

# Recurrent Video Masked Autoencoders

Daniel Zoran, Nikhil Parthasarathy, Yi Yang, Drew A Hudson, João Carreira, Andrew Zisserman  
Google DeepMind

<https://rvm-paper.github.io>

## Abstract

We present *Recurrent Video Masked-Autoencoders (RVM)*: a novel approach to video representation learning that leverages recurrent computation to model the temporal structure of video data. RVM couples an asymmetric masking objective with a transformer-based recurrent neural network to aggregate information over time, training solely on a simple pixel reconstruction loss. This design yields a highly efficient "generalist" encoder: RVM achieves competitive performance with state-of-the-art video models (e.g. VideoMAE, V-JEPA) on video-level tasks like action classification, and point and object tracking, while matching or exceeding the performance of image models (e.g. DINOv2) on tasks that require strong geometric and dense spatial features. Notably, RVM achieves strong performance in the small-model regime without requiring knowledge distillation, exhibiting up to  $30\times$  greater parameter efficiency than competing video masked autoencoders. Finally, we demonstrate that RVM's recurrent nature allows for stable feature propagation over long temporal horizons with linear computational cost, overcoming some of the limitations of standard spatio-temporal attention-based video models. Ablation studies further highlight the factors driving the model's success, with qualitative results showing that RVM learns rich representations of scene semantics, structure, and motion.

## 1. Introduction

It has long been hypothesized that biological systems learn visual representations by predicting the spatio-temporal evolution of the world [9, 53, 62, 64]. Indeed, even limited motion cues are sufficient to drive children's ability to robustly perceive and segment objects [65]. Recent advances in self-supervised learning (SSL) have revived the hope that artificial vision systems might also acquire such predictive world models purely from large-scale unlabeled video [8, 67].

Among the most successful approaches are masked autoencoders (MAEs), which learn by reconstructing randomly masked portions of images or videos [15, 28, 36, 67, 68], and Joint Embedding Predictive Architectures (JEPAs) [4, 8],

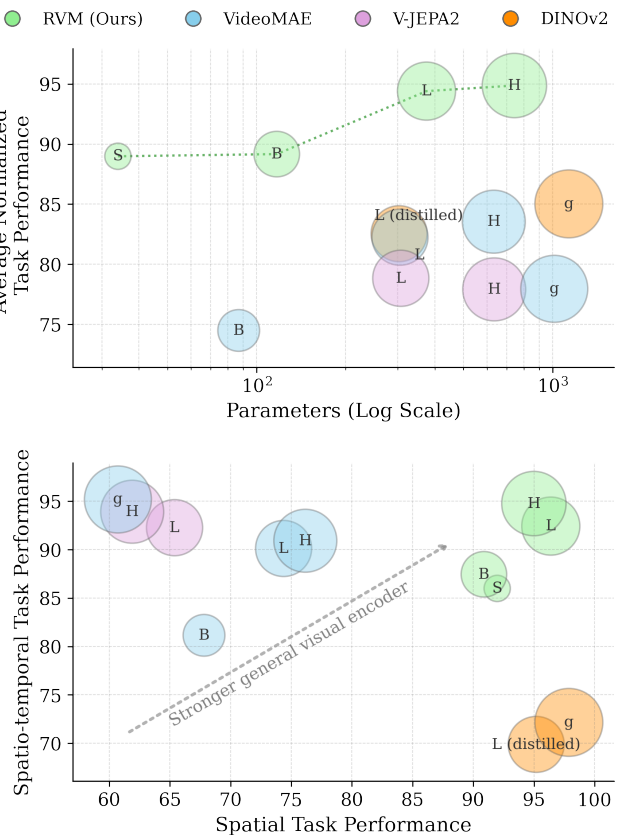


Figure 1. Normalized task performance is calculated for each task (relative to the best model) and averaged across tasks. **Top:** Across a wide range of visual tasks that require strong spatiotemporal features (video) and dense spatial features (image), RVM models set a pareto frontier that outperforms other strong video and image encoders. Spatio-temporal tasks cover: Something Somethingv2, Kinetics, Waymo object tracking, Perception Test TAP; Spatial tasks cover: ScanNet depth and nearest-neighbor correspondence tasks (Davis segmentation, JHMDB, Video Instance Parsing). **Bottom:** RVM models bridge the gap between strong spatial task models (e.g. DINOv2) and strong video task encoders, achieving the best of both worlds. Circle sizes are proportional to model size.

which predict future states in latent space while avoiding collapse via architectural or training heuristics [7, 26, 33]. Latent-space prediction has been argued to encourage learn-

ing task-relevant representations by discarding nuisance factors [32].

For video, both VideoMAE [67, 68] and V-JEPA [4, 8] rely on early-fusion spatio-temporal encoders (with spatio-temporal attention throughout the network) and random masking across entire clips. These designs treat time as uniform and symmetric, both in masking and in attention, neglecting the causal and directional nature of temporal dynamics. As a result, they are less amenable to online or streaming applications such as robotics. Moreover, their chunked offline architectures limit inference to short clips, preventing consistent representation learning over longer horizons.

Conversely, image-based models such as DINO [52, 61] excel at learning semantic representations and provide stable features when unrolled over multiple frames, but as they are still *image* models they are unable to capture motion information in their features. SiamMAE [34] partially addresses these limitations by training image encoders on natural video data, incorporating temporal asymmetry: it conditions on an unmasked “past” (*source*) frame to reconstruct a heavily masked “future” (*target*) frame via a cross-attention decoder. This asymmetric setup provides a strong inductive bias for learning correspondences. Nevertheless, SiamMAE still trains an *image* encoder, and thus it cannot ultimately capture spatio-temporal dependencies needed for true video-centric tasks.

In this work, we propose **Recurrent Video Masked Autoencoders (RVM)**, a family of general visual encoders that (in the spirit of SiamMAE), explicitly model the asymmetry of time through both masking and architecture. The RVM architecture processes videos sequentially by aggregating frame-level representations via a recurrent module. Training only with a pixel reconstruction loss on large natural video data, RVM learns strong representations that set a new pareto frontier in parameter efficiency (Figure 1, top) when evaluated across a wide-range of visual tasks. While SoTA models tend to specialize, RVM is *uniquely general, achieving strong average performance across spatial and video (spatio-temporal) tasks* (Figure 1, bottom). Furthermore, in the small model regime, RVM performs strikingly well *without requiring any form of model distillation*. Finally, owing to its recurrent design, RVM features *show emergent feature stability at long time horizons*, while being able to be unrolled over such sequences with linear compute and memory.

## 2. Related Work

**Self-Supervised Video Models.** Over the recent years, Self-Supervised Learning (SSL) [13, 17, 36, 39] has become a leading paradigm for deriving powerful representations from unlabeled visual data. In the case of videos, diverse learning methods have been proposed that harness the rich spatio-

temporal nature inherent to the domain. Earlier approaches focused on pretext tasks [22] that were designed to encourage learning temporal coherence and dynamics, by predicting frame order [43, 51, 73], motion statistics [2, 55, 70] or playback speed [11, 58, 78]. Other methods leveraged the multi-modal correspondence of video and audio, aiming to predict synchronization between the two [20]. More recently, contrastive learning approaches [35, 45, 54, 60] were developed to encourage consecutive frames’ embeddings to stay closer in the latent space, while pushing frames from different videos apart [74]. Meanwhile, masked modeling approaches [37, 67, 69] have proven both effective and robust in learning rich context-aware video representations, by reconstructing masked spatio-temporal patches from their surroundings.

**Masked Autoencoders.** Within the masked-modeling paradigm [6, 36, 72], a range of works introduce architectural and objective-level extensions to accommodate multiple views or frames [37, 67, 69]. A prominent direction involves the integration of Siamese networks [12, 18] and Masked Image Modeling [6, 36, 72], with examples such as SiamMAE [34], CropMAE [25] and CroCo [71] that respectively reconstruct a heavily-masked frame, crop, or view of a 3D-scene, by conditioning on a second unmasked one. Likewise, Counterfactual World Modeling (CWM) [10] explore temporally-factored masking, where a fully-visible frame informs the prediction of its heavily-occluded subsequent. Guided Future Prediction [16] innovates over standard masking, by replacing a few patches of an input frame with respective patches from a future one, to guide its reconstruction. Alternatively, MotionMAE [76] directly enriches the standard masking objective with the prediction of temporal difference between successive frames, so to encourage modeling of motion and dynamics.

**Recurrent Video Models.** Compared to the methods discussed above, our self-supervised video model stands out as it processes videos recurrently, so to explicitly model their temporal dynamics. It links to prior works about recurrent video architectures [23, 79]. One example is the Recurrent Vision Transformer (RViT) model [77], which forms an aggregated representation of a video by processing it iteratively with attention-based gating. Another notable instance is the Recurrent Convolutional Neural Network (RCNN) [47], which embeds recurrent connections directly into its convolutional layers, enabling the model to learn spatio-temporal features in a unified manner. By adopting a recurrent processing scheme, these models can assimilate the progression and directionality of time, and fully capture the long-range dependencies across frames and the temporal dynamics of videos. Recently, State Space Models (SSMs) such as VideoMamba [44] and VideoMambaPro [49] have also been explored for efficient video understanding. However, unlike our approach, these methods typically rely on

non-causal, bidirectional processing to achieve competitive performance and process videos as a flat sequence of tokens, discarding the spatial information.

### 3. Model

RVM is a recurrent model that encodes frames  $X_t$  sequentially to produce a set of features for each frame. Figure 2 provides an overview of the general architecture. Each input frame is patchified and encoded using a ViT. The resulting tokens are fed into a Recurrent Neural Network (RNN) core, which carries a state from the previous time-step and integrates tokens from the current frame to produce an updated state. This new state serves as the feature representation for the current time-step. By processing sequentially, RVM is able to ingest, discard, and refine information incrementally as it becomes available.

During training, the also model receives a *target* frame  $X_T$  sampled from a (potentially distant) future. We sample the target frame with a random time gap  $\Delta t$  of between 4 and 48 frames from the last source frame. Depending on the exact data used this corresponds to a time gap of 0.15 to 10 seconds. This target frame is then heavily masked and encoded using the same encoder as the input (*source*) frames. The training objective is to reconstruct the target frame using information from the source frames. This is achieved using a cross-attention based decoder [34] by minimizing the  $L_2$  loss between the reconstruction and the target.

#### 3.1. Modules

**Tokenization & Masking** Each frame  $X_t \in \mathbb{R}^{H \times W \times 3}$  (both source and target) is divided into non-overlapping patches of size  $P \times P$ . These patches are embedded via a learnable linear projection, resulting in a feature map of size  $h \times w \times D$  (where  $h = H/P$  and  $w = W/P$ ). These embeddings are then flattened into a sequence of  $N = hw$  tokens of size  $D$ . This tokenization process is applied independently to each frame, with weights shared across all source and target frames. Fourier positional encodings are subsequently added to the tokens.

For the *target* frame, tokens are randomly masked with a ratio  $m$  (defaulting to  $m = 0.95$ ). A learnable [CLS] token is concatenated to the token sequences of both source and target frames. This yields a sequence of  $K$  source token sets, denoted as  $e_1^s, \dots, e_K^s$  where  $e_t^s \in \mathbb{R}^{(N+1) \times D}$ , and a single set of unmasked target tokens  $e^t \in \mathbb{R}^{(M+1) \times D}$ , where  $M = \lfloor (1 - m)N \rfloor$ .

**Encoder** We employ a ViT encoder [24] to process the tokens of each frame independently. Specifically, we utilize standard ViT blocks with pre-normalization and without dropout. Following SiamMAE [34], the encoder weights are shared across all frames, both source and target. We denote the resulting encoded outputs as  $\hat{e}_t^S$  for the source frames

( $t = 1 \dots K$ ) and  $\hat{e}^T$  for the target frame.

**Recurrent Core** The encoded outputs from the source frames are fed into a *recurrent neural network* (RNN) core, formally defined as  $o_t, s_t = R(x_t, s_{t-1})$ . Here,  $x_t$  represents the input at the current time step,  $s_{t-1}$  denotes the state from the previous time step, while  $o_t$  and  $s_t$  represent the output and updated state for the current time step, respectively. We unroll the RNN sequentially over the source frames, producing a sequence of outputs  $o_t \in \mathbb{R}^{(N+1) \times D}$  and states  $s_t \in \mathbb{R}^{(N+1) \times D}$  for  $t = 1 \dots K$ . The initial state  $s_0$  is set to zero. This recurrent mechanism enables the model to aggregate information over time, constructing a temporally-aware representation. We utilize the outputs  $o_t$  as features for downstream tasks. The specific architectural details of the RNN are discussed in Section 3.2.

**Decoder** During training, our objective is to reconstruct the target frame  $X^T$  from its masked tokens  $\hat{e}^T$ , conditioned on the source frame features  $o_t$ . We employ a decoder with both cross- and self-attention mechanisms, similar to [34]. Both target and source frame features are first embedded via a linear layer. Following the standard MAE approach [36], we place the unmasked target tokens into their original grid positions, fill the masked locations with a learnable [MASK] token, and add Fourier positional embeddings. This sequence serves as the input to the decoder.

Each decoder block consists of three sequential components: (1) cross-attention, utilizing target tokens as queries and source tokens (concatenated along the token axis) as keys and values; (2) a feed-forward MLP; and (3) self-attention. All components utilize residual connections and pre-normalization (LayerNorm [5]). Finally, the decoder output is projected to the original patch dimension and reshaped to reconstruct the target frame.

**Loss** We use a simple  $L_2$  loss over the entire reconstructed and target image pixels, with no patch level normalization.

#### 3.2. The Rise of GRU

To effectively aggregate and integrate information over time, our model requires a module capable of maintaining a *state* across time steps. Ideally, this mechanism should retain critical information, discard irrelevant data, and assimilate new inputs as they arrive. Furthermore, we seek to leverage the efficacy of Transformers to facilitate spatiotemporal interactions between tokens. To address these needs, we propose a hybrid architecture combining a Transformer with a Gated Recurrent Unit (GRU).

This RNN core utilizes a combination of cross- and self-attention to integrate information. Specifically, the encoder outputs  $\hat{e}_t$  for the current time step serve as queries, while the keys and values are derived from the previous state  $s_{t-1}$ . To manage this information flow, we adopt the gating mechanism of the standard GRU [19]. The *reset gate*  $r_t$  modulates

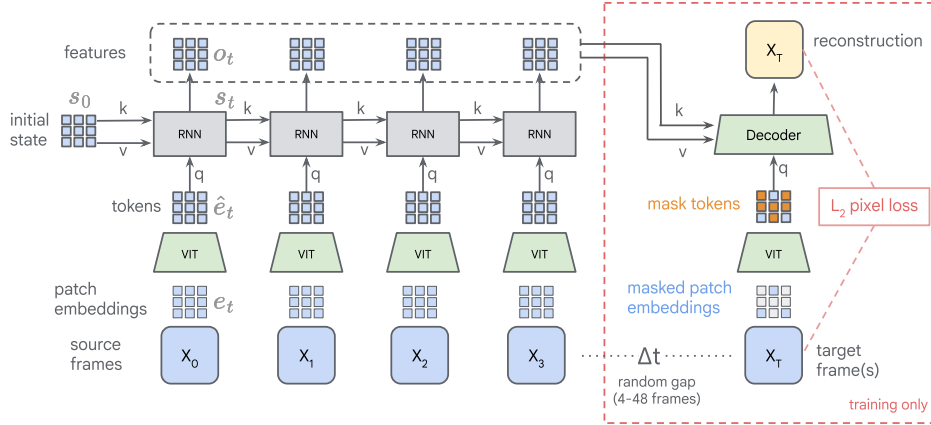


Figure 2. **RVM overview.** The model encodes *source* frames from an input video sequentially. Each frame is independently encoded using a vision transformer and the output tokens are aggregated using a transformer-based RNN to produce a sequence of features. See text for full details. During training, a *target* frame is sampled from a random time gap in the future, masked and encoded using the same ViT encoder. The model is trained to reconstruct the masked target frame using a cross attention decoder, minimizing the  $L_2$  loss between reconstruction and target.

the previous state before it is passed to the attention block, while the *update gate*  $u_t$  determines the balance between the previous state and the new attention output. The module is governed by the following equations:

$$\begin{aligned}
 u_t &= \sigma(W_e^u \hat{e}_t + W_s^u s_{t-1}) \\
 r_t &= \sigma(W_e^r \hat{e}_t + W_s^r s_{t-1}) \\
 \hat{h}_t &= \text{Tx}(q = \hat{e}_t, kv = r_t \odot s_{t-1}) \\
 s_t &= (1 - u_t) \odot s_{t-1} + u_t \odot \hat{h}_t \\
 o_t &= s_t
 \end{aligned}$$

Here,  $\sigma$  denotes the sigmoid function, and  $\text{Tx}$  represents a multi-layer Transformer block utilizing both cross- and self-attention. The weight matrices  $W$  are applied to the feature dimension and shared across all tokens. The state  $s_0$  is initialized to zero. Pseudo-code is provided in the Supplementary Material.

## 4. Experiments

### 4.1. Training

We train the model on a large dataset of a mixture of publicly available web videos. We base the mixture off of the mixture used in [3], containing sampled video clips from HowTo100M [50], Kinetics700 [14], SSV2 [29], YTBB [63], and YT8M [1]. The full dataset contains approximately 8.4M video clips (For more details see Supplementary Material). During training we randomly sample sub-clips, applying random flipping and random resized crop augmentation. The final frames are resized to  $256 \times 256$  resolution. We train several model sizes, scaling the encoder and RNN core accordingly, while following standard masked autoencoder (MAE) practice by keeping the decoder size fixed across

experiments. See Supplementary Material for architectural details (number of layers, hidden dimensions, etc.). All models are trained *from scratch* and no distillation procedure is used.

Unless otherwise stated, models are trained for 1M steps (250k steps for ablations) with a global batch size of 2048, corresponding to about 2B training examples in total. We highlight that the RVM architecture and objective seems to enable training for very long schedules with steady performance increase in all downstream tasks.

Each training example consists of a 64-frame video clip sampled randomly from the dataset mixture. From each clip, we sample 4 consecutive source frames that are processed by the recurrent encoder. We reconstruct 4 target frames, that are sampled independently and uniformly between 4 and 48 frames after the last source frame. Training is distributed across 256 TPU-v6 cores with per-core memory of 32GB. We use `bfloat16` precision for all forward and backward passes, while upcasting the loss and softmax computations to `float32`; model weights are stored in `float32` for stability. To fit large models within device memory, we employ FSDP-like parameter and optimizer sharding. Optimization uses AdamW [48] with a cosine decay learning rate schedule and warm-up phase. Full hyperparameter settings are summarized in the Supplementary Material.

### 4.2. Quantitative Results

**Baselines.** We compare RVM against a set of strong image and video model baselines:

- **Image models:** We compare to DINOv2 [52] as the main baseline for strong spatial task performance. At small model scales, we also include SiamMAE [34] as it is a strong efficient model for dense correspondence tasks such as video segmentation and human keypoint tracking.

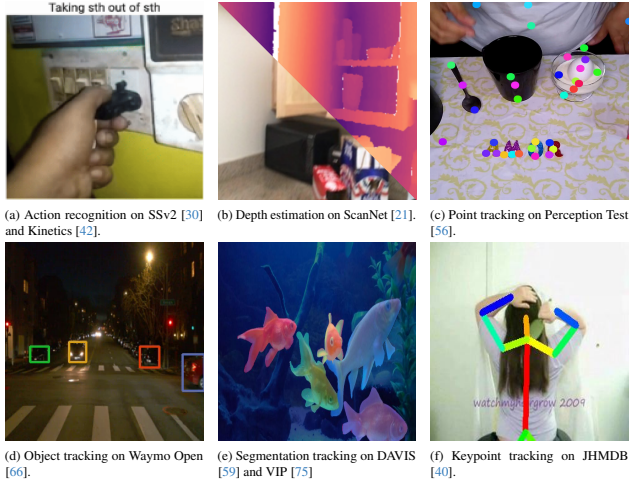


Figure 3. **Evaluation suite.** Individual frames and annotations from some of the evaluation tasks in this paper, covering semantic, geometry and motion perception.

- **Video models:** We evaluate variants of VideoMAE [67], V-JEPA2 [4], and 4DS [15]. These models are designed for large-scale video pretraining and largely represent the current frontier in video self supervised learning.

**Evaluation Suite.** We evaluate RVM on a diverse set of benchmarks (Figure 3 that we group into two primary categories:

- **Spatio-temporal tasks:** where performance generally improves with better modeling of both semantic/spatial and motion features. This set contains action recognition (SSv2 and Kinetics-700) [30, 42], Waymo open object tracking [66], and Perception Test point tracking [56].
- **Spatial tasks:** where strong, dense features extracted independently at the frame-level are generally sufficient. This set contains, Scannet depth estimation [21], DAVIS-2017 segmentation (spatial correspondence across frames) [59], JHMDB human keypoint tracking [40], and VIP human part tracking [75].

**Evaluation Protocol.** From a functional perspective, the tasks above can be also categorized into “readout tasks”, which train readout heads on top of frozen encoder representations, and “nearest-neighbor” (zero-shot) tasks that use pixel-level semantic label-propagation. For readout tasks, on top of the *frozen* pre-trained model we train an attentive readout head that follows the exact protocol described in recent literature [15]. Read-out heads for all models, including external ones, are trained using the same setup, with publicly available checkpoints. Nearest-neighbor tasks perform various forms of label propagation that follow the original protocols from each evaluation dataset. For more details on these benchmarks and the evaluation protocol, see the Supplementary Material.

#### 4.2.1. RVM learns strong generalist vision models

Results for large-scale models (L/H) in Table 1 reveal a clear dichotomy in baseline performance. As an image-encoder, DINOv2 performs well on spatial and semantic tasks but fails on intensive spatio-temporal tasks (e.g., 36.6/39.9 point tracking vs. video models achieving  $> 70$ ). In contrast, native video-encoders (VideoMAE, V-JEPA2) achieve high scores on spatio-temporal benchmarks like SSv2, but trade this off with very poor results on spatial correspondence (e.g.  $< 20$  mIoU on VIP vs. 40 mIoU for DINOv2).

RVM unifies these capabilities, achieving strong performance across both axes. To quantify this balance, we compute a “Normalized avg.” by averaging model scores normalized against the best performance per benchmark. Under this metric, RVM-L and RVM-H not only outperform their direct counterparts by more than 10% but also surpass giant-scale models (DINOv2-g, VideoMAEv2-g) by a similar margin. While the baselines show high variance with task-specific failures, RVM is the only architecture to avoid poor performance across the entire evaluation suite.

#### 4.2.2. RVM learns strong small models without distillation

Table 2 highlights the architectural efficiency of our approach. A key finding is that RVM continues to yield strong performance with very efficient models (ViT-S scale) *without requiring additional knowledge distillation*.

While there are not many model classes that even attempt to train at this scale, RVM-S outperforms competing baselines of similar size, achieving notable gains on SSv2 (+3.7%) and Kinetics-400 (+21.5% over 4DS-S) and outperforming the frame-based SiamMAE model on 3 out of 4 spatial tasks. RVM-S even outperforms a VideoMAE-B model (4x larger) on 6 out of 8 evaluations. This stands in contrast to prior state-of-the-art methods like DINOv2, that rely on distillation from larger teacher models to ensure strong performance in the small-compute regime. While distillation is clearly effective (DINOv2-S (distilled) outperforms RVM-S on 3 benchmarks), it remains the case that RVM-S has the highest average normalized performance across all benchmarks.

In fact, as seen in Figure 1 (Top), because RVM does not perform poorly on any one task, the average normalized accuracy of RVM-S *even when aggregating across all model scales* still outperforms both VideoMAEv2-g and DINOv2-g (*which are 30x larger!*).

#### 4.2.3. Long-term feature consistency

We compare the stability of features generated by different models over extended time horizons. To do this, we utilize the DAVIS segmentation task, specifically filtering the test dataset to include only videos exceeding 80 frames in length. We then evaluate and compare the tracking performance of the models at intervals of 16, 32, 48, 64, and 80 frames.

Table 1. **RVM learns a general visual representation that succeeds at both video-centric tasks that require spatio-temporal representations as well as tasks that primarily require strong dense geometric and spatial features.** While RVM does not outperform all baselines on every benchmark, we see that it provides the strongest general representation across all tasks, indicated by the Avg. normalized accuracy. We compute this by averaging the scores for each model across tasks after normalizing each column by the best model performance.

Model	Size(M)	Spatio-temporal tasks				Spatial Tasks				Average Normalized (†)
		SSv2 Acc. (†%)	Kinetics Acc. (†%)	Waymo mIoU (†)	PT AJ (†)	ScanNet AbsRel (↓)	DAVIS J&F (†%)	JHMDB PCK@0.1 (†%)	VIP mIoU (†)	
4DS-L	310	57.6	45.2	<b>75.9</b>	<b>81.5</b>	1.23	–	–	–	–
VideoMAE-L	305	62.7	52.5	74.9	78.3	1.10	54.3	45.1	18.9	82.2
V-JEPA2-L	307	<b>67.5</b>	<b>59.2</b>	73.1	73.7	1.08	47.5	33.3	17.2	78.8
<b>RVM-L</b>	375	66.7	57.3	73.2	77.3	<b>0.91</b>	<b>66.0</b>	<b>48.4</b>	<b>38.0</b>	<b>94.4</b>
DINOv2-L (distilled)	303	52.2	63.5	51.7	36.6	1.02	61.7	51.4	40.6	82.5
4DS-H	639	60.0	47.5	<b>76.1</b>	<b>81.8</b>	1.14	–	–	–	–
VideoMAE-H	633	64.2	54.5	74.6	77.1	1.04	54.4	43.9	20.8	83.5
V-JEPA2-H	635	<b>68.9</b>	59.1	73.6	77.0	1.05	43.3	29.7	16.0	77.9
<b>RVM-H</b>	743	68.7	<b>60.0</b>	74.2	78.3	<b>0.89</b>	<b>65.6</b>	<b>45.6</b>	<b>37.3</b>	<b>94.9</b>
DINOv2-g	1135	54.8	65.4	50.6	39.9	0.91	62.4	51.1	40.5	85.0
VideoMAEv2-g	1013	65.6	<b>69.7</b>	72.6	73.7	1.08	41.7	28.6	16.9	77.9

Table 2. **RVM enables strong small model performance without distillation.**

Model	Size(M)	Spatio-temporal Tasks				Spatial Tasks				Average Normalized (†)
		SSv2 Acc. (†%)	Kinetics Acc. (†%)	Waymo mIoU (†)	PT AJ (†)	ScanNet AbsRel (↓)	DAVIS J&F (†%)	JHMDB PCK@0.1 (†%)	VIP mIoU (†)	
SiamMAE-S	27	40.0	41.2	55.0	65.0	1.63	62.0	47.0	<b>37.3</b>	80.8
4DS-S	24	39.9	28.1	69.6	75.9	2.05	–	–	–	–
<b>RVM-S</b>	34	<b>59.7</b>	<b>49.6</b>	<b>70.5</b>	<b>76.5</b>	<b>0.97</b>	<b>62.9</b>	<b>47.5</b>	35.9	<b>96.1</b>
DINOv2-S (distilled)	21	48.3	57.1	51.0	33.4	1.17	62.4	48.4	39.4	84.4
VideoMAE-B	87	52.3	38.9	<b>73.1</b>	<b>79.2</b>	1.50	50.6	44.7	19.6	80.4




Figure 4. **RVM features are uniquely stable over long timescales.** We measure temporal stability of visual features by looking at label propagation (feature correspondence) on videos with increasing numbers of frames from the DAVIS 2017 benchmark. RVM performance decays substantially less for long sequences than other SoTA video and image models.

Figure 4 illustrates label-propagation performance as a function of frame count. Results are normalized to each model’s performance on 16 frames. As expected, all models perform worse as the time horizon increases. However, RVM demonstrates a significantly slower decline in performance, outperforming all other video models as well as strong image-based baselines like DINOv2. This indicates that the recurrent core successfully retains temporally useful information

to support long-range correspondence. This result is particularly notable given that RVM is trained with only a 4-frame horizon. In contrast, video models that process video in independent blocks, such as VideoMAE, degrade much faster as the number of frames increases, highlighting the critical importance of carrying state across long intervals. As an added benefit, we also note that as the number of frames processed grows, RVM demonstrates significant latency benefits compared with chunked video models. As seen in the Supplementary Material, the recurrent temporal aggregation shows linear latency scaling as opposed to the quadratic scaling of models that use full spatio-temporal self-attention.

### 4.3. Ablations

To better understand the contributions of different model components, we conduct an extensive ablation analysis. All ablation experiments are performed using the Small (S) version of the model, trained on 500M examples (see Table 3). We specifically investigate the importance of the time aggregation architecture, the number of source frames, and the scaling behavior with respect to training data size. Full details of ablations can be found in the Supplementary Material.

**Number of source frames** We train the model with 1, 2, and 4 source frames while keeping all other settings con-

num frames	SSv2( $\uparrow$ )	Kinetics( $\uparrow$ )	ScanNet( $\downarrow$ )
1	41.020	39.280	1.596
2	47.230	39.020	1.620
<b>4</b>	<b>52.34</b>	<b>39.72</b>	<b>1.50</b>

(a) **Number of source frames.** Having more source frames helps performance. Note that 1 source frame is an apples-to-apples comparison to SiamMAE [34] with exactly the same data, code and architecture.

aggregator	SSv2( $\uparrow$ )	Kinetics( $\uparrow$ )	ScanNet( $\downarrow$ )
SA	49.1	39.6	1.505
<b>RNN</b>	<b>52.34</b>	<b>39.72</b>	<b>1.50</b>

(b) **Encoder architecture.** The RNN performs better than full self-attention, while being more computational efficient (see text).

num steps	SSv2( $\uparrow$ )	Kinetics( $\uparrow$ )	ScanNet( $\downarrow$ )
250M	46.38	33.84	1.75
<b>500M</b>	<b>52.34</b>	<b>39.72</b>	<b>1.50</b>
1B	55.09	44.33	1.32
<b>2B</b>	<b>57.20</b>	<b>47.70</b>	<b>1.200</b>

(c) **Number of training examples and steps.** Even with a small model, our approach benefits from more data and training.

Table 3. **RVM ablation experiments.** We ablate some of the components of the model. We show that (a) using more source frames significantly improves results. (b) Using an RNN to aggregate information across time instead of full self-attention is beneficial, especially with tasks that require motion understanding like SSv2 and (c) that the model benefits from more data and training. Default settings for the ablation are marked in blue .

stant. The single-frame case serves as an "apples-to-apples" comparison with SiamMAE, controlling for the additional RNN layers and other training differences. With two source frames, the model can capture constant velocity motion, though higher-order dynamics (like acceleration) remain out of reach. We observe a consistent improvement in performance across all tasks as the number of source frames increases from 1 to 2, and further to 4.

**Encoder architecture** We compare our proposed RNN temporal aggregator against a classic self-attention Transformer. To ensure a fair comparison, we use a patch size of  $1 \times 16 \times 16$  and match the number of layers and parameters to the RNN core. It is worth noting that the full self-attention mechanism incurs a significantly higher computational cost (FLOPs) compared to the RNN. As shown in the results, the RNN approach is not only more efficient but also performs favorably compared to the self-attention alternative.

**Number of training examples** Finally, we evaluate how the model’s performance scales with the amount of training data. We train four different models on 250M, 500M, 1B, and 2B data samples, respectively. We scale the learning schedules according to the data volume while keeping all other hyperparameters constant. We observe that despite using a relatively small model (34M parameters), our approach continues to benefit from additional data without exhibiting signs of overfitting.

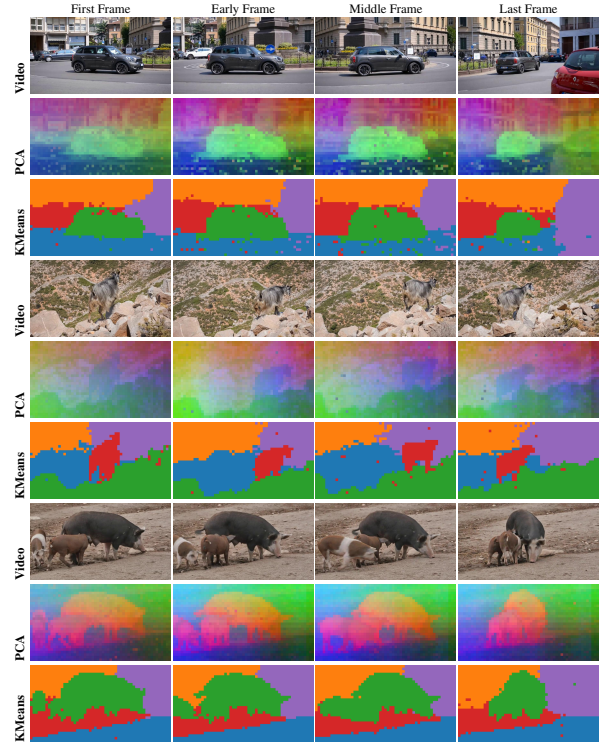


Figure 5. **PCA and K-means of RVM features unrolled on unseen videos.** Despite being trained on only 4 frames the model generalizes to long sequences and unrolls stably over long time horizons. As can be seen, the model learns to extract meaningful features from videos.

#### 4.4. Qualitative Evaluation

We begin by qualitatively evaluating the features learned by our model. Using a trained Large (L) model, we unroll it over various test sequences and aggregate the features across time. First, we observe that although the model was trained on only 4 source frames, it generalizes well to much longer sequences without stability issues.

For our first set of test sequences (Figure 5), we visualize the features using Principal Component Analysis (PCA) and K-means clustering. For PCA, we concatenate features from all frames and spatial locations, compute the principal components, and map the top three components to the RGB channels of an image. The results show that the model captures meaningful video structures. Similarly, for K-means (with  $K = 5$ ), we cluster the concatenated tokens and visualize the resulting segmentation maps by color-coding each cluster. This demonstrates that the model learns to cluster semantically consistent regions in a self-supervised manner. Figure 7 show K-means clustering for other models. RVM produces comparatively clean and stable features.

To show that the model learns meaningful *motion* representation, we test the model on a classic stimulus of a solid white noise square moving on top a static white noise

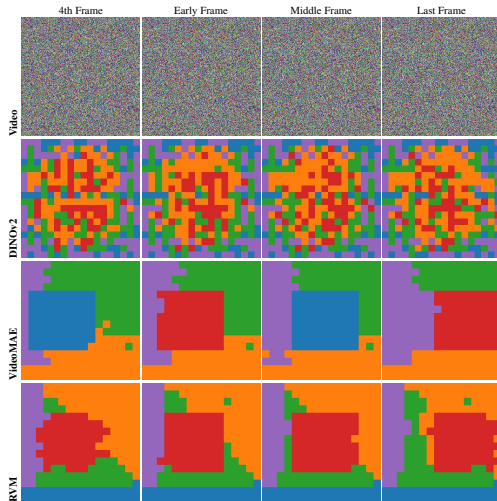


Figure 6. **Detecting a white noise square moving on a white noise background.** From top to bottom: input sequence, RVM K-means visualization, an example feature map. Note that each frame independently in the input sequence is a white noise image, and thus image models like DINO or Siam-MAE cannot extract any useful information from these. RVM however can integrate temporal information and "see" the moving square. It is highly recommended to watch the video which can be found in the supplementary material. All models use the same ViT-L-16 backbone.

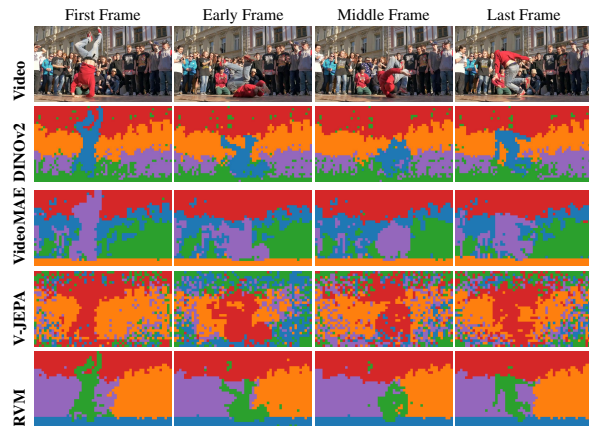


Figure 7. **KMeans visualization on DAVIS video for various ViT-L/16 models.** Unlike RVM, other models produce noisy feature maps lacking structure and consistency.

background. This stimulus is interesting because each frame independently is just a white noise image and contains no meaningful structure (as can be seen in Figure 6 top, compare to the video version in the supplementary material). Hence, *image* encoders like DINO or SiamMAE can not extract any useful information from these (Figure 6 second row). RVM however is able to "see" the resulting structure. Other video models can also capture the underlying structure, but due to their limited temporal support window, they are unable to provide stable features across the whole sequence (note the cluster reassignments in Figure 6 for VideoMAE).

## 5. Limitations

While RVM sets a new frontier for parameter efficiency and enables linear scaling for long-context inference, this recurrent design incurs specific trade-offs. First, unlike spatio-temporal models such as VideoMAE that patchify across time to reduce token counts, RVM processes frames sequentially. This makes RVM computationally heavier for very short sequences where the benefits of recurrence are less pronounced. Second, training requires back-propagation through time with a ViT encoder at every step, which is memory-intensive. Finally, both a benefit and limitation is that we have yet to find the data saturation point for these models. In this work, we train with 2B clips but find that performance continues to improve with more data. It would be beneficial to establish more formal scaling laws for RVM so that we can more efficiently allocate compute.

## 6. Conclusion

We present Recurrent Video Masked-Autoencoders (RVM), a novel framework that leverages recurrent computation as a way to integrate temporal information in self-supervised video representation learning. By coupling an asymmetric masking autonecoder style training objective with a transformer-based recurrent core, RVM effectively aggregates information over time to learn "generalist" visual representations. Our results demonstrate that RVM presents a unique advance in the landscape of current vision models: it matches or exceeds the spatio-temporal capabilities of video-centric models (e.g., VideoMAE, V-JEPA) while retaining the dense, spatial and geometric understanding properties of strong frame-centric models (e.g., DINOv2). Furthermore, RVM introduces a way to train strong small models, trained without the need for knowledge-distillation, that exhibit up to  $30\times$  greater parameter efficiency for the same averaged performance. Finally, we find that this recurrent architecture exhibits superior feature stability over long temporal horizons compared to state-of-the-art "video model" baselines. In sum, our work suggests that bringing back recurrent video processing with a simple pixel-level training objective may be sufficient for learning strong visual models from natural video data without the need for extra tricks like strong augmentation, EMA networks, regularizers etc. Future work will explore further scaling our method and evaluating it in the context of multi-modal and world modeling tasks like robotic control.

## Acknowledgment

We thank Goker Erdogan, Viorica Pătrăucean, Aravindh Mahendran, Miki Rubinstein, Dilara Gokay, Junlin Zhang and Joseph Heyward for helpful discussions and support.

## References

- [1] Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudeendra Vijayanarasimhan. Youtube-8m: A large-scale video classification benchmark. *arXiv preprint arXiv:1609.08675*, 2016. 4, 13
- [2] Pulkit Agrawal, Joao Carreira, and Jitendra Malik. Learning to see by moving. In *Proceedings of the IEEE international conference on computer vision*, pages 37–45, 2015. 2
- [3] Mido Assran, Adrien Bardes, David Fan, Quentin Garrido, Russell Howes, Matthew Muckley, Ammar Rizvi, Claire Roberts, Koustuv Sinha, Artem Zhohus, et al. V-jepa 2: Self-supervised video models enable understanding, prediction and planning. *arXiv preprint arXiv:2506.09985*, 2025. 4, 13
- [4] Mido Assran, Adrien Bardes, David Fan, Quentin Garrido, Russell Howes, Matthew Muckley, Ammar Rizvi, Claire Roberts, Koustuv Sinha, Artem Zhohus, et al. V-jepa 2: Self-supervised video models enable understanding, prediction and planning. *arXiv preprint arXiv:2506.09985*, 2025. 1, 2, 5
- [5] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 3
- [6] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. 2
- [7] Adrien Bardes, Jean Ponce, and Yann LeCun. Mc-jepa: A joint-embedding predictive architecture for self-supervised learning of motion and content features. *arXiv preprint arXiv:2307.12698*, 2023. 1
- [8] 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. *Transactions on Machine Learning Research*, 2024. Featured Certification. 1, 2
- [9] Horace B Barlow et al. Possible principles underlying the transformation of sensory messages. *Sensory communication*, 1(01):217–233, 1961. 1
- [10] Daniel M Bear, Kevin Feigelis, Honglin Chen, Wanhee Lee, Rahul Venkatesh, Klemen Kotar, Alex Durango, and Daniel LK Yamins. Unifying (machine) vision via counterfactual world modeling. *arXiv preprint arXiv:2306.01828*, 2023. 2
- [11] Sagie Benaim, Ariel Ephrat, Oran Lang, Inbar Mosseri, William T Freeman, Michael Rubinstein, Michal Irani, and Tali Dekel. Speednet: Learning the speediness in videos. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9922–9931, 2020. 2
- [12] Luca Bertinetto, Jack Valmadre, Joao F Henriques, Andrea Vedaldi, and Philip HS Torr. Fully-convolutional siamese networks for object tracking. In *European conference on computer vision*, pages 850–865. Springer, 2016. 2
- [13] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. 2, 15
- [14] Joao Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *arXiv preprint arXiv:1907.06987*, 2019. 4, 13
- [15] João Carreira, Dilara Gokay, Michael King, Chuhan Zhang, Ignacio Rocco, Aravindh Mahendran, Thomas Albert Keck, Joseph Heyward, Skanda Koppula, Etienne Pot, Goker Erdogan, Yana Hasson, Yi Yang, Klaus Greff, Guillaume Le Moing, Sjoerd van Steenkiste, Daniel Zoran, Drew A Hudson, Pedro Vélez, Luisa Polanía, Luke Friedman, Chris Duvarney, Ross Goroshin, Kelsey Allen, Jacob Walker, Rishabh Kabra, Eric Aboussouan, Jennifer Sun, Thomas Kipf, Carl Doersch, Viorica Pătrăucean, Dima Damen, Pauline Luc, Mehdi S M Sajjadi, and Andrew Zisserman. Scaling 4D representations. *arXiv [cs.CV]*, 2024. 1, 5, 14, 15, 16
- [16] Joao Carreira, Michael King, Viorica Patraucean, Dilara Gokay, Catalin Ionescu, Yi Yang, Daniel Zoran, Joseph Heyward, Carl Doersch, Yusuf Ayar, et al. Learning from one continuous video stream. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 28751–28761, 2024. 2
- [17] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PmlR, 2020. 2
- [18] Davide Chicco. Siamese neural networks: An overview. *Artificial neural networks*, pages 73–94, 2021. 2
- [19] Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014. 3
- [20] Joon Son Chung and Andrew Zisserman. Out of time: automated lip sync in the wild. In *Asian conference on computer vision*, pages 251–263. Springer, 2016. 2
- [21] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 2017. 5, 14, 15
- [22] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, pages 1422–1430, 2015. 2
- [23] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2625–2634, 2015. 2
- [24] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint*, 2020. 3, 13, 16
- [25] Alexandre Eymaël, Renaud Vandeghen, Anthony Cioppa, Silvio Giancola, Bernard Ghanem, and Marc Van Droogenbroeck. Efficient image pre-training with siamese cropped masked autoencoders. In *European Conference on Computer Vision*, pages 348–366. Springer, 2024. 2
- [26] Zhengcong Fei, Mingyuan Fan, and Junshi Huang. A-jepa: Joint-embedding predictive architecture can listen. *arXiv preprint arXiv:2311.15830*, 2023. 1

- [27] Christoph Feichtenhofer, Haoqi Fan, Bo Xiong, Ross Girshick, and Kaiming He. A large-scale study on unsupervised spatiotemporal representation learning. In *CVPR*, 2021. 15
- [28] Christoph Feichtenhofer, Yanghao Li, Kaiming He, et al. Masked autoencoders as spatiotemporal learners. *Advances in neural information processing systems*, 35:35946–35958, 2022. 1
- [29] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The "something something" video database for learning and evaluating visual common sense. In *ICCV*, 2017. 4, 13, 16
- [30] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The "something something" video database for learning and evaluating visual common sense. In *Proceedings of the IEEE international conference on computer vision*, pages 5842–5850, 2017. 5, 14
- [31] Klaus Greff, Francois Belletti, Lucas Beyer, Carl Doersch, Yilun Du, Daniel Duckworth, David J. Fleet, Dan Gnanaprasgam, Florian Golemo, Charles Herrmann, Thomas Kipf, Abhijit Kundu, Dmitry Lagun, Issam Laradji, Hsueh-Ti (Derek) Liu, Henning Meyer, Yishu Miao, Derek Nowrouzezahrai, Cengiz Oztireli, Etienne Pot, Noha Radwan, Daniel Rebain, Sara Sabour, Mehdi S. M. Sajjadi, Matan Sela, Vincent Sitzmann, Austin Stone, Deqing Sun, Suhani Vora, Ziyu Wang, Tianhao Wu, Kwang Moo Yi, Fangcheng Zhong, and Andrea Tagliasacchi. Kubric: A scalable dataset generator. In *CVPR*, 2022. 15
- [32] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020. 2
- [33] Pierre Guetschel, Thomas Moreau, and Michael Tangemann. S-jepa: Towards seamless cross-dataset transfer through dynamic spatial attention. *arXiv preprint arXiv:2403.11772*, 2024. 1
- [34] Agrim Gupta, Jiajun Wu, Jia Deng, and Fei-Fei Li. Siamese masked autoencoders. *Advances in Neural Information Processing Systems*, 36:40676–40693, 2023. 2, 3, 4, 7
- [35] Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning. *Advances in neural information processing systems*, 33:5679–5690, 2020. 2
- [36] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. 1, 2, 3, 15
- [37] Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. *arXiv preprint arXiv:2205.15868*, 2022. 2
- [38] Allan Jabri, Andrew Owens, and Alexei Efros. Space-time correspondence as a contrastive random walk. *Advances in neural information processing systems*, 33:19545–19560, 2020. 15
- [39] Ashish Jaiswal, Ashwin Ramesh Babu, Mohammad Zaki Zadeh, Debapriya Banerjee, and Fillia Makedon. A survey on contrastive self-supervised learning. *Technologies*, 9(1):2, 2020. 2
- [40] Hueihan Jhuang, Juergen Gall, Silvia Zuffi, Cordelia Schmid, and Michael J Black. Towards understanding action recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 3192–3199, 2013. 5, 14, 15
- [41] 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, 2017. 16
- [42] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 5, 14
- [43] Hsin-Ying Lee, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Unsupervised representation learning by sorting sequences. In *Proceedings of the IEEE international conference on computer vision*, pages 667–676, 2017. 2
- [44] Kunchang Li, Xinhao Li, Yi Wang, Yanan He, Yali Wang, Limin Wang, and Yu Qiao. Videomamba: State space model for efficient video understanding. In *European conference on computer vision*, pages 237–255. Springer, 2024. 2
- [45] Muheng Li, Lei Chen, Yueqi Duan, Zhilan Hu, Jianjiang Feng, Jie Zhou, and Jiwen Lu. Bridge-prompt: Towards ordinal action understanding in instructional videos. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 19880–19889, 2022. 2
- [46] Xueting Li, Sifei Liu, Shalini De Mello, Xiaolong Wang, Jan Kautz, and Ming-Hsuan Yang. Joint-task self-supervised learning for temporal correspondence. *Advances in Neural Information Processing Systems*, 32, 2019. 15
- [47] Ming Liang and Xiaolin Hu. Recurrent convolutional neural network for object recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3367–3375, 2015. 2
- [48] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 4
- [49] Hui Lu, Albert Ali Salah, and Ronald Poppe. Videomamba: A leap forward for mamba in video understanding. *arXiv e-prints*, pages arXiv–2406, 2024. 2
- [50] 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 *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2630–2640, 2019. 4, 13
- [51] Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *European conference on computer vision*, pages 527–544. Springer, 2016. 2
- [52] 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, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024. 2, 4, 16
- [53] Stephanie E Palmer, Olivier Marre, Michael J Berry, and William Bialek. Predictive information in a sensory population. *Proceedings of the National Academy of Sciences*, 112(22):6908–6913, 2015. 1
- [54] Tian Pan, Yibing Song, Tianyu Yang, Wenhao Jiang, and Wei Liu. Videomoco: Contrastive video representation learning with temporally adversarial examples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11205–11214, 2021. 2
- [55] Deepak Pathak, Ross Girshick, Piotr Dollár, Trevor Darrell, and Bharath Hariharan. Learning features by watching objects move. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2701–2710, 2017. 2
- [56] Viorica Patraucean, Lucas Smaira, Ankush Gupta, Adria Recasens Contente, Larisa Markeeva, Dylan Sunil Banarse, Skanda Koppula, Joseph Heyward, Mateusz Malinowski, Yi Yang, Carl Doersch, Tatiana Matejovicova, Yury Sulsky, Antoine Miech, Alexandre Fréchet, Hanna Klimczak, Raphael Koster, Junlin Zhang, Stephanie Winkler, Yusuf Aytar, Simon Osindero, Dima Damen, Andrew Zisserman, and Joao Carreira. Perception test: A diagnostic benchmark for multimodal video models. In *NeurIPS*, 2023. 5, 14, 15
- [57] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 724–732, 2016. 15
- [58] Lyndsey C Pickup, Zheng Pan, Donglai Wei, YiChang Shih, Changshui Zhang, Andrew Zisserman, Bernhard Scholkopf, and William T Freeman. Seeing the arrow of time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2035–2042, 2014. 2
- [59] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*, 2017. 5, 14, 15
- [60] Rui Qian, Tianjian Meng, Boqing Gong, Ming-Hsuan Yang, Huisheng Wang, Serge Belongie, and Yin Cui. Spatiotemporal contrastive video representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6964–6974, 2021. 2
- [61] 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, 2021. 2
- [62] Rajesh PN Rao and Dana H Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1):79, 1999. 1
- [63] Esteban Real, Jonathon Shlens, Stefano Mazzocchi, Xin Pan, and Vincent Vanhoucke. Youtube-boundingboxes: A large high-precision human-annotated data set for object detection in video. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5296–5305, 2017. 4, 13
- [64] Yosef Singer, Yayoi Teramoto, Ben DB Willmore, Jan WH Schnupp, Andrew J King, and Nicol S Harper. Sensory cortex is optimized for prediction of future input. *elife*, 7:e31557, 2018. 1
- [65] Elizabeth S Spelke. Principles of object perception. *Cognitive science*, 14(1):29–56, 1990. 1
- [66] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurélien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in perception for autonomous driving: Waymo open dataset. In *CVPR*, 2020. 5, 14, 15
- [67] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *NeurIPS*, 2022. 1, 2, 5
- [68] 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 *CVPR*, 2023. 1, 2, 16
- [69] Rui Wang, Dongdong Chen, Zuxuan Wu, Yinpeng Chen, Xiyang Dai, Mengchen Liu, Yu-Gang Jiang, Luwei Zhou, and Lu Yuan. Bevt: Bert pretraining of video transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14733–14743, 2022. 2
- [70] Xiaolong Wang and Abhinav Gupta. Unsupervised learning of visual representations using videos. In *Proceedings of the IEEE international conference on computer vision*, pages 2794–2802, 2015. 2
- [71] Philippe Weinzaepfel, Vincent Leroy, Thomas Lucas, Romain Bréquier, Johann Cabon, Vaibhav Arora, Leonid Antsfeld, Boris Chidlovskii, Gabriela Csurka, and Jérôme Revaud. Croco: Self-supervised pre-training for 3d vision tasks by cross-view completion. *Advances in Neural Information Processing Systems*, 35:3502–3516, 2022. 2
- [72] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9653–9663, 2022. 2
- [73] Dejing Xu, Jun Xiao, Zhou Zhao, Jian Shao, Di Xie, and Yueting Zhuang. Self-supervised spatiotemporal learning via video clip order prediction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10334–10343, 2019. 2
- [74] Jiarui Xu and Xiaolong Wang. Rethinking self-supervised correspondence learning: A video frame-level similarity perspective. In *Proceedings of the IEEE/CVF International*

*Conference on Computer Vision*, pages 10075–10085, 2021. [2](#)

- [75] Ning Xu, Linjie Yang, Yuchen Fan, Dingcheng Yue, Yuchen Liang, Jianchao Yang, and Thomas Huang. Youtube-vos: A large-scale video object segmentation benchmark. *arXiv preprint arXiv:1809.03327*, 2018. [5](#), [15](#)
- [76] Haosen Yang, Deng Huang, Bin Wen, Jiannan Wu, Hongxun Yao, Yi Jiang, Xiatian Zhu, and Zehuan Yuan. Motionmae: Self-supervised video representation learning with motion-aware masked auto encoders. *BMVC Proceedings*, 2024. [2](#)
- [77] Jiewen Yang, Xingbo Dong, Liujuan Liu, Chao Zhang, Jiajun Shen, and Dahai Yu. Recurring the transformer for video action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14063–14073, 2022. [2](#)
- [78] Yuan Yao, Chang Liu, Dezhao Luo, Yu Zhou, and Qixiang Ye. Video playback rate perception for self-supervised spatio-temporal representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6548–6557, 2020. [2](#)
- [79] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4694–4702, 2015. [2](#)
- [80] Qixian Zhou, Xiaodan Liang, Ke Gong, and Liang Lin. Adaptive temporal encoding network for video instance-level human parsing. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 1527–1535, 2018. [14](#), [15](#)