

# Diffusion-Based Action Recognition Generalizes to Untrained Domains

Rogério Guimarães\* Frank Xiao\* Pietro Perona Markus Marks  
California Institute of Technology

rogerio@caltech.edu {rogerio, frank, perona, marks}@caltech.edu

## Abstract

Humans can recognize the same actions despite large context and viewpoint variations, such as differences between species (walking in spiders vs. horses), viewpoints (egocentric vs. third-person), and contexts (real life vs movies). Current deep learning models struggle with such generalization. We propose using features generated by a Vision Diffusion Model (VDM), aggregated via a transformer, to achieve human-like action recognition across these challenging conditions. We find that generalization is enhanced by the use of a model conditioned on earlier timesteps of the diffusion process to highlight semantic information over pixel level details in the extracted features. We experimentally explore the generalization properties of our approach in classifying actions across animal species, across different viewing angles, and different recording contexts. Our model sets a new state-of-the-art across all three generalization benchmarks, bringing machine action recognition closer to human-like robustness. Project page: [vision.caltech.edu/actiondiff](https://vision.caltech.edu/actiondiff) Code: [github.com/frankyaoxiao/ActionDiff](https://github.com/frankyaoxiao/ActionDiff)

## 1. Introduction

Analyzing the actions and interactions of agents is potentially one of the most useful applications of computer vision: in autonomous driving [15,87], video gaming [37,44], sports analytics [35], transportation safety [82], neurobiology [114], ethology [67,79] and ecology [66] and myriad other domains. Computer Vision is successful in recognizing agents' actions in large-scale datasets automatically [11,23,46,68,79,85]. However, successful systems are trained with massive supervision, which requires expensive and tedious labeling, and do not generalize well to new viewing conditions and domains. Unsupervised deep learning, which has yielded strong results on many tasks in text and images, has yet to demonstrate similar performance for action recognition in video.

\*Equal contribution.

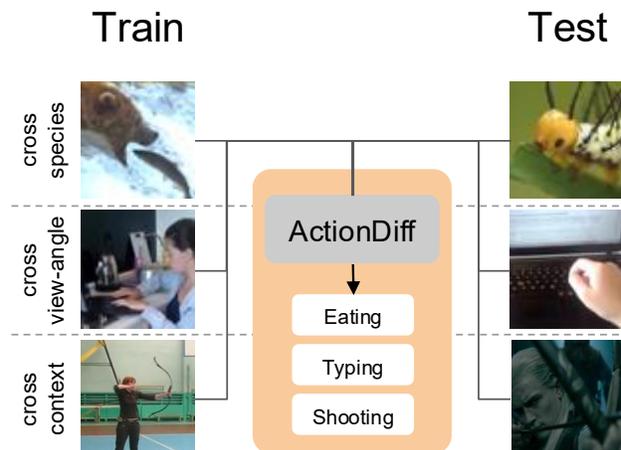


Figure 1. **ActionDiff**. Our method uses the highly semantic features extracted from a frozen Stable Video Diffusion backbone to perform action recognition in tasks that require generalization across different domains. Our model generalizes to new agents (species), view-angles ( $1^{st}$  to  $3^{rd}$ ), and contexts (sports vs. movies) that were not present in the training data.

Amongst the success stories of unsupervised learning for image analysis are semantic segmentation [56], depth estimation [48,112], drawing 2D correspondences [64,88], keypoint estimation [45,49,86,116], and many pixel-level tasks [50,117] and mid-level vision tasks [19]. However, whether the same potential extends beyond the pixel level to a high-level vision task like action recognition is still unclear. We are particularly interested in action classification across species [25,26] and imaging modalities.

Since diffusion models can recognize body parts of animals in images, such as feet and ears, and can generalize such capabilities across species, we are asking if this is also true of actions such as walking or eating? Furthermore, previous literature [21,24,73] shows that earlier steps in the diffusion process focus on low-frequency semantic features, while later steps focus on high-frequency pixel-level details. This indicates that conditioning the diffusion model on these earlier steps can lead to features that are more semantic and robust to generalization. This is in stark contrast

to previous works that use diffusion as a backbone for perception tasks [50, 117], in which the time conditioning aspect is not explored and models are simply conditioned on the last timestep (lowest noise level).

We explore these theoretical notions by comparing features from video diffusion models conditioned on earlier timesteps across various datasets and tasks. We first use these features to train a transformer to recognize actions in one animal species and then use the same classifier to recognize the same behavior in another species. We find that the diffusion model can perform this dramatic semantic shift well, outperforming previously best-performing models. We then test the same strategy on another critical problem in action recognition: generalizing from third-person to first-person video. We find that diffusion model features are more robust to viewpoint change than the previous state-of-the-art (SOTA). Lastly, we test on shifts from recognizing actions from different datasets recorded in different contexts, varying video aspects like framing, view angle, lighting, environment, etc. All in all, we find that our model, based on features of a video diffusion model, handles species, view angle, and context shift better than any of the SOTA models.

Our contributions are as follows:

- We introduce *ActionDiff*, a novel action classification method leveraging video diffusion models as a feature extraction backbone, achieving state-of-the-art performance on video perception tasks.
- We systematically benchmark diffusion-derived features across tasks explicitly designed to evaluate robustness against domain shifts in video data to test if the high semantic level of their features helps in generalization.
- We are the first work to analyze the role of time conditioning in diffusion as a backbone for video, and to experimentally show that features learned in earlier steps are more robust to domain shift. Therefore, we find an application for the theoretical prediction that the features learned at such steps focus on lower-frequency signals and are potentially more semantic and general.

## 2. Related Works

### 2.1. Action Recognition

Action Recognition (AR) is a computer vision task in which a model  $f$  is given a sequence of video frames  $\mathbf{x}$  and must predict  $y = f(\mathbf{x})$ , the label of the action(s) present in the video, given a predefined set of possible actions. Early works on AR used spatial-temporal convolutional architectures [14, 33, 103]. Recent works have used attention-based models like Transformers [2, 62, 89, 95, 102, 111]

and MAMBA [30, 53] to efficiently aggregate long-range spatial-temporal information from different frames.

A more challenging scenario in action recognition is handling domain shift [27, 71, 81], which occurs when a model trained on a source dataset encounters a significantly different target dataset during evaluation. This could be when labeled data is scarce in the test setting or when raw data is unavailable (*e.g.*, when a model will be deployed in an unseen environment). Previous work has tried to address specific domain shifts like different agents (different species performing the same actions) [69, 71], viewpoints (third to first person) [59, 118], and contexts (sports footage vs other sources like movies) [58, 113]. Furthermore, Unsupervised Domain Adaptation (UDA) techniques have been developed to make use of unlabeled data from the target distribution when training the model [16, 78] with labeled data from the source distribution, which helps shift from synthetic to real data. ActionDiff, however, addresses all of these domain shift cases without access to target data and using the same model, relying on the highly semantic features of the pre-trained video diffusion backbone

### 2.2. Video Backbones

Recent works usually employ backbones pre-trained on large-scale datasets that are then fine-tuned on a specific domain. Backbones can be pretrained with supervision in discriminative tasks on large labeled datasets like ImageNet [40, 51, 61, 101] or use self-supervised pretraining techniques like contrastive learning [17, 18, 39, 109], or masked modeling [6, 38, 110]. Recently, a new class of vision-language models like CLIP [75] have also been used, especially for zero-shot tasks like open vocabulary object detection [107, 108, 119].

Video backbones follow a similar trend but advance slower due to challenges such as scarcity of datasets, extra computation and memory required for processing, and the inherent harder challenge of modeling temporal relations. Most of the work in Action Recognition relies on models pretrained on Kinetics-400 [47], using CNNs [14, 32, 57, 63, 90, 96–98] and Transformers [8, 54, 62]. More recent works have used masked autoencoding for self-supervised pretraining using transformers [31, 89, 95, 100, 104].

When training self-supervised models for image or video with a masked modeling reconstruction loss, models need to predict fine-grained details that do not contribute to the overall semantic understanding of an image, like the exact disposition and contours of leaves on a tree branch. It has been suggested that this decreases the semantic level in the features obtained by such models [3]. A few techniques have been developed to increase semantic information in the features of generative models like joint embedding [4, 5, 7, 34] and self-distillation [13, 72]. Still, recent work shows that using features extracted from diffusion backbones can



Figure 2. **Domain Shift Tasks.** We present the three domain-shift tasks we use to measure the generalization performance of ActionDiff. **Left:** samples from the Animal Kingdom dataset [71], which contains examples of actions (eating and swimming) being performed across different animal species. **Middle:** samples from CharadesEgo [81], which contain examples of the same actions (typing and grabbing a pillow) captured from first and third person perspective. **Right:** Samples from UCF-101 (top) and HMDB51 (bottom), which contain examples of the same actions (shooting a bow and kicking a ball) in different contexts. UCF has mostly amateur sports footage, while HMDB also includes other sources (such as movies, TV, and video games.)

be inherently highly semantic in images [88].

### 2.3. Diffusion for Perception Tasks

The quality of diffusion-generated images [76] led to works that tried to use a diffusion model as backbone to extract features from images to be used in downstream tasks like keypoint correspondence, depth estimation, and semantic segmentation [41, 42, 50, 64, 70, 88, 117]. However, given the many design choices that can be made with diffusion, the field is not yet set on the optimal method for extracting such features. Some works extract just the intermediate representations of the denoising U-Net [64, 88], but some also concatenate the cross-attention maps [50, 117], while some extract only such maps [41, 42]. Regarding the text conditioning, some models use no conditioning [64, 88], some use the average embeddings of sentences that are meaningful to the task [117], some use automatic captioners [50], and some optimize the embeddings for the specific task [41, 42]. Some approaches add noise to an image before passing it through the denoiser [88], some do not [50, 117], and some. Questions like what denoiser layer and timestep conditioning should be used have been explored [64] and even whether noise needs to be added to the images before passing them through the denoiser [84].

A natural extension of this work is to use diffusion backbones for video tasks. Several works have tried to process video frames independently using a diffusion backbone to extract features for a downstream task [1, 106]. However, a more suitable approach for video processing is to use a video diffusion model as the backbone—a class of generative models that explicitly incorporates a temporal dimension into their latent space to generate coherent, temporally aligned video frames. Early models used a 3D U-Net to process the extra dimension [43], but most recent works use

spatial-temporal attention blocks, like Stable Video Diffusion [9] and ModelScope [94]. Since they became available, some works have used these pretrained models for downstream video tasks like action recognition [105], semantic segmentation [99], 3D scene understanding [65], and referring video object segmentation [120].

One of the most highlighted properties of diffusion features is their high semantic content, which emerges naturally in the diffusion pretraining process [42, 70, 88]. It has also been shown that this attention to semantics as opposed to fine details increases in the denoiser layers closer to the bottleneck [93] and in earlier timesteps [21, 24, 73]. As mentioned in 2.2, this is a desirable property in pretrained backbones for generalization, but no prior work has taken advantage of these semantic features for this end. Our work shows that the highly semantic features from pretrained Stable Video Diffusion (SVD) outperform SOTA models in tasks that require domain generalization, like adapting to novel agents, viewpoints, and environments.

## 3. Method

### 3.1. Stable Video Diffusion

A Latent Diffusion Model (LDM) consists of four networks: An image encoder  $\mathcal{E}$ , a conditional denoising autoencoder  $\epsilon_\theta$ , a condition encoder  $\tau_\theta$ , and an image decoder  $\mathcal{D}$ . The image encoder and decoder are pretrained to convert an image  $x_0$  into a latent  $z_0$  such that  $\mathcal{E}(x_0) = z_0$  and  $\mathcal{D}(z_0) \approx x_0$ . The condition encoder is trained to encode some type of conditioning  $u$  into a latent  $c = \tau_\theta(u)$  that will guide the diffusion process. The conditional denoising autoencoder is trained to predict the noise  $\epsilon \sim \mathcal{N}(0, 1)$  added to a latent at a given timestep of a forward diffusion process, given the latent  $z_0$ , the timestep  $t$ , and the condi-

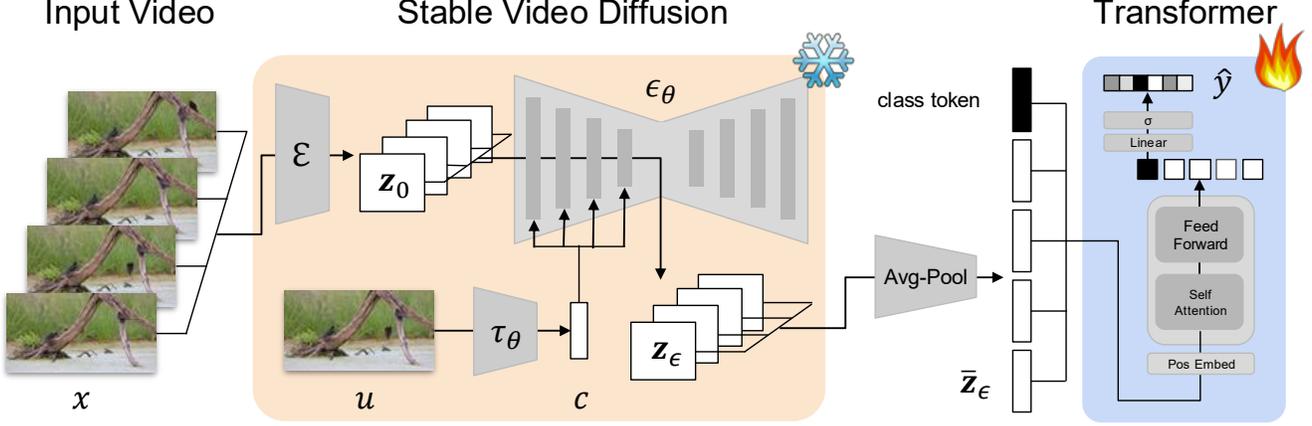


Figure 3. **ActionDiff Architecture.** We split a longer video into shorter segments and extract frame features for each video segment using a frozen Stable Video Diffusion backbone. The video segment frames are encoded into the diffusion latent space by  $\mathcal{E}$ . They are processed together by the denoiser  $\epsilon_\theta$ , in a process guided by a condition  $c$  (the middle frame  $x^{mid}$  embedded by a CLIP encoder  $\tau_\theta$ ) through cross-attention. We extract the outputs of a middle layer  $l$  in the denoiser  $\epsilon_\theta$ , and average pool the outputs across the spatial dimensions to end up with a feature vector for each frame in the video segment. We then collect the sequence of frame features from all video segments and pass them through a learned transformer encoder and a learned class token concatenated to the beginning of the sequence. From the output, we apply a linear layer and a normalization function  $\sigma$  to the class token to obtain probabilities  $\hat{y}$  for each action class.

tioning  $c$ , we minimize the loss:

$$L_{LDM} := \mathbb{E}_{\mathcal{E}(x), \tau_\theta(y), \epsilon \sim \mathcal{N}(0,1), t} \left[ \|\epsilon - \epsilon_\theta(z_t, t, c)\|_2^2 \right], \quad (1)$$

Stable Diffusion (SD) [76] is a pretrained LDM that uses the VQ-GAN encoder and decoder [29] for  $\mathcal{E}$  and  $\mathcal{D}$ , a U-Net [77] for  $\epsilon_\theta$ , and a CLIP text encoder [75] as  $\tau_\theta$  to generate images from text. The U-Net incorporates the conditioning information from  $c$  via cross-attention in each layer in the denoising process.

A latent diffusion model can be modified for video generation by introducing temporal attention layers between the network layers used as  $\epsilon_\theta$  [10]. This way, a sequence of frames  $x$  can have each frame individually encoded through  $\mathcal{E}$  to generate a latent with an added temporal dimension  $z$  and denoised by  $\epsilon_\theta$  through layers that alternate into processing each frame individually and attending to all other frames in the sequence.

Stable Video Diffusion (SVD) [9] uses this architecture to train a video generation diffusion model trained on a large dataset of paired video and captions curated by them. Their model can generate video from text using the CLIP text encoder as  $\tau_\theta$ , but also from a still frame, using the CLIP image encoder instead.

### 3.2. Feature Extraction

Let  $x^m$  be the  $m$ -th frame of a video segment of  $M$  frames, and let the latent associated with that frame be  $z_0^m = \mathcal{E}(x^m)$ . Let us name the whole sequence of latents be  $\mathbf{z}_0 = (z_0^1, z_0^2, \dots, z_0^M)$ . Then we can pass the latent through

the denoiser for a single denoising step conditioned on some  $t$  and  $c$  and we will have  $\mathbf{z}_{t,c} = \epsilon_\theta(\mathbf{z}_0, t, c)$  and we will refer to the  $m$ -th image latent in sequence  $\mathbf{z}_{t,c}$  as  $z_{t,c}^m$ . Notice that after being processed by  $\epsilon_\theta$  the latents are not independent from each other anymore, since the temporal attention of the denoiser will infuse information from all frames in the sequence in each latent.

Since  $\epsilon_\theta$  is usually a deep neural network (A U-Net in the case of SVD), let  $L$  be its number of layers. We can refer to the output of the  $l$ -th layer of the denoiser when it's calculating  $\epsilon_\theta(\mathbf{z}_0, t, c)$  as  $\mathbf{z}_{t,c}^l$  and to the  $m$ -th latent in this sequence of intermediate latents as  $z_{t,c}^{m,l}$ . For a predefined set of  $t, c$ , and  $l$ , we will use  $z_{t,c}^{m,l}$  is a latent of dimensions  $(h_l, w_l, c_l)$ , in which the height, width, and number of channels depend on the layer  $l$ . Let the sequence of latents  $z_{t,c}^{m,l}$  be  $\mathbf{z}_\epsilon$ . We then apply average pooling on the height and width of the feature vector of  $c_l$  channels representing the  $m$ -th frame in our original video segment. We will refer to the average-pooled latents sequence for all video segment frames as  $\bar{\mathbf{z}}_\epsilon$ .

There are memory constraints on GPUs that will not allow the processing of long videos all at once by  $\epsilon_\theta$ . In that case, we extract the features by grouping frames in windows of length  $w$  that fit the GPU and processing each window individually.

### 3.3. Action Recognition Head

Given our sequence of latents  $\bar{\mathbf{z}}_\epsilon$ , we need a model  $f_\phi$  that we will train to predict whether an action is present in

Model	Full Dataset (mAP)	Unseen Species (acc)
<i>Published Methods</i>		
CARe [71]	25.25	39.7
MSQNet [69]	73.10	42.5
Mamba-MSQNet [30]	74.60	-
<i>Self-Supervised Frozen Backbones (our impl.)</i>		
VideoMAEv2 [95]	63.90	37.25
V-JEPA [7]	78.64	51.40
SDv2 [76]	78.66	41.96
ActionDiff (ours)	<b>80.79</b>	<b>51.49</b>

Table 1. **Results on Animal Kingdom dataset.** ActionDiff beats the SOTA and other self-supervised frozen backbones in action recognition both on the full dataset and the unseen species partition, in which the species used at test time were not seen during training.

the video. As mentioned in 3.2, longer videos can’t have all the features being processed at once by  $\epsilon$ , and therefore, earlier features are not injected with information from later features that could refine their representation.

To allow for more information mixing, we process the sequence of features using a transformer encoder  $f_\phi$  [91]. The transformer receives a sequence of tokens as input, and we will use the average-pooled latents  $\bar{z}_\epsilon$  of our Video Diffusion backbone plus a learned class token [28] as the input sequence. In each of its layers, the transformer decoder will perform multi-headed self-attention between the tokens. In this case, the output  $f_\phi(\bar{z}_\epsilon)$  will be a sequence of tokens with the same length as the input sequence ( $\bar{z}_\epsilon + 1$ ).

To predict the final output  $\hat{y}$ , the probability that each action is present in the video, we extract the class token from the transformer output and pass it through a linear layer to obtain logits for each action, then through a normalization function  $\sigma$  to obtain the final probabilities. When on a multi-label task setting, we use a sigmoid as  $\sigma$  and binary cross-entropy loss on each probability. We use softmax as  $\sigma$  and cross entropy loss in a single-label task setting. During training, we also use Focal Loss [60] and Mix-Up [115].

## 4. Experiments

To show the semantic content of the features extracted by SVD, we use it for action recognition in tasks in which domain gaps make it crucial for the model to learn the semantic content of the video. These tasks involve recognizing an action in a different setting than the model was trained on: testing on drastically different agents (different species), view angles (first and third person), and framings (fixed surveillance cameras vs intentional framing on

movies and YouTube).

In all of our experiments, in addition to ActionDiff, we report the previously published results and the performance of training our classifier (the transformer decoder) using features obtained by two other frozen video backbones pretrained on self-supervised tasks. These backbones are VideoMAEv2 [95] pretrained on their *UnlabeledHybrid* dataset and V-JEPA [7] pretrained on their *VideoMix2M* dataset. This ensures we can directly compare the quality of the backbone features against each other since all other experimental details are the same across backbones. Specifically, V-JEPA uses pretraining tasks that aim to learn semantic features to help generalization, which makes it a great comparison for evaluating the semantic level of video diffusion features.

All experiments use the SVD-XT backbone, which is trained to generate clips of up to 25 frames, so we use an extraction window  $w = 25$ . The commonly used number of generative steps of the backbone is 30, so we condition on timestep 20. Notice that we index our timesteps in the direction of the generative process (reverse diffusion), so earlier steps mean conditioning on higher noise levels, and later time steps are closer to a clean image. Therefore, timestep 20 is around 1/3 of the process earlier than the final generative timestep usually used in diffusion-based methods [50, 117] for higher semantic value. We extract features from layer 9, before the last downsampling operation.

We also report Stable Diffusion v2 (SDv2) [76] as a frozen pretrained backbone by running SVD but ignoring the temporal layers, using only the s layers from the Image Diffusion backbone that SVD builds upon and processing each frame without temporal context.

We condition the denoiser on timestep  $t = 20$  and extract the features from the layer that obtained the best result in the experiments with the temporal context. Comparing ActionDiff to SDv2 shows us the contribution of the temporal context encoded in the features by the SVD backbone. For other experimental details an hyperparameters, see Sec. A.1

### 4.1. Cross-Species Generalization

The Animal Kingdom dataset [71] provides 30,100 short videos of 140 actions performed by 650 different animal species. Videos can contain multiple agents, species, and actions. The frequency of an action in this dataset is highly unbalanced across different species. Therefore, to perform well, models must be as close to actor-agnostic as possible to capitalize on training data from actions from one species to another. In Tab. 1, we show that ActionDiff outperforms the SOTA model (Mamba-MSQNet) significantly, and SVD features outperform other unsupervised frozen backbones.

Furthermore, we can directly test the cross-species generalization abilities of our model by training and testing on two separate groups of species. To make our results compa-

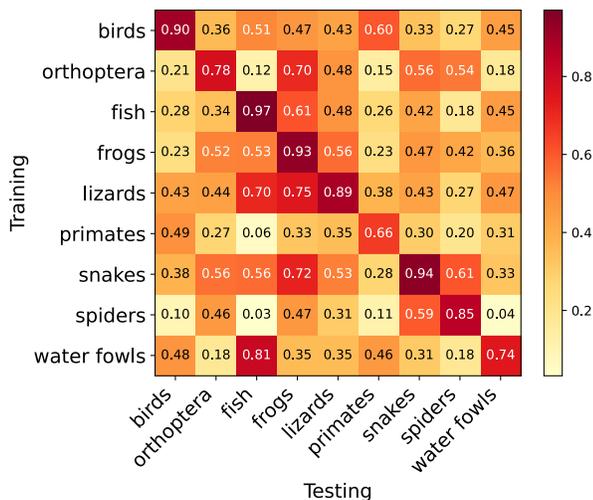


Figure 4. **Species-to-species performance.** We test ActionDiff species-to-species generalization by training and testing in multiple animal type pairs. Results on the main diagonal *e.g.* birds-birds show the testing performance in the training data for reference. Overall, we do not find that similarity between species (physical or taxonomic distance helps the model. Conversely, performance is correlated to shared behaviors that cause similar action distributions in the data. For example, a model trained on water fowls (duck and geese) performs much better on fish than on birds.

able to those reported in [71], we use the same set of four animal types for training (*i.e.*, birds, fishes, frogs, snakes) and five animal types for testing (*i.e.*, lizards, primates, spiders, orthopteran insects, water fowls). We also use only videos labeled with a single action (so we report accuracy instead of mAP). The actions include moving, eating, attending, swimming, sensing, and remaining still. As shown in Tab. 1, ActionDiff significantly outperforms the previous state-of-the-art model, MSQNet (50.53 vs. 45.2), as well as other self-supervised frozen backbone methods.

To further analyze what influences our model performance in cross-species generalization, we measured model accuracy when training in a single species and testing on a different single species so that we can compare how well a species transfers to another and assert whether similarity (or proximity in the taxonomic tree) predicts performance (Fig. 4). Cross-species results correlate more with similar action frequencies between the species than their similarity. For example, training on water fowls works better on fish than in birds (more *swimming*), and training on lizards performs better on frogs than snakes (less *moving*, more *standing still*), which can also be observed by comparing the accuracy matrix (Fig. 4) with the action frequency correlation matrix (Fig. S12) across species. The correlation between accuracy and action frequency correlation is  $\rho = 0.74$  (Fig. S13). ActionDiff performance improvements are not

Model	1 <sup>st</sup> → 1 <sup>st</sup> mAP	3 <sup>rd</sup> → 1 <sup>st</sup> mAP
<i>Fine-tuned Methods (published)</i>		
SSDA [22]	25.8	16.6
AON [80]	25.9	-
ResNet-152 [81]	27.1	19.5
Ego-Exo [55]	30.1	-
EgoVLP [59]	32.1	-
LaViLa [118]	36.1	-
<i>Self-Supervised Frozen Backbones (our impl.)</i>		
VideoMAEv2 [95]	16.9	15.3
V-JEPA [7]	23.4	17.3
SDv2 [76]	35.0	29.4
ActionDiff (ours)	<b>36.5</b>	<b>30.2</b>

Table 2. **Results on the Charades-Ego dataset.** ActionDiff beats the SOTA and other self-supervised frozen backbones in action recognition on 1<sup>st</sup> to 1<sup>st</sup> person and 3<sup>rd</sup> to 1<sup>st</sup> person viewpoint.

Model	U → H Accuracy	H → U Accuracy
<i>Published Methods</i>		
ADA [92]	56.9	72.2
M-ADA [74]	55.6	71.5
Jigsaw [12]	55.2	72.4
VideoDG [113]	59.1	74.9
STDN [58]	60.2	77.1
<i>Self-Supervised Frozen Backbones (our impl.)</i>		
VideoMAEv2 [95]	23.8	26.7
V-JEPA [7]	51.75	60.58
SDv2 [76]	76.5	79.5
ActionDiff (ours)	<b>77.6</b>	<b>81.5</b>

Table 3. **Results on the UCF101 (U) to HMDB51 (H) domain shift and vice-versa.** ActionDiff beats the previous SOTA and other self-supervised frozen backbones when trained on any of the datasets and tested on the other. Results for ADA, M-ADA, and Jigsaw were obtained by [113].

due to simply predicting the most likely action, as it significantly outperforms a Naïve Bayes approach that predicts the most frequent action in the train set (Fig. S16), and it beats the other methods in the unseen species task. Action frequencies across species are shown in Fig. S14.

## 4.2. Cross-Viewpoint Generalization

The Charades-Ego dataset [81] consists of 68.8 hours of real-world video footage containing 68,536 annotated action instances captured simultaneously from paired first-person and third-person perspectives. As shown in Tab. 2,

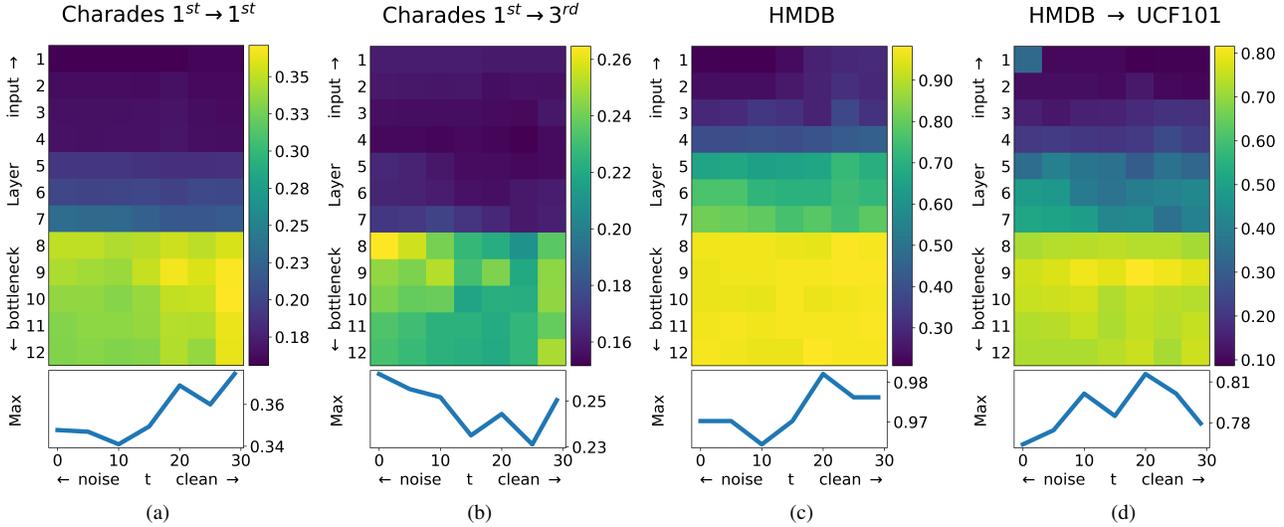


Figure 5. **Diffusion Layer and Timestep Conditioning.** We test the performance of our model with diffusion features extracted from different layers and conditioned on different timesteps to analyze the difference between the best features for in- and out-of-domain tasks. We use two example tasks where the model can train on the same train set for in and out-of-domain tasks. Each heatmap shows the results (mAP for CharadesEgo and acc for HMDB to UCF) for each layer (y-axis) and timestep (x-axis) on a task, and the plot below each heatmap shows the best result across layers obtained for each timestep. When we analyze which features are best for in-domain vs out-of-domain testing, we see a shift toward those obtained with earlier timesteps when we test out-of-domain. Layers are indexed in the direction from the input to the bottleneck, and timesteps are indexed in the direction of the generative process (reverse diffusion).

our ActionDiff model establishes a new state-of-the-art when trained solely on first-person videos, despite utilizing an SVD backbone that was not fine-tuned specifically on egocentric data. This result is particularly noteworthy since the previous state-of-the-art, LaViLa [118], relied extensively on pretraining with Ego4D [36], one of the largest egocentric video datasets available. Furthermore, we also show how the semantic knowledge of diffusion can be used to generalize across viewpoints since our model also beats SOTA results when training the classifier on third-person views and testing on first-person videos.

### 4.3. Cross-Context Generalization

UCF101 [83] is a dataset containing 101 classes of user-uploaded human actions. The HMDB51 [52] is also a human action detection dataset over 51 classes with videos sourced from movies and YouTube videos. In our experiments, we use videos that contain one of the twelve action classes shared between the two datasets, as assembled by [16]. The videos from UCF101 are user uploaded and captured in often similar contexts, mostly footage from amateur sports. The videos from HMDB51 are sourced from more unconstrained contexts (such as TV, movies, and video games) that vary environment, lighting, camera framing and angle, as shown in Fig. 2.

Tab. 3 shows our model beats the previous SOTA when adapting from UCF101 to HMDB51, which was obtained by a video domain generalization method [58], and the other

self-supervised frozen backbones we tested.

### 4.4. Diffusion Layer and Timestep Conditioning

For analysis, we measure the effect of layer and timestep conditioning on the features extracted by the diffusion backbone by grid-searching across all layers  $l$  and multiple timesteps  $t$ . Results in Fig. 5 show the results of evaluating those features in two pairs of in-domain (Charades Ego 1<sup>st</sup> to 1<sup>st</sup> person (a) and HMDB51 (b)) and out-of-domain (Charades Ego 1<sup>st</sup> to 3<sup>rd</sup> person (c) and HMDB51 to UCF101 (d)). When we compare the testing in-domain vs out-of-domain ((a) to (b) and (c) to (d)), we see that earlier steps perform better than the last steps out of domain, a trend not observed in-domain, where the last steps are slightly better. This corroborates previous work showing that earlier diffusion steps focus more on coarse shapes and low-frequency signals, while later steps focus more on fine details and high-frequency signals [21, 24, 73].

Our findings indicate optimal performance at earlier timesteps than reported in previous studies, which identified timesteps around 80% to 90% of the diffusion process as optimal for keypoint correspondence in images [64]. This difference suggests that earlier timesteps are particularly beneficial for enhancing domain generalization in out-of-domain tasks. Heatmaps (a) and (c) show that timestep conditioning play a very minor role in in-domain tasks across the same layer. It is only in heatmaps (b) and (d), out of domain, that timestep makes a difference within the same layer. This is

potentially why effects of timestep conditioning have not been previously reported.

Different tasks can vary in the location of the optimal UNet layer, but all tasks show a trend of better performance in deeper layers. In Tab. S5 we also verify that in all the tasks reported in our experiments, the deeper layer we are using (9) performs better than earlier layers (3 and 6).

#### 4.5. Ablations

To test the importance of different parts of our full model, we conduct ablations on the Animal Kingdom dataset with respect to the classifier, extraction window length, and diffusion conditioning. Results are reported in Tab. 4 and were performed using an SVD backbone pre-trained for generating smaller videos (14 frames) than the SVD-XT backbone we use in the other experiments.

In the first part of Tab. 4, we see that the transformer encoder has great gains compared to a simple linear or MLP classifier. This shows the importance of the attention mechanism to aggregate the information from the multiple latents representing all the input video frames. In the linear and MLP setting, these frames are concatenated.

Then, we analyze the effect of using different extraction window lengths, which is the number of frames encoded together by the backbone and provide all the temporal context that video diffusion can encode in the features. Throughout all the experiments in this paper, SVD outperforms SDv2, which shows that the video backbone’s temporal context provides an advantage compared to the image backbone. However, this advantage saturates quickly with the number of frames in the extraction window, and we see no significant difference in using 10, 16, or 21 frames.

Next, we measure the effect of two training techniques we used to address train imbalance and help generalization. Due to class imbalance in Animal Kingdom among actions, focal loss [60] provides gains in performance to our model. Mix-Up [115] also helps since this dataset requires generalization across agents, and thus, memorization is hurtful.

Finally, we tested different conditioning schemes during extraction using our video diffusion backbone. By default, SVD uses a single frame to condition its video generation process, and the results in our previous experiments use the OpenCLIP image embedding [20] of the middle frame in a window as the conditioning during feature extraction. This method is better than unconditional extraction or conditioning on the OpenCLIP text embedding of the names of all the possible actions we want to recognize in a dataset.

### 5. Discussion

We present ActionDiff, a novel method that utilizes features from video diffusion models for action recognition. To achieve this, we explore the best way to extract and integrate features from a video diffusion model over time.

Classifier	Window	Loss	Cond	mAP
Linear	10	Focal	Frame	73.0
MLP	10	Focal	Frame	75.2
Transformer	10	Focal	Frame	80.8
Transformer	10	Focal	Frame	80.8
Transformer	16	Focal	Frame	80.9
Transformer	21	Focal	Frame	80.8
Transformer	10	BCE	Frame	79.9
Transformer	10	Focal	Frame	80.7
Transformer	10	Focal	None	77.4
Transformer	10	Focal	Action	74.0
Transformer	10	Focal	Frame	80.8

Table 4. **Ablation Studies.** We ablate different parts of our model and vary hyperparameters on the Animal Kingdom Full dataset to show their impact on performance. We test different classifier architectures, extraction window lengths for the backbone, the role of focal loss in an unbalanced dataset, and different conditioning choices available to the Stable Video Diffusion backbone.

We identify the optimal diffusion timestep and layer to extract video diffusion features for action recognition. We then use these features as input to a transformer. We use this architecture to test performance on some major action recognition challenges. We first test on *Cross-Species Generalization*, i.e., how well actions can be recognized across different species, using the Animal Kingdom dataset. We find that video diffusion features are exceptionally suited for this task, and our approach sets the current state-of-the-art for this challenging task. Next, using the Charades-Ego dataset, we test how well our approach copes with *View Angle Shifts*. Again, features from diffusion models perform very well, beating SOTA methods. Finally, using the UCF-HMDB dataset, we show that the early features from diffusion models also generalize well, identifying the same actions being performed in different contexts, beating SOTA methods available. Overall, our method sets the current SOTA in multiple action recognition tasks, thereby showing the utility of using diffusion model features. It is noteworthy that, in the Animal Kingdom dataset, V-JEPA closely matches our results despite being built upon fundamentally different design principles, highlighting the potential of diverse approaches in representation learning. A notable limitation of our approach is its substantial computational cost, primarily due to the diffusion-based feature extraction, which may hinder scalability and practical deployment in resource-constrained scenarios. Future research could explore diffusion-based features for other video-related tasks beyond the scope of our current experiments, such as video grounding, temporal localization, or activity recognition. Additionally, integrating parameter-efficient fine-tuning techniques, such as Low-Rank Adaptation (LoRA), to reduce computational overhead represents an exciting direction for making our model more accessible and efficient in practical scenarios.

## References

- [1] Amirhossein Alimohammadi, Sauradip Nag, Saeid Asgari Taghanaki, Andrea Tagliasacchi, Ghassan Hamarneh, and Ali Mahdavi-Amiri. SMITE: segment me in time. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. 3
- [2] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lucic, and Cordelia Schmid. ViViT: A Video Vision Transformer. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6816–6826, Montreal, QC, Canada, Oct. 2021. IEEE. 2
- [3] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Mike Rabbat, and Nicolas Ballas. Masked siamese networks for label-efficient learning. In *European conference on computer vision*, pages 456–473. Springer, 2022. 2
- [4] Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15619–15629, 2023. 2
- [5] Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatao Gu, and Michael Auli. Data2vec: A general framework for self-supervised learning in speech, vision and language. In *International conference on machine learning*, pages 1298–1312. PMLR, 2022. 2
- [6] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: BERT pre-training of image transformers. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. 2
- [7] Adrien Bardes, Quentin Garrido, Jean Ponce, Xinlei Chen, Michael Rabbat, Yann LeCun, Mido Assran, and Nicolas Ballas. Revisiting feature prediction for learning visual representations from video. *Trans. Mach. Learn. Res.*, 2024, 2024. 2, 5, 6
- [8] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 813–824. PMLR, 2021. 2
- [9] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, Varun Jampani, and Robin Rombach. Stable Video Diffusion: Scaling Latent Video Diffusion Models to Large Datasets, Nov. 2023. arXiv:2311.15127 [cs]. 3, 4
- [10] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 22563–22575. IEEE, 2023. 4
- [11] Kristin Branson, Alice A Robie, John Bender, Pietro Perona, and Michael H Dickinson. High-throughput ethomics in large groups of drosophila. *Nature methods*, 6(6):451–457, 2009. 1
- [12] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2229–2238, 2019. 6
- [13] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 9630–9640. IEEE, 2021. 2
- [14] João Carreira and Andrew Zisserman. Quo vadis, action recognition? A new model and the kinetics dataset. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 4724–4733. IEEE Computer Society, 2017. 2
- [15] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, et al. Argoverse: 3d tracking and forecasting with rich maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8748–8757, 2019. 1
- [16] Min-Hung Chen, Zsolt Kira, Ghassan Alregib, Jaekwon Yoo, Ruxin Chen, and Jian Zheng. Temporal attentive alignment for large-scale video domain adaptation. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 6320–6329. IEEE, 2019. 2, 7
- [17] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR, 2020. 2
- [18] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 15750–15758. Computer Vision Foundation / IEEE, 2021. 2
- [19] Xuweiyi Chen, Markus Marks, and Zezhou Cheng. Probing the mid-level vision capabilities of self-supervised learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025*, pages 30095–30105. Computer Vision Foundation / IEEE, 2025. 1
- [20] Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for contrastive language-image

- learning. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2829, June 2023. arXiv:2212.07143 [cs]. 8
- [21] Jooyoung Choi, Jungbeom Lee, Chaehun Shin, Sungwon Kim, Hyunwoo Kim, and Sungroh Yoon. Perception prioritized training of diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 11462–11471. IEEE, 2022. 1, 3, 7
- [22] Jinwoo Choi, Gaurav Sharma, Manmohan Chandraker, and Jia-Bin Huang. Unsupervised and semi-supervised domain adaptation for action recognition from drones. In *IEEE Winter Conference on Applications of Computer Vision, WACV 2020, Snowmass Village, CO, USA, March 1-5, 2020*, pages 1706–1715. IEEE, 2020. 6
- [23] Heiko Dankert, Liming Wang, Eric D Hoopfer, David J Anderson, and Pietro Perona. Automated monitoring and analysis of social behavior in drosophila. *Nature methods*, 6(4):297–303, 2009. 1
- [24] Giannis Daras and Alexandros G. Dimakis. Multiresolution Textual Inversion, Nov. 2022. arXiv:2211.17115 [cs]. 1, 3, 7
- [25] Sandeep Robert Datta, David J Anderson, Kristin Branson, Pietro Perona, and Andrew Leifer. Computational neuroethology: a call to action. *Neuron*, 104(1):11–24, 2019. 1
- [26] Anthony I Dell, John A Bender, Kristin Branson, Iain D Couzin, Gonzalo G de Polavieja, Lucas PJJ Noldus, Alfonso Pérez-Escudero, Pietro Perona, Andrew D Straw, Martin Wikelski, et al. Automated image-based tracking and its application in ecology. *Trends in ecology & evolution*, 29(7):417–428, 2014. 1
- [27] Andong Deng, Taojiannan Yang, and Chen Chen. A large-scale study of spatiotemporal representation learning with a new benchmark on action recognition. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, pages 20462–20474. IEEE, 2023. 2
- [28] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and Thamar Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics, 2019. 5
- [29] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 12873–12883. Computer Vision Foundation / IEEE, 2021. 4
- [30] Edoardo Fazzari, Donato Romano, Fabrizio Falchi, and Cesare Stefanini. Selective state models are what you need for animal action recognition. *Ecological Informatics*, 85:102955, Mar. 2025. 2, 5
- [31] Christoph Feichtenhofer, Haoqi Fan, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. 2
- [32] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 6201–6210. IEEE, 2019. 2
- [33] Christoph Feichtenhofer, Axel Pinz, and Richard P. Wildes. Spatiotemporal multiplier networks for video action recognition. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 7445–7454. IEEE Computer Society, 2017. 2
- [34] Quentin Garrido, Mahmoud Assran, Nicolas Ballas, Adrien Bardes, Laurent Najman, and Yann LeCun. Learning and Leveraging World Models in Visual Representation Learning, Mar. 2024. arXiv:2403.00504 [cs]. 2
- [35] Indrajeet Ghosh, Sreenivasan Ramasamy Ramamurthy, Avijoy Chakma, and Nirmalya Roy. Sports analytics review: Artificial intelligence applications, emerging technologies, and algorithmic perspective. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 13(5):e1496, 2023. 1
- [36] Kristen Grauman, Andrew Westbury, Eugene Byrne, Vincent Cartillier, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Devansh Kukreja, Miao Liu, Xingyu Liu, Miguel Martin, Tushar Nagarajan, Ilija Radosavovic, Santhosh Kumar Ramakrishnan, Fiona Ryan, Jayant Sharma, Michael Wray, Mengmeng Xu, Eric Zhongcong Xu, Chen Zhao, Siddhant Bansal, Dhruv Batra, Sean Crane, Tien Do, Morrie Doulaty, Akshay Erappalli, Christoph Feichtenhofer, Adriano Fragomeni, Qichen Fu, Abraham Gebreselasie, Cristina González, James Hillis, Xuhua Huang, Yifei Huang, Wenqi Jia, Weslie Khoo, Jáchym Kolár, Satwik Kottur, Anurag Kumar, Federico Landini, Chao Li, Yanghao Li, Zhenqiang Li, Karttikeya Mangalam, Raghava Modhugu, Jonathan Munro, Tullie Murrell, Takumi Nishiyasu, Will Price, Paola Ruiz Puentes, Merey Ramazanova, Leda Sari, Kiran K. Somasundaram, Audrey Southerland, Yusuke Sugano, Ruijie Tao, Minh Vo, Yuchen Wang, Xindi Wu, Takuma Yagi, Ziwei Zhao, Yunyi Zhu, Pablo Arbeláez, David Crandall, Dima Damen, Giovanni Maria Farinella, Christian Fuegen, Bernard Ghanem, Vamsi Krishna Ithapu, C. V. Jawahar, Hanbyul Joo, Kris Kitani, Haizhou Li, Richard A. Newcombe, Aude Oliva, Hyun Soo Park, James M. Rehg, Yoichi Sato, Jianbo Shi, Mike Zheng Shou, Antonio Torralba, Lorenzo Torresani, Mingfei Yan, and Jitendra Malik. Ego4d: Around the world in 3,600 hours of egocentric video. *IEEE Trans. Pattern Anal. Mach. Intell.*, 47(11):9468–9509, 2025. 7
- [37] William H. Guss, Brandon Houghton, Nicholay Topin, Phillip Wang, Cayden R. Codel, Manuela Veloso, and

- Ruslan Salakhutdinov. Minerl: A large-scale dataset of minecraft demonstrations. In Sarit Kraus, editor, *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 2442–2448. ijcai.org, 2019. 1
- [38] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked Autoencoders Are Scalable Vision Learners, Dec. 2021. arXiv:2111.06377 [cs]. 2
- [39] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 9726–9735. Computer Vision Foundation / IEEE, 2020. 2
- [40] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society, 2016. 2
- [41] Eric Hedlin, Gopal Sharma, Shweta Mahajan, Xingzhe He, Hossam Isack, Abhishek Kar, Helge Rhodin, Andrea Tagliasacchi, and Kwang Moo Yi. Unsupervised keypoints from pretrained diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 22820–22830. IEEE, 2024. 3
- [42] Eric Hedlin, Gopal Sharma, Shweta Mahajan, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. Unsupervised semantic correspondence using stable diffusion. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. 3
- [43] Jonathan Ho, Tim Salimans, Alexey A. Gritsenko, William Chan, Mohammad Norouzi, and David J. Fleet. Video diffusion models. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. 3
- [44] Katja Hofmann. Minecraft as ai playground and laboratory. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*, pages 1–1, 2019. 1
- [45] Tomas Jakab, Ankush Gupta, Hakan Bilen, and Andrea Vedaldi. Self-supervised learning of interpretable keypoints from unlabelled videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8787–8797, 2020. 1
- [46] Hueihan Jhuang, Thomas Serre, Lior Wolf, and Tomaso A. Poggio. A biologically inspired system for action recognition. In *IEEE 11th International Conference on Computer Vision, ICCV 2007, Rio de Janeiro, Brazil, October 14-20, 2007*, pages 1–8. IEEE Computer Society, 2007. 1
- [47] 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, May 2017. arXiv:1705.06950 [cs]. 2
- [48] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9492–9502, 2024. 1
- [49] Daniel Khalil, Christina Liu, Pietro Perona, Jennifer J. Sun, and Markus Marks. Learning keypoints for multi-agent behavior analysis using self-supervision. In *IEEE/CVF Winter Conference on Applications of Computer Vision, WACV 2025, Tucson, AZ, USA, February 26 - March 6, 2025*, pages 578–588. IEEE, 2025. 1
- [50] Neehar Kondapaneni, Markus Marks, Manuel Knott, Rogério Guimarães, and Pietro Perona. Text-image alignment for diffusion-based perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13883–13893, 2024. 1, 2, 3, 5
- [51] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In F. Pereira, C. J. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012. 2
- [52] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *2011 International conference on computer vision*, pages 2556–2563. IEEE, 2011. 7
- [53] Kunchang Li, Xinhao Li, Yi Wang, Yinan He, Yali Wang, Limin Wang, and Yu Qiao. Videomamba: State space model for efficient video understanding. In Ales Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol, editors, *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part XXVI*, volume 15084 of *Lecture Notes in Computer Science*, pages 237–255. Springer, 2024. 2
- [54] Kunchang Li, Yali Wang, Peng Gao, Guanglu Song, Yu Liu, Hongsheng Li, and Yu Qiao. Uniformer: Unified transformer for efficient spatial-temporal representation learning. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. 2
- [55] Yanghao Li, Tushar Nagarajan, Bo Xiong, and Kristen Grauman. Ego-exo: Transferring visual representations from third-person to first-person videos. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 6943–6953. Computer Vision Foundation / IEEE, 2021. 6
- [56] Ziyi Li, Qinye Zhou, Xiaoyun Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. Open-vocabulary object segmentation with diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7667–7676, 2023. 1

- [57] Ji Lin, Chuang Gan, and Song Han. TSM: temporal shift module for efficient video understanding. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 7082–7092. IEEE, 2019. 2
- [58] Kun-Yu Lin, Jia-Run Du, Yipeng Gao, Jiaming Zhou, and Wei-Shi Zheng. Diversifying spatial-temporal perception for video domain generalization. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. 2, 6, 7
- [59] Kevin Qinghong Lin, Jinpeng Wang, Mattia Soldan, Michael Wray, Rui Yan, Eric Zhongcong Xu, Difei Gao, Rong-Cheng Tu, Wenzhe Zhao, Weijie Kong, Chengfei Cai, Hongfa Wang, Dima Damen, Bernard Ghanem, Wei Liu, and Mike Zheng Shou. Egocentric video-language pretraining. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. 2, 6
- [60] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 2999–3007. IEEE Computer Society, 2017. 5, 8
- [61] 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 *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9992–10002, Montreal, QC, Canada, Oct. 2021. IEEE. 2
- [62] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video Swin Transformer. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3192–3201, New Orleans, LA, USA, June 2022. IEEE. 2
- [63] Zhaoyang Liu, Limin Wang, Wayne Wu, Chen Qian, and Tong Lu. TAM: temporal adaptive module for video recognition. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 13688–13698. IEEE, 2021. 2
- [64] Grace Luo, Lisa Dunlap, Dong Huk Park, Aleksander Holynski, and Trevor Darrell. Diffusion hyperfeatures: Searching through time and space for semantic correspondence. *Advances in Neural Information Processing Systems*, 36:47500–47510, 2023. 1, 3, 7
- [65] Yunze Man, Shuhong Zheng, Zhipeng Bao, Martial Hebert, Liangyan Gui, and Yu-Xiong Wang. Lexicon3d: Probing visual foundation models for complex 3d scene understanding. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang, editors, *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. 3
- [66] Kelly Ray Mannion, Elizabeth F Ballare, Markus Marks, and Thibaud Gruber. A multicomponent approach to studying cultural propensities during foraging in the wild. *Journal of Animal Ecology*, (8), 2023. 1
- [67] Markus Marks, Qiuhan Jin, Oliver Sturman, Lukas von Ziegler, Sepp Kollmorgen, Wolfger von der Behrens, Valerio Mante, Johannes Bohacek, and Mehmet Fatih Yanik. Deep-learning-based identification, tracking, pose estimation and behaviour classification of interacting primates and mice in complex environments. *Nature Machine Intelligence*, 4(4):331–340, 2022. 1
- [68] Alexander Mathis, Pranav Mamidanna, Kevin M Cury, Taiga Abe, Venkatesh N Murthy, Mackenzie Weyand Mathis, and Matthias Bethge. Deeplabcut: markerless pose estimation of user-defined body parts with deep learning. *Nature neuroscience*, 21(9):1281–1289, 2018. 1
- [69] Anindya Mondal, Sauradip Nag, Joaquin M. Prada, Xiatian Zhu, and Anjan Dutta. Actor-agnostic Multi-label Action Recognition with Multi-modal Query. In *2023 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pages 784–794, Oct. 2023. arXiv:2307.10763 [cs]. 2, 5
- [70] Koichi Namekata, Amirmojtaba Sabour, Sanja Fidler, and Seung Wook Kim. Emerdiff: Emerging pixel-level semantic knowledge in diffusion models. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. 3
- [71] Xun Long Ng, Kian Eng Ong, Qichen Zheng, Yun Ni, Si Yong Yeo, and Jun Liu. Animal kingdom: A large and diverse dataset for animal behavior understanding. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 19001–19012. IEEE, 2022. 2, 3, 5, 6
- [72] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jégou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision. *Trans. Mach. Learn. Res.*, 2024, 2024. 2
- [73] Yong-Hyun Park, Mingi Kwon, Jaewoong Choi, Junghyo Jo, and Youngjung Uh. Understanding the latent space of diffusion models through the lens of riemannian geometry. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. 1, 3, 7
- [74] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *Proceedings of*

- the *IEEE/CVF conference on computer vision and pattern recognition*, pages 12556–12565, 2020. 6
- [75] 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 Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR, 2021. 2, 4
- [76] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 10674–10685. IEEE, 2022. 3, 4, 5, 6
- [77] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells III, and Alejandro F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2015 - 18th International Conference Munich, Germany, October 5 - 9, 2015, Proceedings, Part III*, volume 9351 of *Lecture Notes in Computer Science*, pages 234–241. Springer, 2015. 4
- [78] Aadarsh Sahoo, Rutav Shah, Rameswar Panda, Kate Saenko, and Abir Das. Contrast and mix: Temporal contrastive video domain adaptation with background mixing. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 23386–23400, 2021. 2
- [79] Cristina Segalin, Jalani Williams, Tomomi Karigo, May Hui, Moriel Zelikowsky, Jennifer J Sun, Pietro Perona, David J Anderson, and Ann Kennedy. The mouse action recognition system (mars) software pipeline for automated analysis of social behaviors in mice. *Elife*, 10:e63720, 2021. 1
- [80] Gunnar A. Sigurdsson, Abhinav Gupta, Cordelia Schmid, Ali Farhadi, and Karteek Alahari. Actor and observer: Joint modeling of first and third-person videos. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 7396–7404. Computer Vision Foundation / IEEE Computer Society, 2018. 6
- [81] Gunnar A. Sigurdsson, Abhinav Gupta, Cordelia Schmid, Ali Farhadi, and Karteek Alahari. Charades-ego: A large-scale dataset of paired third and first person videos. *CoRR*, abs/1804.09626, 2018. 2, 3, 6
- [82] Philippe Barbosa Silva, Michelle Andrade, and Sara Ferreira. Machine learning applied to road safety modeling: A systematic literature review. *Journal of traffic and transportation engineering (English edition)*, 7(6):775–790, 2020. 1
- [83] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 7
- [84] Nick Stracke, Stefan Andreas Baumann, Kolja Bauer, Frank Fundel, and Björn Ommer. Cleandiff: Diffusion features without noise. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025*, pages 117–127. Computer Vision Foundation / IEEE, 2025. 3
- [85] Jennifer J Sun, Markus Marks, Andrew Wesley Ulmer, Dipam Chakraborty, Brian Geuther, Edward Hayes, Heng Jia, Vivek Kumar, Sebastian Oleszko, Zachary Partridge, et al. Mabe22: a multi-species multi-task benchmark for learned representations of behavior. In *International Conference on Machine Learning*, pages 32936–32990. PMLR, 2023. 1
- [86] Jennifer J Sun, Serim Ryou, Roni H Goldshmid, Brandon Weissbourd, John O Dabiri, David J Anderson, Ann Kennedy, Yisong Yue, and Pietro Perona. Self-supervised keypoint discovery in behavioral videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2171–2180, 2022. 1
- [87] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2446–2454, 2020. 1
- [88] Luming Tang, Menglin Jia, Qianqian Wang, Cheng Perng Phoo, and Bharath Hariharan. Emergent correspondence from image diffusion. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. 1, 3
- [89] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. 2
- [90] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 6450–6459. Computer Vision Foundation / IEEE Computer Society, 2018. 2
- [91] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Informa-*

- tion Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, 2017. 5
- [92] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. *Advances in neural information processing systems*, 31, 2018. 6
- [93] Andrey Voynov, Qinghao Chu, Daniel Cohen-Or, and Kfir Aberman. P+: Extended Textual Conditioning in Text-to-Image Generation, July 2023. arXiv:2303.09522 [cs]. 3
- [94] Jiuniu Wang, Hangjie Yuan, Dayou Chen, Yingya Zhang, Xiang Wang, and Shiwei Zhang. ModelScope Text-to-Video Technical Report, Aug. 2023. arXiv:2308.06571 [cs]. 3
- [95] Limin Wang, Bingkun Huang, Zhiyu Zhao, Zhan Tong, Yinan He, Yi Wang, Yali Wang, and Yu Qiao. Videomae V2: scaling video masked autoencoders with dual masking. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 14549–14560. IEEE, 2023. 2, 5, 6
- [96] Limin Wang, Wei Li, Wen Li, and Luc Van Gool. Appearance-and-relation networks for video classification. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 1430–1439. Computer Vision Foundation / IEEE Computer Society, 2018. 2
- [97] Limin Wang, Zhan Tong, Bin Ji, and Gangshan Wu. TDN: temporal difference networks for efficient action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 1895–1904. Computer Vision Foundation / IEEE, 2021. 2
- [98] 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 Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part VIII*, volume 9912 of *Lecture Notes in Computer Science*, pages 20–36. Springer, 2016. 2
- [99] Qian Wang, Abdelrahman Eldesokey, Mohit Mendiratta, Fangneng Zhan, Adam Kortylewski, Christian Theobalt, and Peter Wonka. Zero-Shot Video Semantic Segmentation based on Pre-Trained Diffusion Models, May 2024. arXiv:2405.16947 [cs]. 3
- [100] Rui Wang, Dongdong Chen, Zuxuan Wu, Yinpeng Chen, Xiyang Dai, Mengchen Liu, Yu-Gang Jiang, Luowei Zhou, and Lu Yuan. BEVT: BERT pretraining of video transformers. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 14713–14723. IEEE, 2022. 2
- [101] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid Vision Transformer: A Versatile Backbone for Dense Prediction without Convolutions. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 548–558, Montreal, QC, Canada, Oct. 2021. IEEE. 2
- [102] Yi Wang, Kunchang Li, Yizhuo Li, Yinan He, Bingkun Huang, Zhiyu Zhao, Hongjie Zhang, Jilan Xu, Yi Liu, Zun Wang, Sen Xing, Guo Chen, Junting Pan, Jiashuo Yu, Yali Wang, Limin Wang, and Yu Qiao. InternVideo: General Video Foundation Models via Generative and Discriminative Learning, Dec. 2022. arXiv:2212.03191 [cs]. 2
- [103] Yunbo Wang, Mingsheng Long, Jianmin Wang, and Philip S. Yu. Spatiotemporal pyramid network for video action recognition. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 2097–2106. IEEE Computer Society, 2017. 2
- [104] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan L. Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 14648–14658. IEEE, 2022. 2
- [105] Zejia Weng, Xitong Yang, Zhen Xing, Zuxuan Wu, and Yu-Gang Jiang. Genrec: Unifying video generation and recognition with diffusion models. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang, editors, *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. 3
- [106] Lehong Wu, Lilang Lin, Jiahang Zhang, Yiyang Ma, and Jiaying Liu. MacDiff: Unified Skeleton Modeling with Masked Conditional Diffusion. In Aleš Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol, editors, *Computer Vision – ECCV 2024*, volume 15084, pages 110–128. Springer Nature Switzerland, Cham, 2025. Series Title: Lecture Notes in Computer Science. 3
- [107] Size Wu, Wenwei Zhang, Lumin Xu, Sheng Jin, Xiangtai Li, Wentao Liu, and Chen Change Loy. Clipself: Vision transformer distills itself for open-vocabulary dense prediction. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. 2
- [108] Xiaoshi Wu, Feng Zhu, Rui Zhao, and Hongsheng Li. CORA: adapting CLIP for open-vocabulary detection with region prompting and anchor pre-matching. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 7031–7040. IEEE, 2023. 2
- [109] Zhirong Wu, Yuanjun Xiong, Stella X. Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 3733–3742. Computer Vision Foundation / IEEE Computer Society, 2018. 2
- [110] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*,

- CVPR 2022, New Orleans, LA, USA, June 18-24, 2022, pages 9643–9653. IEEE, 2022. [2](#)
- [111] Shen Yan, Xuehan Xiong, Anurag Arnab, Zhichao Lu, Mi Zhang, Chen Sun, and Cordelia Schmid. Multiview Transformers for Video Recognition. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3323–3333, New Orleans, LA, USA, June 2022. IEEE. [2](#)
- [112] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10371–10381, 2024. [1](#)
- [113] Zhiyu Yao, Yunbo Wang, Jianmin Wang, Philip S Yu, and Mingsheng Long. Videodg: Generalizing temporal relations in videos to novel domains. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):7989–8004, 2021. [2](#), [6](#)
- [114] Moriel Zelikowsky, May Hui, Tomomi Karigo, Andrea Choe, Bin Yang, Mario R Blanco, Keith Beadle, Viviana Gradinaru, Benjamin E Deverman, and David J Anderson. The neuropeptide tac2 controls a distributed brain state induced by chronic social isolation stress. *Cell*, 173(5):1265–1279, 2018. [1](#)
- [115] Hongyi Zhang, Moustapha Cissé, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. [5](#), [8](#)
- [116] Yuting Zhang, Yijie Guo, Yixin Jin, Yijun Luo, Zhiyuan He, and Honglak Lee. Unsupervised discovery of object landmarks as structural representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2694–2703, 2018. [1](#)
- [117] Wenliang Zhao, Yongming Rao, Zuyan Liu, Benlin Liu, Jie Zhou, and Jiwen Lu. Unleashing text-to-image diffusion models for visual perception. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5729–5739, 2023. [1](#), [2](#), [3](#), [5](#)
- [118] Yue Zhao, Ishan Misra, Philipp Krähenbühl, and Rohit Girdhar. Learning video representations from large language models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 6586–6597. IEEE, 2023. [2](#), [6](#), [7](#)
- [119] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luwei Zhou, Xiyang Dai, Lu Yuan, Yin Li, and Jianfeng Gao. Regionclip: Region-based language-image pretraining. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 16772–16782. IEEE, 2022. [2](#)
- [120] Zixin Zhu, Xuelu Feng, Dongdong Chen, Junsong Yuan, Chunming Qiao, and Gang Hua. Exploring pre-trained text-to-video diffusion models for referring video object segmentation. In Ales Leonardis, Elisa Ricci, Stefan Roth,

Olga Russakovsky, Torsten Sattler, and Gül Varol, editors, *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part XII*, volume 15070 of *Lecture Notes in Computer Science*, pages 452–469. Springer, 2024. [3](#)