7
+MVM	RGB Pixel Values	51.0 (+1.8)	27.4 (+1.4)	58.0 (+1.4)	69.8 (+0.4)	51.7 (+1.0)
	Histogram of Oriented Gradients [10]	50.1 (+0.9)	27.4 (+1.4)	57.7 (+1.1)	70.2 (+0.8)	51.8 (+1.1)
	Depth Maps (DPT-L [60])	50.3 (+1.1)	28.0 (+2.0)	57.4 (+0.8)	70.6 (+0.8)	52.0 (+1.3)
	Optical Flow (RAFT-L [69])	49.7 (+0.5)	25.8 (-0.2)	55.8 (-0.8)	69.4	50.3 (-0.4)
	Spatial-focused Image Features (Swin-B [49])	51.1 (+1.9)	29.4 (+3.4)	59.9 (+3.6)	73.1 (+3.7)	54.1 (+3.4)
	Temporal-aware Video Features (VidSwin-L [50])	49.8 (+0.6)	29.9 (+3.9)	58.1 (+1.5)	70.2 (+0.8)	52.7 (+2.0)
	Discrete Visual Tokens (DALL-E [59])	50.7 (+1.5)	27.3 (+1.3)	58.3 (+1.7)	70.0 (+0.6)	51.9 (+1.2)
Multimodal Features (CLIP-ViT-B [58])	50.2 (+1.0)	30.0 (+4.0)	58.8 (+2.2)	71.1 (+1.7)	53.3 (+2.6)	

Table 19. Comparing target features for MVM applied to video-text data. All variants are pre-trained on WebVid [3] for 5 epochs. Masking is performed randomly (RM) with a ratio of 15%. The final pre-training setting is highlighted in gray .

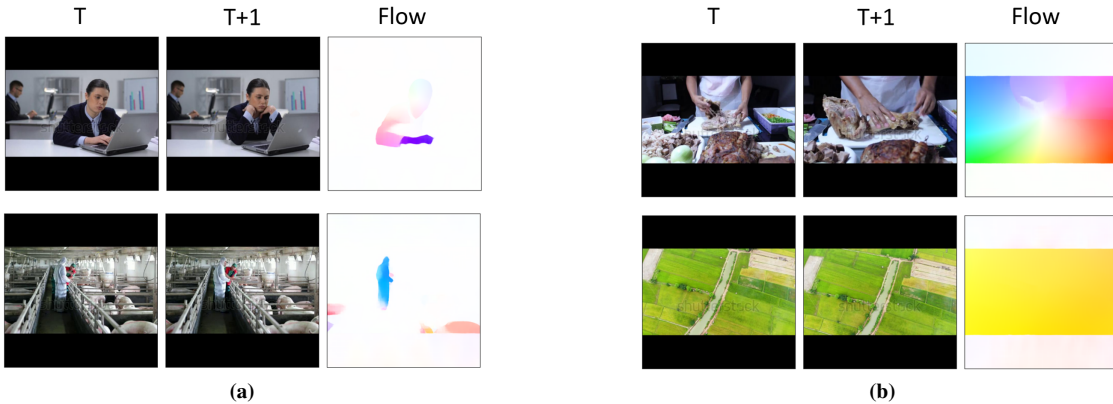


Figure 2. Visualization of optical flow (Flow) predictions by RAFT-L [69] with sparsely sampled frames. We show examples of good cases in (a) and bad cases in (b).

main reason behind the failure in optical flow estimation.

We also show the visualizations of zero-shot text-to-video retrieval on MSRVTT (Figure 3), DiDeMo (Figure 4), and LSMDC (Figure 5) to demonstrate that MVM can help video understanding from different domains, such as gaming, animation, human activity, or movie scene.

B. Additional Pre-training Details

Video-Text Matching (VTM). VTM enhances the cross-modal fusion via modeling the alignments between visual and textual inputs. At each training step, we randomly replace the corresponding text \mathcal{X}_{pos} for a given video \mathcal{V} with the text description \mathcal{X}_{neg} from a different video in the same batch. Both the positive pair $(\mathcal{V}, \mathcal{X}_{\text{pos}})$ and negative pair $(\mathcal{V}, \mathcal{X}_{\text{neg}})$ are modeled by Cross-modal Transformer (CT), and VTM is to tell them apart from the global VidL representation h^c of the [CLS] token. In particular, h^c will be processed by a fully-connected layer (FC^{VTM}) to learn contrastively through classification:

$$b^{\text{pos}} = \text{FC}^{\text{VTM}}(h_{\text{pos}}^c), b^{\text{neg}} = \text{FC}^{\text{VTM}}(h_{\text{neg}}^c),$$

$$\mathcal{L}_{\text{VTM}} = -\frac{1}{B} \sum_i \log \frac{b_i^{\text{pos}}}{b_i^{\text{pos}} + \sum_i b_i^{\text{neg}}}, \quad (2)$$

where h_{pos}^c or h_{neg}^c is h^c of positive or negative pairs.

Masked Language Modeling (MLM). In MLM, we ran-

domly mask out some word tokens with a probability of 15%.⁹ The goal is to recover these masked word tokens x from the joint VidL features h modeled by CT. Specifically, the corresponding h^x for these masked tokens are fed in a fully-connected layer (FC^{MLM}) and projected to the discrete token space for classification:

$$x'_i = \text{FC}^{\text{MLM}}(h_i^x),$$

$$\mathcal{L}_{\text{MLM}} = -\mathbb{E} \left[\frac{1}{|\mathcal{M}^{\text{MLM}}|} \sum_{i \in \mathcal{M}^{\text{MLM}}} \log P(x_i | x'_i) \right], \quad (3)$$

where \mathcal{M}^{MLM} denotes the index set of masked word tokens.

Implementation Details. Our VIOLETv2 is implemented based on PyTorch [56]. As discussed in the main text and supported by the additional experimental results above, our final pre-training setting is (i) VTM+MLM+MVM (with MVM target as spatial-focused image features from Swin-B [49], applied on video-text inputs only) as the pre-training tasks; (ii) 2-layer MLP as the MVM prediction head and l_1 regression as the MVM loss; and (iii) blockwise masking + random masking with masking ratio of 30% as the masking strategy. We adopt AdamW [51] as the optimizer with a warmup learning rate schedule of 5e-5 peak learning rate, betas of (0.9, 0.98), and weight decay of 1e-3 for all pre-training experiments. We pre-train our model on 32

⁹Following BERT [12], We replace 80% of masked word tokens as the [MASK] token, 10% as a random token, and 10% as its original token.

NVIDIA V100 GPUs with a batch size of 28 per GPU. Pre-training with 10 epochs on WebVid2.5M [3] + CC3M [65] takes about 27 hours to finish. We present the training settings for all finetuning experiments in the next section.

C. Experimental Setup of Downstream Tasks

We test our pre-trained models on 3 popular VidL tasks across 13 downstream datasets, including text-to-video retrieval, video question answering, and video captioning. For text-to-video retrieval, we report downstream performance on MSRVTT [84], DiDeMo [26], and LSMDC [62] and use Recall at K (R@K, K=1,5,10) as the evaluation metric. For video question answering, we consider datasets in both multiple-choice and open-ended settings, including TGIF-Action, TGIF-Transition, TGIF-Frame [27], MSRVTT-MC [87], MSRVTT-QA, MSVD-QA [83], LSMDC-MC and LSMDC-FiB [71]. We evaluate our models using accuracy. For video captioning, we report CIDER scores on MSRVTT and MSVD.

We follow the standard training/validation/testing splits of the original datasets. If not otherwise stated, we sparsely sample $T = 5$ video frames and adopt video frame size 224 with patch size $H = W = 32$. Similar to pre-training, we use AdamW [51] to fine-tune our model for each downstream task with a warmup learning rate schedule of $2e-5$ peak learning rate, betas of (0.9, 0.98), and weight decay of $1e-3$. All finetuning experiments are conducted on Microsoft Azure [1] adopting mixed-precision training with DeepSpeed [61].¹⁰ All video data are pre-processed by evenly extracting 32 frames to avoid expensive decoding on-the-fly. During training, we randomly sample T frames from 32 frames, resize the shorter side of all frames to 224, and random crop (224x224) at the same location for all the frames in a given video. During inference, we evenly sample T frames from 32 frames and center crop (224x224) for all the sampled video frames.

C.1. Text-To-Video Retrieval

For text-to-video retrieval, similar to visual-text matching (VTM) during pre-training, we treat corresponding video-text pairs in the same batch as positives and all other pairwise combinations as negatives. We adopt a fully-connected (FC) layer (FC^{T2V}) over the VidL representation h^c of the [CLS] token to learn through classification:

$$b^{\text{pos}} = FC^{T2V}(h_{\text{pos}}^c), b^{\text{neg}} = FC^{T2V}(h_{\text{neg}}^c),$$

$$\mathcal{L}_{T2V} = -\frac{1}{B} \sum_i \log \frac{b_i^{\text{pos}}}{b_i^{\text{pos}} + \sum b_i^{\text{neg}}}, \quad (4)$$

where h_{pos}^c or h_{neg}^c is h^c of positive or negative pairs. In particular, we use pre-trained FC^{VTM} for zero-shot text-to-

¹⁰We conduct retrieval finetuning on 8 80GB A100 GPUs to enable larger batch size, while all other finetuning experiments are conducted on 8 32GB V100 GPUs.

VideoQA	Task	#Option
Multiple-Choice	TGIF-Action [27]	5
	TGIF-Transition [27]	5
	MSRVTT-MC [87]	5
	LSMDC-MC [71]	5
Open-Ended	TGIF-Frame [27]	-
	MSRVTT-QA [83]	-
	MSVD-QA [8]	-
	LSMDC-FiB [71]	-

Table 20. Summary of **video question answering** tasks. For open-ended Video QA, we do not limit the answer vocabulary to a fixed candidate set.

video retrieval and to initialize FC^{T2V} for further fine-tuning on each downstream text-to-video retrieval task.

MSRVTT [84] contains 10K YouTube videos with 200K human annotations. For fair comparison [3, 36], we train on 9K training+validation splits and evaluate on the 1K-A testing split. We adopt batch size 20 per GPU and train for 10 epochs.

DiDeMo [26] consists of 10K videos annotated with 40K sentences from Flickr. Following [3, 36], we concatenate all sentences from the same video into a paragraph and perform paragraph-to-video retrieval for DiDeMo. We adopt batch size 16 per GPU and train for 10 epochs.

LSMDC [62] contains 118K video clips from 202 movies. Each clip has a caption from movie scripts or descriptive video services. Following [3, 54], we evaluate on 1K testing clips that disjoint from the training+validation splits. We adopt batch size 20 per GPU and train for 5 epochs.

C.2. Video Question Answering

We test our model on video question answering (QA) tasks in both multiple-choice and open-ended settings as Table 20. We follow LAVENDER [42] to formulate Video QA as Masked Language Modeling due to its superior performance. For multiple-choice QA tasks, we concatenate question with all answer options and add a [MASK] to form the input text (Q+A0+A1+A2+A3+A4+[MASK]). We treat the same Masked Language Modeling (MLM) layer as used in pre-training upon h^x to predict the word token corresponding to the answer index (e.g., 0, 1, 2, 3, 4). Similarly, for open-ended QA tasks, we apply MLM over the input (Q+[MASK]). Cross-entropy loss is used to supervise the downstream finetuning over the whole word vocabulary.

TGIF-Action, TGIF-Transition, and TGIF-Frame [27] require spatial-temporal reasoning to answer questions regarding GIF videos in TGIF-QA. Specifically, we aim to test our model along three dimensions: (i) **Action**: to recognize the repeated action; (ii) **Transition**: to identify the transition between the before and after states; (iii) **Frame**: to answer questions about a specific frame from the GIF video. Among them, TGIF-Action and TGIF-Transition are collected under a multiple-choice setting, and TGIF-Frame is

an open-ended video QA task with free-form answers. We adopt batch size 24 and train for 56/20/10 epochs for Action/Transition/Frame, respectively.

MSRVTT-MC [87] and MSRVTT-QA [83] are created based on videos and captions in MSRVTT [84]. MSRVTT-MC is a multiple-choice task with videos as questions, and captions as answers. Each video contains 5 captions, with only one positive match. This setting can be viewed as video-to-text retrieval, hence we simply evaluate the model trained on MSRVTT-Retrieval. MSRVTT-QA contains 243K open-ended questions over 10K videos. We adopt batch size 24 per GPU and training epochs 8.

MSVD-QA [83] consists of 47K open-ended questions over 2K videos, based on video-caption pairs from MSVD [8]. We adopt batch size 24 per GPU and train for 10 epochs.


LSMDC-MC and LSMDC-FiB [71] are built from the LSMDC dataset [62]. Similar to MSRVTT-MC, LSMDC-MC requires the model to select the only positive caption that describes the video from 5 caption candidates, and can be formulated as video-to-text retrieval. LSMDC-FiB replaces a word in the question sentence with the [BLANK] token, and requires the model to recover the missing word. We regard LSMDC-FiB as an open-ended Video QA task. In particular, we replace the [BLANK] token with [MASK] token, and use the MLM prediction head over the representation h_x of the [MASK] token to predict the correct answer. We adopt batch size 24 per GPU and train for 10 epochs.

C.3. Video Captioning

For video captioning, we evaluate on MSRVTT [84] and MSVD [7]. **MSRVTT** consists of 10K videos with 20 captions per video, and **MSVD** contains 2K videos, with 40 captions per video. We follow the standard captioning splits to train/evaluate with VIOLETv2. The captioning finetuning is formulated as masked language modeling (MLM) with a causal attention mask so that the current word token only attends to the tokens before it, following SwinBERT [44]. During training, we set the probability of random masking caption tokens to 0.15, the same as what is used in MLM during pre-training. We adopt batch size 24 per GPU and train for 20 epochs. During inference, we generate the captions auto-regressively. At each generation step, a [MASK] token is appended to the previously generated tokens, and the model will predict the current tokens based on the learned embedding at the [MASK] token position. We perform generation until the model outputs a [SEP], which is defined as the sentence ending token or when it reaches the maximum generation step 50.

"There is someone playing a game in a computer."

"A man is driving a car through the countryside."



"Cartoon show for kids."

"A video of a rock group performing one of their songs."


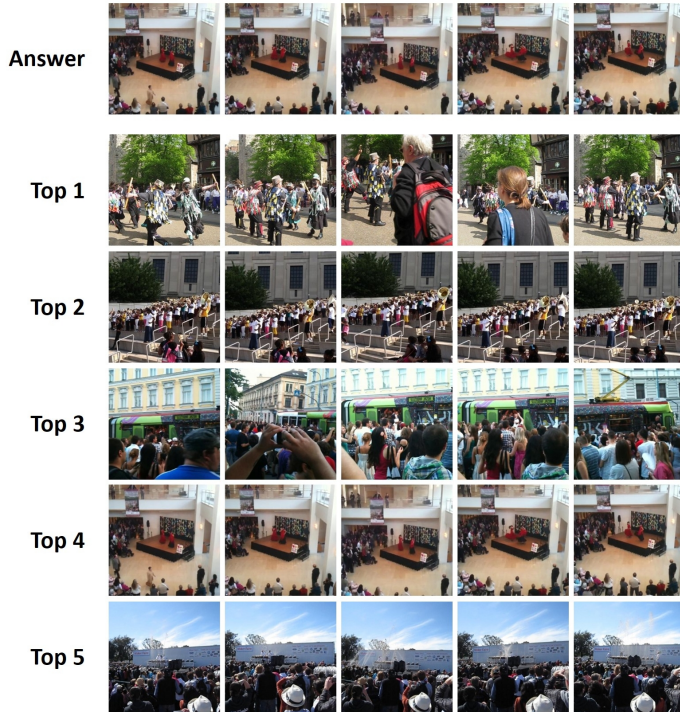


Figure 3. Qualitative examples of zero-shot text-to-video retrieval on MSRVT [84].

"We first see people walking and the crowd. We see two people walk in front of audience. A man ' lady are seen walking through the scene."



"The baby's face is the only face in the frame for a brief time. Then the two girls on either side of him appear. A blond hair girl kneels down beside a baby."



"There is a parade and a man goes by with a tuba. Camera zooms in on a statue the crowd is carrying. A parade of people in purple carrying something passes."



"Camera zooms in to stage. Blue light flashes up from the stage for first time. Lights change from red to green the second time."




Figure 4. Qualitative examples of zero-shot text-to-video retrieval on DiDeMo [26].

"She looks over at SOMEONE."

"SOMEONE turns to SOMEONE with a cold stare."



"SOMEONE gives a perplexed smile."

"Meeting SOMEONE's gaze, she gives a nod."

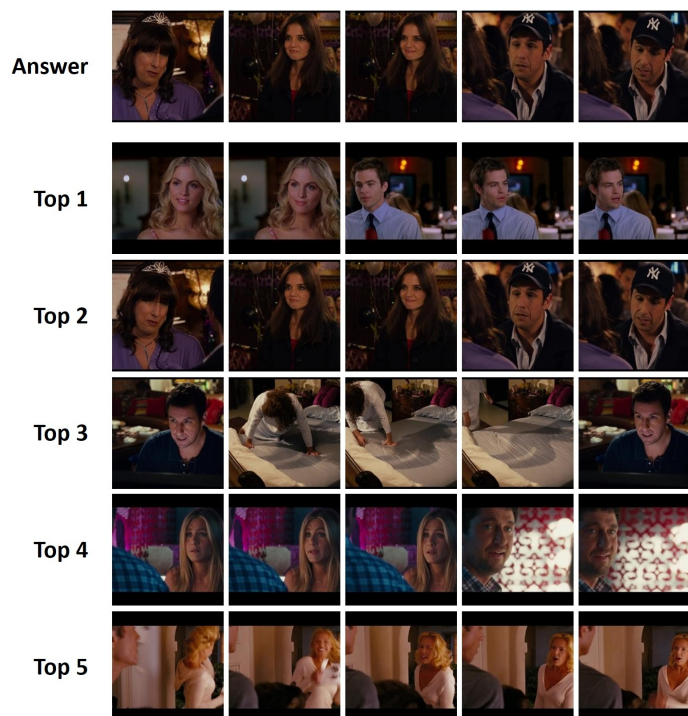


Figure 5. Qualitative examples of zero-shot text-to-video retrieval on LSMDC [62].

References

- [1] Microsoft Azure. <https://azure.microsoft.com/>. 12
- [2] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-Up and Top-Down Attention for Image Captioning and Visual Question Answering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [3] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in Time: A Joint Video and Image Encoder for End-to-End Retrieval. In *International Conference on Computer Vision (ICCV)*, 2021. 2, 4, 5, 6, 7, 8, 9, 10, 11, 12
- [4] Hangbo Bao, Li Dong, and Furu Wei. BEiT: BERT Pre-Training of Image Transformers. In *International Conference for Learning Representations (ICLR)*, 2022. 1, 2, 3, 4, 6
- [5] Shyamal Buch, Cristóbal Eyzaguirre, Adrien Gaidon, Jiajun Wu, Li Fei-Fei, and Juan Carlos Niebles. Revisiting the "Video" in Video-Language Understanding. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 5, 10
- [6] Joao Carreira and Andrew Zisserman. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [7] David L. Chen and William B. Dolan. Collecting Highly Parallel Data for Paraphrase Evaluation. In *ACL*, 2011. 1, 8, 13
- [8] David L. Chen and William B. Dolan. Collecting Highly Parallel Data for Paraphrase Evaluation. In *Annual Meetings of the Association for Computational Linguistics (ACL)*, 2011. 12, 13
- [9] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. UNITER: UNiversal Image-TExt Representation Learning. In *European Conference on Computer Vision (ECCV)*, 2020. 2
- [10] Navneet Dalal and Bill Triggs. Histograms of Oriented Gradients for Human Detection. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005. 3, 4, 7, 11
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: a Large-Scale Hierarchical Image Database. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 2, 6
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2019. 1, 2, 4, 11
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference for Learning Representations (ICLR)*, 2021. 4
- [14] Zi-Yi Dou, Yichong Xu, Zhe Gan, Jianfeng Wang, Shuohang Wang, Lijuan Wang, Chenguang Zhu, Pengchuan Zhang, Lu Yuan, Nanyun Peng, Zicheng Liu, and Michael Zeng. An Empirical Study of Training End-to-End Vision-and-Language Transformers. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [15] Chenyou Fan, Xiaofan Zhang, Shu Zhang, Wensheng Wang, Chi Zhang, and Heng Huang. Heterogeneous Memory Enhanced Multimodal Attention Model for Video Question Answering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [16] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. SlowFast Networks for Video Recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [17] Tsu-Jui Fu, Linjie Li, Zhe Gan, Kevin Lin, William Yang Wang, Lijuan Wang, and Zicheng Liu. VIOLET: End-to-End Video-Language Transformers with Masked Visual-token Modeling. In *arXiv:2111.1268*, 2021. 1, 2, 3, 4, 8, 9
- [18] Valentin Gabeur, Chen Sun, Karteek Alahari, and Cordelia Schmid. Multi-modal Transformer for Video Retrieval. In *European Conference on Computer Vision (ECCV)*, 2020. 2
- [19] Zhe Gan, Linjie Li, Chunyuan Li, Lijuan Wang, Zicheng Liu, and Jianfeng Gao. Vision-Language Pre-training: Basics, Recent Advances, and Future Trends. In *Foundations and Trends in Computer Graphics and Vision*, 2022. 2
- [20] Jiyang Gao, Runzhou Ge, Kan Chen, and Ram Nevatia. Motion-Appearance Co-Memory Networks for Video Question Answering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [21] Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. TALL: Temporal Activity Localization via Language Query. In *International Conference on Computer Vision (ICCV)*, 2017. 2
- [22] Yuying Ge, Yixiao Ge, Xihui Liu, Dian Li, Ying Shan, Xiaohu Qie, and Ping Luo. BridgeFormer: Bridging Video-text Retrieval with Multiple Choice Questions. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 8
- [23] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked Autoencoders Are Scalable Vision Learners. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 3
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 5, 6
- [25] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing Moments in Video with Natural Language. In *International Conference on Computer Vision (ICCV)*, 2017. 2
- [26] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing Moments in Video with Natural Language. In *International Conference on Computer Vision (ICCV)*, 2017. 4, 8, 12, 15
- [27] Yunseok Jang, Yale Song, Youngjae Yu, Youngjin Kim, and Gunhee Kim. TGIF-QA: Toward Spatio-Temporal Reason-

- ing in Visual Question Answering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2, 4, 8, 12
- [28] Jianwen Jiang, Ziqiang Chen, Haojie Lin, Xibin Zhao, and Yue Gao. Divide and Conquer: Question-Guided Spatio-Temporal Contextual Attention for Video Question Answering. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2020. 2
- [29] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The Kinetics Human Action Video Dataset. In *arXiv:1705.06950*, 2017. 2, 4, 5
- [30] Seonhoon Kim, Seohyeong Jeong, Eunbyul Kim, Inho Kang, and Nojun Kwak. Self-supervised Pre-training and Contrastive Representation Learning for Multiple-choice Video QA. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2021. 2
- [31] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-Captioning Events in Videos. In *International Conference on Computer Vision (ICCV)*, 2017. 1, 2
- [32] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Fei-Fei Li. Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. In *International Journal of Computer Vision (IJCV)*, 2017. 2
- [33] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. In *International Conference for Learning Representations (ICLR)*, 2020. 2
- [34] Thao Minh Le, Vuong Le, Svetha Venkatesh, and Truyen Tran. Hierarchical Conditional Relation Networks for Video Question Answering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [35] Jie Lei, Tamara L Berg, and Mohit Bansal. QVHighlights: Detecting Moments and Highlights in Videos via Natural Language Queries. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2021. 2
- [36] 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 *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 7, 8, 12
- [37] Jie Lei, Licheng Yu, Mohit Bansal, and Tamara L. Berg. TVQA: Localized, Compositional Video Question Answering. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018. 2
- [38] Jie Lei, Licheng Yu, Tamara L. Berg, and Mohit Bansal. TVQA+: Spatio-Temporal Grounding for Video Question Answering. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020. 2
- [39] Jie Lei, Licheng Yu, Tamara L. Berg, and Mohit Bansal. TVR: A Large-Scale Dataset for Video-Subtitle Moment Retrieval. In *European Conference on Computer Vision (ECCV)*, 2020. 2
- [40] Dongxu Li, Junnan Li, Hongdong Li, Juan Carlos Niebles, and Steven C.H. Hoi. Align and Prompt: Video-and-Language Pre-training with Entity Prompts. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 7, 8
- [41] 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 *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020. 1, 2, 8, 9, 10
- [42] Linjie Li, Zhe Gan, Kevin Lin, Chung-Ching Lin, Zicheng Liu, Ce Liu, and Lijuan Wang. Lavender: Unifying video-language understanding as masked language modeling. *arXiv preprint arXiv:2206.07160*, 2022. 12
- [43] Linjie Li, Jie Lei, Zhe Gan, Licheng Yu, Yen-Chun Chen, Rohit Pillai, Yu Cheng, Luowei Zhou, Xin Eric Wang, William Yang Wang, Tamara Lee Berg, Mohit Bansal, Jingjing Liu, Lijuan Wang, and Zicheng Liu. VALUE: A Multi-Task Benchmark for Video-and-Language Understanding Evaluation. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2021. 2
- [44] Kevin Lin, Linjie Li, Chung-Ching Lin, Faisal Ahmed, Zhe Gan, Zicheng Liu, Yumao Lu, and Lijuan Wang. SwinBERT: End-to-End Transformers with Sparse Attention for Video Captioning. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 7, 8, 13
- [45] Ziyi Lin, Shijie Geng, Renrui Zhang, Peng Gao, Gerard de Melo, Xiaogang Wang, Jifeng Dai, Yu Qiao, and Hongsheng Li. Frozen Clip Models are Efficient Video Learners. In *European Conference on Computer Vision (ECCV)*, 2022. 10
- [46] Song Liu, Haoqi Fan, Shengsheng Qian, Yiru Chen, Wenkui Ding, and Zhongyuan Wang. HiT: Hierarchical Transformer with Momentum Contrast for Video-Text Retrieval. In *arXiv:2103.15049*, 2021. 2
- [47] Yang Liu, Samuel Albanie, Arsha Nagrani, and Andrew Zisserman. Use What You Have: Video Retrieval Using Representations From Collaborative Experts. In *British Machine Vision Conference (BMVC)*, 2020. 2
- [48] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. RoBERTa: A Robustly Optimized BERT Pretraining Approach. In *arXiv:1907.11692*, 2019. 2
- [49] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. In *International Conference on Computer Vision (ICCV)*, 2021. 4, 5, 6, 7, 9, 11
- [50] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video Swin Transformer. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 4, 5, 10, 11
- [51] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *International Conference for Learning Representations (ICLR)*, 2019. 11, 12
- [52] Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. CLIP4Clip: An Empirical

- Study of CLIP for End to End Video Clip Retrieval. In *arXiv:2104.08860*, 2021. 8
- [53] 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 *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [54] 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 *International Conference on Computer Vision (ICCV)*, 2019. 1, 2, 12
- [55] Bolin Ni, Houwen Peng, Minghao Chen, Songyang Zhang, Gaofeng Meng, Jianlong Fu, Shiming Xiang, and Haibin Ling. Expanding Language-Image Pretrained Models for General Video Recognition. In *European Conference on Computer Vision (ECCV)*, 2022. 10
- [56] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019. 11
- [57] Mandela Patrick, Po-Yao Huang, Yuki Asano, Florian Metze, Alexander Hauptmann, Joao Henriques, and Andrea Vedaldi. Support-set bottlenecks for video-text representation learning. In *International Conference for Learning Representations (ICLR)*, 2021. 2
- [58] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning Transferable Visual Models From Natural Language Supervision. In *International Conference on Machine Learning (ICML)*, 2021. 2, 4, 5, 7, 11
- [59] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-Shot Text-to-Image Generation. In *International Conference on Machine Learning (ICML)*, 2021. 4, 5, 7, 11
- [60] Rene Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision Transformers for Dense Prediction. In *International Conference on Computer Vision (ICCV)*, 2021. 3, 4, 5, 7, 11
- [61] Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. DeepSpeed: System Optimizations Enable Training Deep Learning Models with Over 100 Billion Parameters. In *Knowledge Discovery in Database (KDD)*, 2020. 12
- [62] Anna Rohrbach, Marcus Rohrbach, Niket Tandon, and Bernt Schiele. A Dataset for Movie Description. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 1, 2, 8, 12, 13, 16
- [63] Andrew Rouditchenko, Angie Boggust, David Harwath, Brian Chen, Dhiraj Joshi, Samuel Thomas, Kartik Adhikari, Hilde Kuehne, Rameswar Panda, Rogerio Feris, Brian Kingsbury, Michael Picheny, Antonio Torralba, and James Glass. AVLnet: Learning Audio-Visual Language Representations from Instructional Videos. In *INTER-SPEECH*, 2021. 2
- [64] Paul Hongsuck Seo, Arsha Nagrani, Anurag Arnab, and Cordelia Schmid. End-to-end Generative Pretraining for Multimodal Video Captioning. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 8
- [65] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual Captions: A Cleaned, Hypernymed, Image Alt-text Dataset For Automatic Image Captioning. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2018. 7, 8, 9, 12
- [66] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. VideoBERT: A Joint Model for Video and Language Representation Learning. In *International Conference on Computer Vision (ICCV)*, 2019. 1, 2
- [67] Hao Tan and Mohit Bansal. LXMERT: Learning Cross-Modality Encoder Representations from Transformers. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2019. 2
- [68] Hao Tan, Jie Lei, Thomas Wolf, and Mohit Bansal. VIM-PAC: Video Pre-Training via Masked Token Prediction and Contrastive Learning. In *arXiv:2106.11250*, 2021. 2, 6
- [69] Zachary Teed and Jia Deng. RAFT: Recurrent All-Pairs Field Transforms for Optical Flow. In *European Conference on Computer Vision (ECCV)*, 2020. 3, 4, 5, 10, 11
- [70] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. VideoMAE: Masked Autoencoders are Data-Efficient Learners for Self-Supervised Video Pre-Training. In *arXiv:2203.12602*, 2022. 2
- [71] Atousa Torabi, Niket Tandon, and Leonid Sigal. Learning Language-Visual Embedding for Movie Understanding with Natural-Language. In *arXiv:1609.08124*, 2016. 8, 12, 13
- [72] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training Data-efficient Image Transformers & Distillation through Attention. In *International Conference on Machine Learning (ICML)*, 2021. 10
- [73] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural Discrete Representation Learning. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2017. 4
- [74] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2017. 2
- [75] Alex Jinpeng Wang, Yixiao Ge, Rui Yan, Yuying Ge, Xudong Lin, Guanyu Cai, Jianping Wu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. All in One: Exploring Unified Video-Language Pre-training. In *arXiv:2203.07303*, 2022. 8
- [76] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal Segment Networks: Towards Good Practices for Deep Action Recognition. In *European Conference on Computer Vision (ECCV)*, 2016. 2

- [77] Rui Wang, Dongdong Chen, Zuxuan Wu, Yinpeng Chen, Xiyang Dai, Mengchen Liu, Yu-Gang Jiang, Luwei Zhou, and Lu Yuan. BEVT: BERT Pretraining of Video Transformers. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [78] Xin Wang, Jiawei Wu, Junkun Chen, Lei Li, Yuan-Fang Wang, and William Yang Wang. VATEX: A Large-Scale, High-Quality Multilingual Dataset for Video-and-Language Research. In *International Conference on Computer Vision (ICCV)*, 2019. 2
- [79] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked Feature Prediction for Self-Supervised Visual Pre-Training. In *arXiv:2112.09133*, 2022. 1, 2, 3
- [80] Longhui Wei, Lingxi Xie, Wengang Zhou, Houqiang Li, and Qi Tian. MVP: Multimodality-guided Visual Pre-training. In *arXiv:2203.05175*, 2022. 1, 2
- [81] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking Spatiotemporal Feature Learning: Speed-Accuracy Trade-offs in Video Classification. In *European Conference on Computer Vision (ECCV)*, 2018. 2
- [82] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. SimMIM: A Simple Framework for Masked Image Modeling. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 3
- [83] Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. Video Question Answering via Gradually Refined Attention over Appearance and Motion. In *ACM Multimedia (ACMMM)*, 2017. 1, 2, 8, 12, 13
- [84] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. MSR-VTT: A Large Video Description Dataset for Bridging Video and Language. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1, 2, 8, 12, 13, 14
- [85] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Just Ask: Learning to Answer Questions from Millions of Narrated Videos. In *International Conference on Computer Vision (ICCV)*, 2021. 2, 8
- [86] Zekun Yang, Noa Garcia, Chenhui Chu, Mayu Otani, Yuta Nakashima, and Haruo Takemura. BERT Representations for Video Question Answering. In *Winter Conference on Applications of Computer Vision (WACV)*, 2020. 2
- [87] Youngjae Yu, Jongseok Kim, and Gunhee Kim. A Joint Sequence Fusion Model for Video Question Answering and Retrieval. In *European Conference on Computer Vision (ECCV)*, 2018. 8, 12, 13
- [88] Rowan Zellers, Ximing Lu, Jack Hessel, Youngjae Yu, Jae Sung Park, Jize Cao, Ali Farhadi, and Yejin Choi. MERLOT: Multimodal Neural Script Knowledge Models. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2021. 1, 2, 6, 8, 9
- [89] Bowen Zhang, Hexiang Hu, and Fei Sha. Cross-Modal and Hierarchical Modeling of Video and Text. In *European Conference on Computer Vision (ECCV)*, 2018. 2
- [90] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. iBOT: Image BERT Pre-Training with Online Tokenizer. In *International Conference on Learning Representations (ICLR)*, 2012. 1, 2, 3
- [91] Luwei Zhou, Chenliang Xu, and Jason J. Corso. Towards Automatic Learning of Procedures from Web Instructional Videos. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2018. 2
- [92] Linchao Zhu and Yi Yang. ActBERT: Learning Global-Local Video-Text Representations. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2