

# BEVT: BERT Pretraining of Video Transformers

Rui Wang<sup>1,2</sup> Dongdong Chen<sup>3</sup> Zuxuan Wu<sup>1,2\*</sup> Yinpeng Chen<sup>3</sup> Xiyang Dai<sup>3</sup>  
 Mengchen Liu<sup>3</sup> Yu-Gang Jiang<sup>1,2\*</sup> Luwei Zhou<sup>3</sup> Lu Yuan<sup>3</sup>

<sup>1</sup>Shanghai Key Lab of Intelligent Info. Processing, School of Computer Science, Fudan University

<sup>2</sup>Shanghai Collaborative Innovation Center on Intelligent Visual Computing

<sup>3</sup>Microsoft Cloud + AI

## Abstract

This paper studies the BERT pretraining of video transformers. It is a straightforward but worth-studying extension given the recent success from BERT pretraining of image transformers. We introduce BEVT which decouples video representation learning into spatial representation learning and temporal dynamics learning. In particular, BEVT first performs masked image modeling on image data, and then conducts masked image modeling jointly with masked video modeling on video data. This design is motivated by two observations: 1) transformers learned on image datasets provide decent spatial priors that can ease the learning of video transformers, which are often times computationally-intensive if trained from scratch; 2) discriminative clues, i.e., spatial and temporal information, needed to make correct predictions vary among different videos due to large intra-class and inter-class variations. We conduct extensive experiments on three challenging video benchmarks where BEVT achieves very promising results. On Kinetics 400, for which recognition mostly relies on discriminative spatial representations, BEVT achieves comparable results to strong supervised baselines. On Something-Something-V2 and Diving 48, which contain videos relying on temporal dynamics, BEVT outperforms by clear margins all alternative baselines and achieves state-of-the-art performance with a 71.4% and 87.2% Top-1 accuracy respectively. Code is available at <https://github.com/xyzforever/BEVT>.

## 1. Introduction

Transformers [54, 59] have become the dominant network structures in the natural language processing (NLP) field and have demonstrated tremendous success in different NLP tasks. Recently, the pioneering work ViT [19] tokenizes one image into a series of patch-based tokens

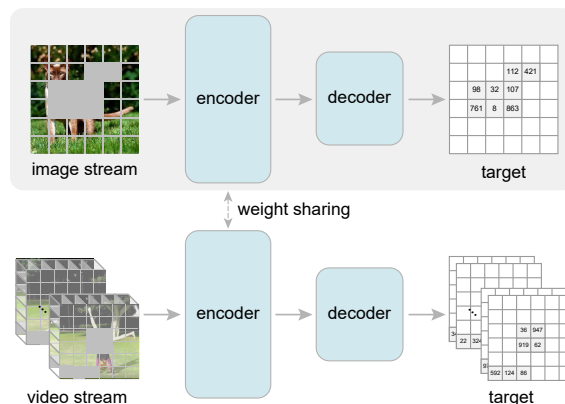


Figure 1. A conceptual overview of BEVT. With a decoupled design, BEVT **first** conducts masked image modeling on image data, and **then** conducts jointly masked image modeling and masked video modeling on image&video data by weight sharing.

and applies the transformer architecture for image recognition. Many approaches [12, 17, 36, 60] further demonstrate the power of transformers as generic vision backbones and achieve impressive performance on various vision tasks. Beyond image tasks, there are also a few studies showing the promise of transformers for video understanding [2, 37].

The key to the success of Transformers in NLP is BERT pretraining [4, 14, 35], one of the most successful pretraining tasks, which predicts masked tokens in corrupted texts. This motivates a few recent studies to explore the BERT-style pretraining for image representation learning by recovering raw pixels [29] or latent codes [3, 18] of masked image patches. However, how to leverage such a strategy for video understanding has never been explored.

In this paper, we study BERT pretraining of video transformers. Unlike static images, videos depict how objects move and interact over time. Such dynamic nature brings additional difficulty for representation learning. It is often found that learning representations from scratch on videos is computationally expensive and requires extremely large-scale datasets with millions of samples [24], if not hun-

\*Corresponding authors.

dreds of millions of samples [1]. Instead of training from scratch, a few methods demonstrate that self-supervised models pretrained on image datasets benefit video recognition under both supervised [2, 37] and unsupervised settings [5]. These approaches simply leverage pretrained models as better initializations to learn spatial-temporal features in videos. While widely used and sometimes effective, the spatial context relationships learned from the image pre-training phase are likely to be drastically modified during video feature learning.

We argue that spatial priors encoded in pretrained self-supervised models should be explicitly preserved when performing video representation learning. The intuition behind is that there are large inter-class variations among different videos and their dependencies on what discriminative information to use (*i.e.*, spatial and temporal clues) to make correct predictions differ. For instance, for actions like “applying lipstick”, spatial knowledge is generally sufficient, as evidenced by the fact simply using 2D features offers decent results on datasets like Kinetics [9]. On the other hand, temporal dynamics are crucial for differentiating actions between two fine-grained diving sequences [33]. This highlights the importance of considering the differences among video samples during feature learning.

In light of this, we introduce BEVT, which decouples video representation learning into spatial representation learning and temporal dynamics learning. More specifically, BEVT builds upon the Video Swin Transformer [37] due to their computationally efficient architectures<sup>\*</sup>, and is trained with a BERT-style objective to fully unleash the power of transformers for representation learning. BEVT contains an image stream for spatial modeling and a video stream for temporal modeling, interacting with each other for video modeling. In particular, the image stream, operating on RGB images, learns spatial priors first in an unsupervised fashion by predicting masked image patches in the form of latent codes derived from a pretrained VQ-VAE [3]. It is then used to initialize the attention weight matrices of the video stream, whose inputs are sampled video clips, so as to save computation for video transformers. The video stream, on the other hand, learns temporal dynamics through predicting masked 3D tubes represented by latent codes. The two streams, taking image and video pairs as inputs, are then jointly trained on video data through a weight sharing strategy. Such a design not only maintains spatial knowledge learned from image datasets to ensure decent results for static video samples but also learns temporal information to guarantee correct predictions for samples that contain dynamic movements. Finally, BEVT is finetuned on targeted datasets for downstream evaluation.

We conduct extensive experiments on three challenging

video datasets, *i.e.*, Kinetics-400 (K400) [9], Something-Something-v2 (SSv2) [26], and Diving-48 (DIVING-48) [33]. On K400, BEVT offers 81.1% Top-1 accuracy, which is better than the strong supervised baseline 80.6% [37]. On SSv2 and DIVING48, BEVT achieves 71.4% and 87.2% Top-1 accuracy outperforming state-of-the-art methods [2, 5, 23, 37] by clear margins. To further analyze the performance difference among these three datasets, we further provide the temporal dependency analysis and demonstrate that videos in K400 mainly rely on spatial clues for correct predictions while videos from SSv2 and DIVING48 require more temporal information.

Our main contributions are summarized as follows: (1) We explore the BERT-style training objective to fully unleash the power of transformers to learn discriminative video representations; (2) We introduce a novel two-stream network that decouples spatial representation learning and temporal dynamics learning; (3) We demonstrate different video samples have different preferences towards spatial and temporal clues; (4) We conduct extensive experiments on three challenging video benchmarks and achieve comparable or better results with state-of-the-art methods.

## 2. Related Work

**Video understanding with CNNs.** There is a plethora of work on video understanding with CNNs, most of which focus on learning spatial-temporal features [9, 22, 23, 34, 50, 53, 57, 61, 62]. These approaches can be divided into two categories: (1) temporal aggregation and (2) 3D CNNs. In particular, temporal aggregation methods typically extract image features/scores frame-by-frame and then combine frame-level information to achieve video-level predictions through recurrent networks [15, 66] or average pooling [47, 57]. On the other hand, 3D CNNs extend 2D convolutions into the time domain by using 3D convolutions on stacked RGB frames for the joint-modeling of spatial-temporal relationships [9, 22, 23, 50, 53]. 3D CNNs are generally computationally expensive, and this motivates a line of research on efficient video recognition [13, 22, 34, 51, 53, 72, 74]. Instead of using CNNs, we explore transformers for video understanding due to their strong results on image recognition tasks.

**Vision transformers.** Motivated by the impressive performance of transformers in a wide range of NLP tasks, there is a growing interest in using transformers for computer vision tasks [19, 20, 28, 36, 38, 43, 49, 56, 73]. More specifically, ViT [19] generalizes transformers to the image domain by splitting images into patches which are further embedded with a linear layer as inputs to a Transformer. While demonstrating great potential in image recognition tasks, ViT relies on pretraining on substantially large-scale datasets and the training process is computationally expensive. To mit-

<sup>\*</sup>Note that we only use the architecture and do not load the pretrained weights.

igate these issues, extensive studies have been introduced. For example, DeiT [49] uses a distillation loss to speed up the training process. CSwin Transformer [17] proposes multi-head grouping and performs attention within cross-shaped windows. MobileFormer [12] further introduces lightweight transformers. There are also a few very recent studies that extend image transformers for video understanding [2, 20, 37, 68]. Fan *et al.* use a multi-scale design to generate spatial-temporal tokens in different sizes for action recognition [20]. Liu *et al.* extend Swin Transformers into the video domain [37]. In contrast, we study BERT pretraining of video transformer in a self-supervised way, which is orthogonal to transformer network design.

**Self-supervised representation learning.** At the core of many computer vision tasks is how to learn discriminative features for targeted datasets. Since collecting labeled datasets is labor-intensive and costly, there is an ever-increasing trend in learning representations in a self-supervised manner [6, 7, 16, 27, 32, 67]. The main idea is to design surrogate tasks including inpainting [44], colorization [70, 71], jigsaw predictions [42], rotation predictions [25], *etc.*, as a form of supervisory signals in lieu of manual labels. More recently, contrastive learning has been a popular paradigm for feature learning by forcing images to be closer to their augmented copies than other samples [10, 30, 40, 63, 65]. In contrast to these approaches using CNNs as backbones, there are a few very recent studies leveraging contrastive learning [8, 11] for transformers.

**BERT pretraining.** In contrast to contrastive learning widely used in vision, BERT pretraining [14] is extremely popular and extensively studied [4, 35] in NLP. As an effort that unifies vision and NLP under the same BERT pretraining framework, the recent work BEiT [3] and ICT [55] utilizes the masked image modeling task to do BERT pretraining of image transformers and achieves great success for different tasks. And one concurrent work PeCo [18] further proposes a perceptual codebook to improve the performance. Another concurrent work [29] extends it from recovering patch tokens to raw pixels. In this paper, we study BERT pretraining for video transformers as an orthogonal unifying effort. Different from BERT pretraining of image transformers and the concurrent effort [48], we decouple video pretraining into spatial representation learning and temporal dynamics learning so as to accommodate the varying need of distinct salient clues for different videos.

### 3. Method

The goal of BEVT is to learn video representations effectively for both relatively static videos and dynamic videos in a self-supervised manner. Here, “relatively static videos” mean the videos only requiring discriminative spatial representation for recognition, while “dynamic videos” mean

that videos that also require temporal dynamics for recognition. Besides the effectiveness, another key problem to consider in video pretraining is efficiency. Compared to image pretraining, video pretraining is more computationally expensive, thus making pretraining on large-scale video data from scratch inefficient or even inapplicable without massive computational resources.

To this end, BEVT decouples the video pretraining into spatial representation learning and temporal dynamics learning. And the spatial representation learning is only conducted on image data, while the temporal dynamics learning is conducted on video data. To implement this idea, our BEVT contains two streams, operating on images and videos, respectively. In the following, we introduce different components of our framework. Figure 2 gives an overview of our framework.

**Image and video patches.** For the video stream, given a video clip  $X_{vid} \in \mathbb{R}^{T \times H \times W \times 3}$  with  $T$  frames, we follow VideoSwin [37] and convert it into  $\frac{T}{2} \times \frac{H}{4} \times \frac{W}{4}$  3D patches, each with a size of  $2 \times 4 \times 4 \times 3$ ; each 3D patch contains a 96-dimensional features. For the image stream, given an input image  $X_{img} \in \mathbb{R}^{H \times W \times 3}$ , we consider each patch with a size of  $4 \times 4 \times 3$  as a token and set the feature dimension of each token as 48. We then project each token to a token embedding vector of dimension  $C$  by a linear embedding layer. Then the sequence of token embedding is input into the following transformer architectures.

**Masked image and video tokens.** Motivated by the great success of BERT in NLP tasks, BEVT is optimized to simultaneously perform masked image modeling (MIM) and masked video modeling (MVM) by predicting “corrupted” image and video tokens, respectively. The MIM is designed to capture spatial priors while the MVM is used to capture temporal dynamics in videos. In particular, for the image stream, since input images are divided into non-overlapping patches, we randomly mask several patches and the image stream is trained to recover them as in [3]. More specifically, the embedded feature of each masked patch is replaced by a learnable mask token embedding. For the video stream, we randomly mask 3D tokens and train the video stream to predict those masked tokens. The set of masked image and video tokens and the remaining patch features are sent to encoders, as will be introduced below.

**Mask strategy.** For masked image modeling, following [3], we use blockwise masking instead of randomly selecting each masked patch. When generating masked positions for an image, we mask a block of patches each time and set the minimum number of patches for each block. The position, the aspect ratio and the size of each block are randomly selected under a preset range. We repeat masking blocks until the ratio of masked patches exceeds the preset lower bound. For masked video modeling, we employ

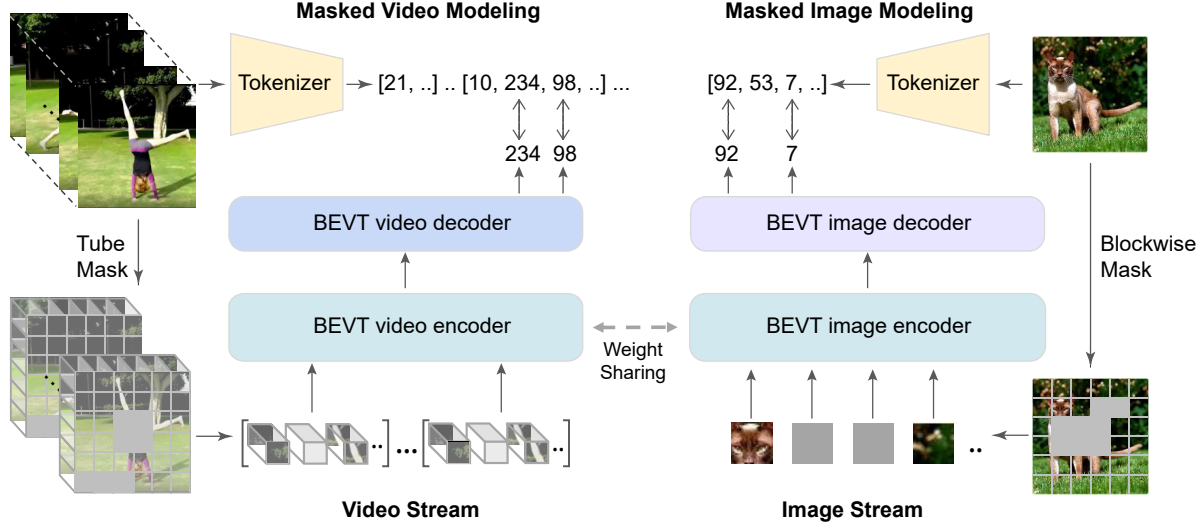


Figure 2. An overview of our framework. BEVT contains an image stream and a video stream that learns video representations jointly using a BERT-style objective. In particular, the image and video stream, operating on single images and video cubes, respectively, predict masked image patches and 3D cubes derived from a tokenizer.

a tube masking strategy that is a straightforward extension of blockwise masking. Given an input video clip of length  $T$ , we first randomly choose the number of masked frames (tube length)  $l$  and the start frame  $t$ . Then we employ blockwise masking to generate a 2D mask, and apply this 2D mask to each frame from  $t$  to  $t + l$ . In other words, for each masked frame, the set of masked positions is the same and the shape of the whole 3D mask is a tube. The range of the masked tube length is  $[0.5T, T]$  and the masking ratio of each masked frame is 0.5.

**BEVT encoders.** BEVT contains two encoders, one for the image stream and one for the video stream. Both encoders are instantiated with the Video Swin Transformer [37] due to its strong performance with a moderate computational cost. Note that in contrast to [37] that performs fully-supervised training, we use the Video Swin Transformer as our backbone for self-supervised learning. In particular, Video Swin Transformer [37] follows the design of Swin Transformer [36] and is a hierarchical architecture consisting of four stages. Between every two stages, spatial downsampling is performed by patch merging layers, which concatenates the features of each group of  $2 \times 2$  spatially neighboring patches. After downsampling, a linear layer maps the features of each concatenated token to half of their dimension. A series of Swin attention blocks comes after to apply feature transformation.

Given a sequence of tokens as inputs, the video encoder outputs a feature map with the size of  $\frac{T}{2} \times \frac{H}{32} \times \frac{W}{32} \times 8C$ . Since Video Swin Transformer only performs temporal downsampling in the beginning linear embedding layer, it degrades to a 2D architecture when the temporal dimension of the input is 1. As a result, for the image encoder, the

output feature map has a size of  $\frac{H}{32} \times \frac{W}{32} \times 8C$ .

**Tokenizer.** Following [3], we use the visual tokens generated by a pretrained image VQ-VAE [46] as the groundtruth tokens and our pretraining task is to predict the tokens for masked patches. The pretrained VQ-VAE tokenizer maps image patch into discrete tokens  $z$  by searching the closest latent codes in its pre-learned visual codebook. Given an input image  $X_{img} \in \mathbb{R}^{H \times W \times 3}$ , it will be tokenized into the visual token map  $Z_{img} \in V^{\frac{H}{16} \times \frac{W}{16}}$ . Similarly for an input video  $X_{vid} \in \mathbb{R}^{T \times H \times W \times 3}$ , it is tokenized into visual token map  $Z_{vid} \in V^{T \times \frac{H}{16} \times \frac{W}{16}}$ . Note that, considering the pretrained VQ-VAE only downsamples  $8 \times 8$  patch into one token, we downsample the input images/frames by 1/2 before feeding into the tokenizer so that the output token map has the spatial resolution of  $\frac{H}{16} \times \frac{W}{16}$ .

**BEVT decoders.** To learn meaningful representations by predicting the tokens for the masked image and video patches, BEVT has an image decoder and a video decoder as the auxiliary prediction heads, which will be discarded in finetuning stage. Existing modern vision transformers including Swin Transformer follow the hierarchical design and downsample the input into decreased spatial/temporal resolutions. Taking the VideoSwin of video stream shown in Figure 3 as an example, it consists of four stages, and the feature maps  $F_4$  in the last stage have the dimension of  $\frac{T}{2} \times \frac{H}{32} \times \frac{W}{32}$ . In order to match the dimension of feature maps to the number of groundtruth visual tokens, we design a lightweight decoder for the video stream in BEVT. As shown in Figure 3, it first spatially upsamples the stage-4 feature  $F_4$  by using a transposed convolutional layer, and then concatenate the upsampled stage-4 feature  $\hat{F}_4$  with



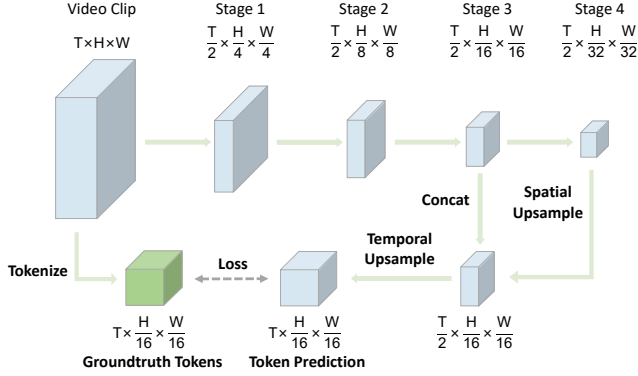


Figure 3. BEVT encoder and decoder for masked video modeling.

stage-3 feature  $F_3$  together and fuse them with a simple linear layer. Finally, the fused feature  $F$  will be temporally upsampled with another transposed convolution layer to obtain  $\hat{F}$ . To predict the token for each position  $(t, i, j)$ , a simple softmax based classifier is applied upon  $\hat{F}$ .

$$\mathbf{p}_{t,i,j} = \text{softmax}(W f_{t,i,j} + b) \quad (1)$$

where  $f_{t,i,j}$  is the feature vector of the output feature map  $\hat{F}$  at position  $(t, i, j)$ ,  $\mathbf{p}_{t,i,j}$  denotes the corresponding probability vector.  $W$  and  $b$  are the weight and the bias of a linear layer. For the decoder in the image stream, it follows a similar design, and the only difference is without the temporally upsampling part.

**Training objectives.** Denote the positions of masked patches in the input images and videos as  $M_I$  and  $M_V$ , the objective of masked image modeling is to maximize the log-likelihood of the groundtruth token  $z_{i,j}$  for each mask position  $(i, j)$ .

$$L_{MIM} = -\frac{1}{|M_I|} \sum_{(i,j) \in M_I} \log \mathbf{p}_{i,j}^{z_{i,j}} \quad (2)$$

where the superscript of  $\mathbf{p}$  denotes indexing the probability value of one specific position. Similarly, the target of masked video modeling can be denoted as:

$$L_{MVM} = -\frac{1}{|M_V|} \sum_{(t,i,j) \in M_V} \log \mathbf{p}_{t,i,j}^{z_{t,i,j}} \quad (3)$$

The objective of two-stream joint training is a simple combination of two objectives:

$$L = L_{MIM} + \lambda L_{MVM} \quad (4)$$

where  $\lambda$  is the hyper-parameter that balances the weights of the image stream and the video stream.

**Training strategies.** Following our decoupled design, we first train the image stream on ImageNet with the masked

image modeling task to learn discriminative spatial representation. The resulting model is then used to initialize the video stream, and *both streams are jointly trained* by optimizing Equation 4 such that the objective  $L_{MIM}$  preserves spatial information while  $L_{MVM}$  learns to capture temporal dynamics in videos. Such a strategy not only makes BEVT much more efficient than pretraining video transformers on large-scale video data from scratch, but also satisfies the need of learning different discriminative clues for different types of video samples.

**Weight sharing between streams.** When jointly training the image and video stream, instead of learning two sets of model weights for the two streams independently, we design a weight sharing strategy so that they can share model weights for the encoder except some image/video specific parts. This is motivated by the good property of transformer networks, *i.e.*, most operators (including multi-head attention and FFN) are oriented to tokens but not specific input types. Taking the Video Swin transformer as an example, we use the following strategies for weight sharing: (1) We use independent 2D patch partitioning layers instead of 3D patch partitioning, and add a linear embedding layer in the first stage for projecting image tokens to the same dimension as the original 3D video patch embedding; (2) We adapt 3D shifted local window to the 2D scenario. This is fulfilled by reusing the submatrix of the original 3D relative positional embedding where the relative temporal distance is 0 as the 2D relative positional embedding. With such a design, the image stream and the video stream can help each other by optimizing one “mostly-unified” encoder.

**Finetuning and inference.** Once pretrained, BEVT provides decent video representations that can be transferred for downstream tasks. On targeted datasets, we simply use the 3D patch embedding layers and the video encoder, to which a few task-specific layers (e.g., classification head for video recognition) are appended, for finetuning. The resulting model can then be readily used for inference.

## 4. Experiments

### 4.1. Experimental Setup

**Datasets and evaluation metrics.** We evaluate our method on three representative video recognition datasets: Kinetics-400 (K400) [9], Something-Something-v2 (SSv2) [26], and Diving-48 (DIVING48) [33]. K400 contains video clips from YouTube with an average duration of 10 seconds and the videos are manually labeled into 400 categories. Following [23], we use  $\sim 240K$  videos for training and  $\sim 20K$  videos for testing. SSv2 is also a large-scale video dataset that contains  $\sim 160K$  videos for training and  $\sim 20K$  videos for testing. The videos in SSv2 are labeled into 174 classes the average duration is 4 seconds. DIVING48 contains  $\sim 17K$  fine-grained diving sequences, which are fur-

ther split into a training set with around  $\sim 15K$  clips and a testing set with  $\sim 2K$  clips. Compared to K400, recognizing videos in SSV2 and DIVING48 requires more temporal information, as will be introduced below. Following official instructions, we report Top-1 accuracy on all three datasets. And the default resolution  $224 \times 224$  is used.

**Implementation Details.** We use Video Swin-Base for experiments throughout the paper unless mentioned otherwise. For pretraining the image stream BEVT-I alone, we train for 800 epochs on ImageNet-1K with a batch size of 2048. For pretraining the video stream BEVT-V alone or the two-stream BEVT, we train for 150 epochs on K400 with a batch size of 256 and the clip length  $T$  is 16. For two stream pretraining, IMAGENET images are used to train the image stream with a batch size of 2048, and the loss weight  $\lambda$  is simply set to 1. We use the DALL-E tokenizer [46] unless explicitly stated. The image stream pretraining takes about 4 days on 32 V100 GPUs. The two-stream pretraining of 150 epochs takes about 4 days on 32 V100 GPUs.

## 4.2. Main Results

**Effectiveness of BEVT for video transformer pretraining.** To demonstrate the effectiveness of BEVT, we compare it with the four image transformer pretraining baselines: (1) *Image Sup*: pretraining the image Swin Transformer on Imagenet-1K in a supervised way. Similar strategies are commonly used in existing video transformer papers [2, 5, 37]. (2) *Image CL*: pretraining the image Swin Transformer on Imagenet-1K with a self-supervised contrastive learning method [64]. (3) *BEVT-I*: pretraining the image Swin Transformer only with the image stream, which is similar to BEiT [3]. (4) *BEVT-V*: pretraining the video Swin Transformer only with the video stream. The pretrained weights from *Image Sup*, *Image CL* and *BEVT-I* used as the initialization of the Video Swin Transformer for finetuning. For the video stream, we design two baselines by conducting BERT pretraining on K400 and HowTo100M [39] from scratch, *i.e.*, *BEVT-V*, which is our framework without the decoupled design. As emphasized before, because video pretraining is more computationally expensive than image pretraining, pretraining on HowTo100M dataset with many epochs is not applicable. For fair comparisons, we also use 32 V100 GPUs to pre-train HowTo100M about 8 days (about 2 epochs).

The comparison results are summarized in Table 1. We observe that: (1) BEVT outperforms the *Image Sup* baseline by clear margins (4.3% and 2.7%) on SSV2 and DIVING48, respectively. This not only suggests that learning representations with BEVT using a BERT-style training objective is promising without the need for manual labels, but also shows that only image-based pretraining is not enough for these two datasets. On K400, the performance of BEVT is on par with the *Image Sup* baseline. (2) We also see that

Method	Pretrain	SSv2	DIVING48	K400
Image Sup	IN-1K	66.3	84.0	80.6
Image CL	IN-1K	67.1	85.5	<b>80.9</b>
BEVT-I	IN-1K	69.2	81.2	80.5
BEVT-V	K400	67.1	83.7	76.2
BEVT-V	HowTo100M	64.2	82.3	75.1
BEVT	IN-1K+K400	<b>70.6</b>	<b>86.7</b>	80.6

Table 1. Comparison of different pretraining methods. Video Swin-Base is used here.

Dataset	Normal	Single-frame	Random-Shuffling
K400	80.6	65.3 $\downarrow$ 15.3	77.8 $\downarrow$ 2.8
SSv2	66.3	6.3 $\downarrow$ 60.0	19.0 $\downarrow$ 47.3
DIVING48	84.0	13.8 $\downarrow$ 70.2	50.4 $\downarrow$ 33.6

Table 2. Effects of removing temporal information for different video datasets. Video Swin-Base models trained with labeled video data are used. We show top-1 accuracy for evaluation.

BEVT offers comparable or better results compared to *Image CL* on these three datasets. (3) Compared to BEVT-I, BEVT is better on SSV2 and DIVING48 by 1.4% and 5.5% respectively, highlighting the gains brought by the video stream. Similarly, BEVT obtains similar results as BEVT-I on K400. (4) Compared to BEVT, BEVT-V pretraining on K400 or HowTo100M from scratch under the similar computation budget achieves much worse results. We hypothesize it may be because the data diversity of K400 is not as good as ImageNet. And for HowTo100M, pretraining much more epochs may help learn better video representation, but it is too costly. This also further justifies the decoupled design in our BEVT.

**Deeper dataset analysis.** To further understand the performance variations of BEVT among three datasets, we perform a temporal dependency study to investigate the amount of temporal information required for correct predictions. Specifically, we use the following two testing strategies: (1) Single-frame, where we randomly sample a frame and replace all other frames with this one, leading to a static video; (2) Random-Shuffling, where a random shuffling is performed along the temporal axis. The results are summarized in 2. We observe that both strategies have a relatively small impact on K400 compared to SSV2 and DIVING48 where there is a 60% and 70% performance drop when using the Single-frame strategy. This suggests most videos in K400 can be recognized by discriminative spatial clues whereas temporal dynamics is particularly important for SSV2 and DIVING48. Comparing across Table 1 and Table 2, we make the following conclusions: (1) On datasets like K400 where spatial clues are dominant, finetuning a model with spatial priors, *e.g.* pretrained on IMAGENET, can achieve

Method	Pretrain	Top-1	GFLOPs × crops
TimeSformer-HR [5]	IN-21K	62.5	$1703 \times 3$
SlowFast R101 [23]	K400	63.1	$106 \times 3$
TSM-RGB [34]	K400	63.3	$62 \times 6$
MSNet [31]	IN-21K	64.7	$67 \times 1$
blVNet [21]	SSv2	65.2	$129 \times 1$
ViViT-L [2]	-	65.4	$903 \times \text{N/A}$
MViT-B [20]	K400	67.7	$455 \times 3$
Mformer-L [45]	IN-21K+K400	68.1	$1185 \times 3$
Swin-B [37]	IN-1K	66.3	$321 \times 3$
Swin-B [37]	IN-21K+K400	69.6	$321 \times 3$
BEVT	IN-1K+K400	<b>70.6</b>	$321 \times 3$
BEVT †	IN-1K+K400	<b>71.4</b>	$321 \times 3$

Table 3. Comparison to state-of-the-art on SSv2. † denotes that we use the IN-1K pretrained PeCo tokenizer [18] instead of the DALL-E tokenizer [46] during pretraining.

decent performance. Additional video modeling brings little effect to the overall performance; (2) The use of the video stream in BEVT is crucial to learn necessary temporal information for datasets like SSv2 and DIVING48. This confirms our hypothesis that different videos rely on different discriminative clues for accurate predictions due to the large intra-class and inter-class variations among videos.

**Comparisons with State-of-the-art methods.** We compare BEVT with alternative methods on SSv2, DIVING48 and K400. On SSv2 and DIVING48, we see from Table 3 and Table 4 that our approach achieves the best performance by clear margins when compared to existing SOTA methods, including supervised models. It is worth mentioning that, on SSv2, a common practice to achieve better results is to perform two rounds of pretraining—a model is pre-trained on both IMAGENET and K400 in a fully supervised fashion before finetuning the model on SSv2. Instead, we pretrain on IMAGENET and K400 without using any manual labels, yet our performance is still better. On K400, we see from Table 5 that BEVT achieves competitive results with SOTA methods using similar or less computation measured by GFLOPs.

### 4.3. Ablation Study

**Importance of image stream pretraining.** In our BEVT, we first conduct the image stream pretraining only on the large-scale image data to efficiently learn spatial representations, then use it as the initialization for joint pretraining. To show its importance, we provide some ablation results in Table 6, where the column “Init” means whether to use the image stream pretrained weights as initialization or not. We have some interesting findings: (1) Using the image stream pretrained weights as initialization

Method	Pretrain	Top-1	Params
SlowFast R101 [23]	K400	77.6	53.3M
TimeSformer-L [5]	IN-21K	81.0	121.4M
TQN [69]	K400	81.8	N/A
Swin-B [37]	IN-1K	84.0	88.1M
BEVT	IN-1K+K400	<b>86.7</b>	88.1M
BEVT †	IN-1K+k400	<b>87.2</b>	88.1M

Table 4. Comparison to state-of-the-art on DIVING48. † denotes that we use the IN-1K pretrained PeCo tokenizer [18] instead of the DALL-E tokenizer [46] during pretraining.

Method	Pretrain	Top-1	GFLOPs × crops
R(2+1)D [53]	-	72.0	$75 \times 10$
I3D [9]	IN-1K	72.1	$108 \times \text{N/A}$
NL I3D-101 [58]	IN-1K	77.7	$359 \times 30$
ip-CSN-152 [52]	-	77.8	$109 \times 30$
SlowFast R101 [23]	-	79.8	$234 \times 30$
X3D-XXL [22]	-	80.4	$144 \times 30$
MViT-B, $32 \times 3$ [20]	-	80.2	$170 \times 5$
MViT-B, $64 \times 3$ [20]	-	81.2	$455 \times 9$
Mformer [45]	IN-21K	79.7	$369.5 \times 30$
ViT-B-VTN [41]	IN-21K	78.6	$4218 \times 1$
TimeSformer-L [5]	IN-21K	80.7	$2380 \times 3$
ViViT-L/ $16 \times 2$ [2]	IN-21K	80.6	$1446 \times 12$
Swin-B [37]	IN-1K	80.6	$282 \times 12$
BEVT	IN-1K+K400	80.6	$282 \times 12$
BEVT †	IN-1K+K400	81.1	$282 \times 12$

Table 5. Comparison to state-of-the-art on K400. † denotes that we use the IN-1K pretrained PeCo tokenizer [18] instead of the DALL-E tokenizer [46] during pretraining.

can benefit both pure video stream pretraining (*i.e.*, “BEVT-V”) and the following joint pretraining of image stream and video stream (*i.e.*, “BEVT”). (2) Even with the initializations, jointly training the image stream with the video stream is still necessary and can bring desirable performance gain.

**Image data in joint training.** By default, when jointly learning spatial and temporal representations, the image stream in BEVT continues to use the IMAGENET images as training images. In this ablation, we also experiment with a variant that uses K400 frames for the image stream. The results are shown in Table 7. We see that images from IMAGENET are slightly better than those from K400, *i.e.* less than 0.3% on all three datasets. This suggests that the image stream, designed to preserve spatial knowledge, is not very sensitive to data domains.

**Different Pretrained Tokenizer.** We also experiment with the PeCo tokenizer [18] instead of the DALL-E tok-

Method	Init	SSv2	DIVING48	K400
BEVT-I	-	69.2	81.2	80.5
BEVT-V	×	67.1	83.7	76.2
BEVT-V	✓	70.0	85.2	79.6
BEVT	×	67.9	85.1	78.5
BEVT	✓	<b>70.6</b>	<b>86.7</b>	<b>80.6</b>

Table 6. Ablation study to show the importance of image stream pretraining. Init means models are initialized from image transformers pretrained with the image stream on ImageNet-1K.

Method	Video	Image	SSv2	Diving48	K400
BEVT-V	K400	-	70.0	85.2	79.6
BEVT	K400	K400	70.3	86.6	80.5
BEVT	K400	IN-1K	<b>70.6</b>	<b>86.7</b>	<b>80.6</b>

Table 7. Ablation study on the image data of joint pretraining. Models are initialized from image transformers pretrained with the image stream on ImageNet-1K.

Method	Pretrain	SSv2	DIVING48	K400
Image Sup	IN-1K	57.6	77.4	75.4
BEVT-I	IN-1K	61.8	77.2	76.5
BEVT	IN-1K+K400	<b>65.7</b>	<b>82.3</b>	<b>77.1</b>

Table 8. Comparison of different pretraining methods with TimeSformer [5]. PeCo tokenizer is used here.

enizer [46] in BEVT. PeCo is only pretrained on ImageNet-1K and uses the same codebook size as in DALL-E. As shown in Table 3-5, PeCo outperforms the DALL-E on all three datasets and pushes BEVT to better results with 71.4% and 87.2% Top-1 accuracy on SSv2 and DIVING48 respectively. This demonstrates that better performance can be achieved with a better visual tokenizer.

**Extension to different model architectures.** We also instantiate the BEVT framework with TimeSformer [5]. The results in Table 8 demonstrate that BEVT offers better results compared to IN-1K supervised pretraining and image stream pretraining on three video downstream tasks.

**Effect of masking strategies.** We evaluate the BEVT-V with different masking strategies for the video stream, *i.e.* the temporal length to be masked and the ratio of masking. The experiments are conducted with Video Swin Tiny for time consideration. In addition to tube masking strategies, we compare with: (1) *Random-3D*: which samples random patches and masks them following a uniform distribution. (2) *Frame-Diff*: which uses the same strategy to choose masked frames as tube masking, but applies blockwise masking independently for each frame. 2D masks may be different for different masked frames. (3) *Ran-*

Strategy	Length	Ratio	SSv2	K400
Tube	0.5T-T	40%	61.5	70.9
Tube	0.5T-T	50%	63.3	71.6
Tube	0.5T-T	60%	63.6	71.4
Tube	0.5T-T	70%	63.5	71.4
Tube	0.25T-0.75T	50%	62.8	69.5
Tube	0.75T-1.0T	50%	63.1	71.6
Tube	T	50%	62.6	71.2
Random-3D	-	50%	59.1	67.4
Frame-Diff	0.5T-T	50%	62.9	70.6
Random-Frame	0.5T-T	50%	62.6	70.9

Table 9. Ablation study on the mask strategy. Video Swin-Tiny is used for this study.

*dom-Frame*: which samples random frames and masks them with the same 2D mask generated by blockwise masking. The temporal positions of masked frames may not be consecutive. The results are summarized in Table 9. We have several observations: (1) Masking tubes offers the better results compared to other masking methods like Random-3D and Random-Frame. (2) Setting too small tube temporal length (e.g., [0.25T, 0.75T]) or too large temporal length (e.g., T) will both incur inferior results on SSv2. We guess it is because the former setting will make the masked video modeling too easy while the later will degrade to masked image model to some extent. (3) Applying different block masks for different frames (“Frame-Diff”) is also not good, which possibly shares the similar reason as the small temporal length, *i.e.*, making masked video modeling too easy because information can be easily borrowed from adjacent/short-term frames.

## 5. Conclusion and Discussion

Benefited from the BERT pretraining strategy, transformers have dominated various NLP tasks. Recent success of transformers in computer vision motivates the use of BERT pretraining for feature learning. We study how to explore BERT pretraining for video transformers. We introduced BEVT that learns both discriminative spatial and temporal representations using a carefully designed decoupling strategy, which is not only efficient but also effective. BEVT achieved SOTA performance on three video recognition datasets. Although our decoupled design is efficient, BEVT is still computationally intensive and exploring training with constrained resources is left as future work.

**Acknowledgement** This project was supported by National Key R&D Program of China (No. 2021ZD0112805). Y.-G. Jiang was sponsored in part by “Shuguang Program” supported by Shanghai Education Development Foundation and Shanghai Municipal Education Commission (No. 20SG01). Z. Wu was supported by NSFC under Grant No. 62102092.



## 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. In *NeurIPS*, 2021. 2
- [2] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *ICCV*, 2021. 1, 2, 3, 6, 7
- [3] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. BEiT: BERT pre-training of image transformers. In *ICLR*, 2022. 1, 2, 3, 4, 6
- [4] Hangbo Bao, Li Dong, Furu Wei, Wenhui Wang, Nan Yang, Xiaodong Liu, Yu Wang, Jianfeng Gao, Songhao Piao, Ming Zhou, et al. Unilmv2: Pseudo-masked language models for unified language model pre-training. In *ICML*, 2020. 1, 3
- [5] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, 2021. 2, 6, 7, 8
- [6] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *ECCV*, 2018. 3
- [7] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, 2020. 3
- [8] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 3
- [9] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017. 2, 5, 7
- [10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 3
- [11] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *ICCV*, 2021. 3
- [12] Yinpeng Chen, Xiyang Dai, Dongdong Chen, Mengchen Liu, Xiaoyi Dong, Lu Yuan, and Zicheng Liu. Mobileformer: Bridging mobilenet and transformer. *arXiv preprint arXiv:2108.05895*, 2021. 1, 3
- [13] Yunpeng Chen, Yannis Kalantidis, Jianshu Li, Shuicheng Yan, and Jiashi Feng. Multi-fiber networks for video recognition. In *ECCV*, 2018. 2
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2019. 1, 3
- [15] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015. 2
- [16] Jeff Donahue and Karen Simonyan. Large scale adversarial representation learning. In *NeurIPS*, 2019. 3
- [17] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows. *arXiv preprint arXiv:2107.00652*, 2021. 1, 3
- [18] 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. 1, 3, 7
- [19] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua 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*, 2021. 1, 2
- [20] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *ICCV*, 2021. 2, 3, 7
- [21] Quanfu Fan, Chun-Fu Richard Chen, Hilde Kuehne, Marco Pistoia, and David Cox. More is less: learning efficient video representations by big-little network and depthwise temporal aggregation. In *NeurIPS*, 2019. 7
- [22] Christoph Feichtenhofer. X3d: Expanding architectures for efficient video recognition. In *CVPR*, 2020. 2, 7
- [23] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, 2019. 2, 5, 7
- [24] Christoph Feichtenhofer, Haoqi Fan, Bo Xiong, Ross B. Girshick, and Kaiming He. A large-scale study on unsupervised spatiotemporal representation learning. In *CVPR*, 2021. 1
- [25] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *ICLR*, 2018. 3
- [26] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fründ, Peter Yianilos, Moritz Mueller-Freitag, Florian Hoppe, Christian Thureau, Ingo Bax, and Roland Memisevic. The "something something" video database for learning and evaluating visual common sense. In *ICCV*, 2017. 2, 5
- [27] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, koray kavukcuoglu, Remi Munos, and Michal Valko. Bootstrap your own latent - a new approach to self-supervised learning. In *NeurIPS*, 2020. 3
- [28] Dongyoon Han, Sangdoo Yun, Byeongho Heo, and Youngjoon Yoo. Rethinking channel dimensions for efficient model design. In *CVPR*, 2021. 2
- [29] 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. 1, 3
- [30] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 3

- [31] Heeseung Kwon, Manjin Kim, Suha Kwak, and Minsu Cho. Motionsqueeze: Neural motion feature learning for video understanding. In *ECCV*, 2020. 7
- [32] Junnan Li, Pan Zhou, Caiming Xiong, and Steven Hoi. Prototypical contrastive learning of unsupervised representations. In *ICLR*, 2021. 3
- [33] Yingwei Li, Yi Li, and Nuno Vasconcelos. Resound: Towards action recognition without representation bias. In *ECCV*, 2018. 2, 5
- [34] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *ICCV*, 2019. 2, 7
- [35] Yinhan Liu, Mylène 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. *arXiv preprint arXiv:1907.11692*, 2019. 1, 3
- [36] 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 *ICCV*, 2021. 1, 2, 4
- [37] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. *arXiv preprint arXiv:2106.13230*, 2021. 1, 2, 3, 4, 6, 7
- [38] Lingchen Meng, Hengduo Li, Bor-Chun Chen, Shiyi Lan, Zuxuan Wu, Yu-Gang Jiang, and Ser-Nam Lim. Adavit: Adaptive vision transformers for efficient image recognition. In *CVPR*, 2022. 2
- [39] 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*, 2019. 6
- [40] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *CVPR*, 2020. 3
- [41] Daniel Neimark, Omri Bar, Maya Zohar, and Dotan Aselsmann. Video transformer network. *arXiv preprint arXiv:2102.00719*, 2021. 7
- [42] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016. 3
- [43] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In *ICML*, 2018. 2
- [44] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016. 3
- [45] Mandela Patrick, Dylan Campbell, Yuki M. Asano, Ishan Misra Florian Metze, Christoph Feichtenhofer, Andrea Vedaldi, and Joao F. Henriques. Keeping your eye on the ball: Trajectory attention in video transformers. In *NeurIPS*, 2021. 7
- [46] 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. 4, 6, 7
- [47] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NeurIPS*, 2014. 2
- [48] Hao Tan, Jie Lei, Thomas Wolf, and Mohit Bansal. Vimpac: Video pre-training via masked token prediction and contrastive learning. *arXiv preprint arXiv:2106.11250*, 2021. 3
- [49] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training data-efficient image transformers & distillation through attention. In *ICML*, 2021. 2, 3
- [50] Du Tran, Lubomir D Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. C3d: Generic features for video analysis. In *ICCV*, 2015. 2
- [51] Du Tran, Heng Wang, Lorenzo Torresani, and Matt Feiszli. Video classification with channel-separated convolutional networks. In *ICCV*, 2019. 2
- [52] Du Tran, Heng Wang, Lorenzo Torresani, and Matt Feiszli. Video classification with channel-separated convolutional networks. In *ICCV*, 2019. 7
- [53] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, 2018. 2, 7
- [54] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 1
- [55] Ziyu Wan, Jingbo Zhang, Dongdong Chen, and Jing Liao. High-fidelity pluralistic image completion with transformers. In *ICCV*, 2021. 3
- [56] Junke Wang, Zuxuan Wu, Jingjing Chen, Xintong Han, Abhinav Shrivastava, Yu-Gang Jiang, and Ser-Nam Lim. Objectformer for image manipulation detection and localization. In *CVPR*, 2022. 2
- [57] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks for action recognition in videos. *IEEE TPAMI*. 2
- [58] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018. 7
- [59] Thomas Wolf, Julien Chaumond, Lysandre Debut, Victor Sanh, Clement Delangue, Anthony Moi, Pierric Cistac, Morgan Funtowicz, Joe Davison, Sam Shleifer, et al. Transformers: State-of-the-art natural language processing. In *EMNLP*, 2020. 1
- [60] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. Cvt: Introducing convolutions to vision transformers. In *ICCV*, 2021. 1
- [61] Zuxuan Wu, Hengduo Li, Caiming Xiong, Yu-Gang Jiang, and Larry Steven Davis. A dynamic frame selection framework for fast video recognition. *IEEE TPAMI*, 2022. 2
- [62] Zuxuan Wu, Caiming Xiong, Yu-Gang Jiang, and Larry S Davis. Liteeval: A coarse-to-fine framework for resource efficient video recognition. In *NeurIPS*, 2019. 2
- [63] Zhirong Wu, Yuanjun Xiong, Stella Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance-level discrimination. In *CVPR*, 2018. 3
- [64] Zhenda Xie, Yutong Lin, Zhuliang Yao, Zheng Zhang, Qi Dai, Yue Cao, and Han Hu. Self-supervised learning with swin transformers. *arXiv preprint arXiv:2105.04553*, 2021. 6

- [65] Mang Ye, Xu Zhang, Pong C Yuen, and Shih-Fu Chang. Un-supervised embedding learning via invariant and spreading instance feature. In *CVPR*, 2019. 3
- [66] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *CVPR*, 2015. 2
- [67] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stephane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *CVPR*, 2021. 3
- [68] Bowen Zhang, Jiahui Yu, Christopher Fifty, Wei Han, Andrew M Dai, Ruoming Pang, and Fei Sha. Co-training transformer with videos and images improves action recognition. *arXiv preprint arXiv:2112.07175*, 2021. 3
- [69] Chuhan Zhang, Ankush Gupta, and Andrew Zisserman. Temporal query networks for fine-grained video understanding. In *CVPR*, 2021. 7
- [70] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *ECCV*, 2016. 3
- [71] Richard Zhang, Phillip Isola, and Alexei A Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *CVPR*, 2017. 3
- [72] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. In *ECCV*, 2018. 2
- [73] Luowei Zhou, Yingbo Zhou, Jason J Corso, Richard Socher, and Caiming Xiong. End-to-end dense video captioning with masked transformer. In *CVPR*, 2018. 2
- [74] Mohammadreza Zolfaghari, Kamaljeet Singh, and Thomas Brox. Eco: Efficient convolutional network for online video understanding. In *ECCV*, 2018. 2