

DropMAE: Masked Autoencoders with Spatial-Attention Dropout for Tracking Tasks

Qiangqiang Wu¹ Tianyu Yang^{2*} Ziquan Liu¹ Baoyuan Wu⁴ Ying Shan³ Antoni B. Chan¹

¹Department of Computer Science, City University of Hong Kong

²International Digital Economy Academy ³Tencent AI Lab

⁴School of Data Science, The Chinese University of Hong Kong, Shenzhen

{qiangqw2-c, ziquanliu2-c}@my.cityu.edu.hk, tianyu-yang@outlook.com

wubaoyuan@cuhk.edu.cn, yingsshan@tencent.com, abchan@cityu.edu.hk

Abstract

In this paper, we study masked autoencoder (MAE) pre-training on videos for matching-based downstream tasks, including visual object tracking (VOT) and video object segmentation (VOS). A simple extension of MAE is to randomly mask out frame patches in videos and reconstruct the frame pixels. However, we find that this simple baseline heavily relies on spatial cues while ignoring temporal relations for frame reconstruction, thus leading to sub-optimal temporal matching representations for VOT and VOS. To alleviate this problem, we propose DropMAE, which adaptively performs spatial-attention dropout in the frame reconstruction to facilitate temporal correspondence learning in videos. We show that our DropMAE is a strong and efficient temporal matching learner, which achieves better fine-tuning results on matching-based tasks than the ImageNet-based MAE with $2\times$ faster pre-training speed. Moreover, we also find that motion diversity in pre-training videos is more important than scene diversity for improving the performance on VOT and VOS. Our pre-trained DropMAE model can be directly loaded in existing ViT-based trackers for fine-tuning without further modifications. Notably, DropMAE sets new state-of-the-art performance on 8 out of 9 highly competitive video tracking and segmentation datasets. Our code and pre-trained models are available at <https://github.com/jimmy-dq/DropMAE.git>.

1. Introduction

Recently, transformers have achieved enormous success in many research areas, such as natural language processing (NLP) [6, 23], computer vision [100] and audio generation [45, 76]. In NLP, masked autoencoding is commonly used to train large-scale generalizable NLP transformers contain-




Figure 1. State-of-the-art comparison on GOT-10k [40] following the one-shot protocol. Our DropTrack with the proposed DropMAE pre-training achieves state-of-the-art performance without using complicated pipelines such as online updating.

ing billions of parameters. Inspired by the great success of self-supervised learning in NLP, recent advances [37, 92] in computer vision suggest that training large-scale vision transformers may undertake a similar trajectory with NLP. The seminal work MAE [37] reconstructs the input image from a small portion of patches. The learned representation in this masked autoencoder has been demonstrated to be effective in many computer vision tasks, such as image classification, object detection and semantic segmentation.

In video object tracking (VOT), recently two works, SimTrack [10] and OSTRack [100], explore using an MAE pre-trained ViT model as the tracking backbone. Notably, these two trackers achieve state-of-the-art performance on existing tracking benchmarks without using complicated tracking pipelines. The key to their success is the robust pre-training weights learned by MAE on ImageNet [70]. In addition, [10, 100] also demonstrate that, for VOT, MAE *unsupervised* pre-training on ImageNet is more effective than *supervised* pre-training using class labels – this is mainly because MAE pre-training learns more fine-grained local

*Corresponding Author




Figure 2. Visualization of the attention maps of the TwinMAE baseline and our DropMAE in the reconstruction of a random masked patch, which is denoted as a red bounding box in the left input frame. TwinMAE leverages the spatial cues (within the same frame) more than temporal cues (between frames) for reconstruction. Our proposed DropMAE improves the baseline by effectively alleviating co-adaptation between spatial cues in the reconstruction, focusing more on temporal cues, thus achieving better learning of temporal correspondences for VOT and VOS.

structures that are useful for accurate target localization required for VOT, whereas supervised training learns high-level class-related features that are invariant over appearance changes. Despite the promising performance achieved by [10, 100], the MAE pre-training on ImageNet could still be sub-optimal for the tracking task due to the natural gap between images and videos, i.e., no prior temporal correspondence information can be learned in static images. However, previous tracking methods [3, 44, 86] have shown that temporal correspondence learning is the key in developing a robust and discriminative tracker. Thus there is an opportunity to further develop the MAE framework specifically for matching-base video tasks, such as VOT and VOS.

One simple way to extend MAE to videos is to randomly mask out frame patches in a video clip (i.e., video frame pairs) and then reconstruct the video clip. We denote this simple baseline as *twin MAE* (TwinMAE). Given a masked patch query, as illustrated in Figs. 2 & 4, we find that TwinMAE heavily relies on the spatially neighbouring patches *within the same frame* to reconstruct the masked patch, which implies a heavy co-adaptation of spatial cues (within-frame tokens) for reconstruction and may cause learning of sub-optimal temporal representations for matching-based downstream tasks like video object tracking and segmentation.

To address this issue with the TwinMAE baseline, we propose DropMAE specifically designed for pre-training a masked autoencoder for matching-based video downstream tasks (e.g., VOT and VOS). Our DropMAE adaptively performs spatial-attention dropout to break up co-adaptation between spatial cues (within-frame tokens) during the frame reconstruction, thus encouraging more temporal interactions and facilitating temporal correspondence learning in

the pre-training stage. Interestingly, we obtain several important findings with DropMAE: 1) DropMAE is an effective and efficient temporal correspondence learner, which achieves better fine-tuning results on matching-based tasks than the ImageNet-based MAE with $2\times$ faster pre-training speed. 2) Motion diversity in pre-training videos is more important than scene diversity for improving the performance on VOT and VOS.

We conduct downstream task evaluation on 9 competitive VOT and VOS benchmarks, achieving state-of-the-art performance on these benchmarks. In particular, our trackers with DropMAE pre-training obtain 75.9% AO on GOT-10k, 52.7% AUC on LaSOT_{ext}, 56.9% AUC on TNL2K and 92.1%/83.0% $\mathcal{J}\&\mathcal{F}$ scores on DAVIS-16/17, w/o using complicated online updating or memory mechanisms.

In summary, the main contributions of our work are:

- To the best of our knowledge, we are the first to investigate masked autoencoder video pre-training for temporal matching-based downstream tasks. Specifically, we explore various video data sources for pre-training and build a TwinMAE baseline to study its effectiveness on temporal matching tasks. Since none exists, we further build a ViT-based VOS baseline for fine-tuning.
- We propose DropMAE, which adaptively performs spatial-attention dropout in the frame reconstruction to facilitate effective temporal correspondence learning in videos.
- Our trackers with DropMAE pre-training sets new state-of-the-art performance on 8 out of 9 highly competitive video tracking and segmentation benchmarks without complicated tracking pipelines.

2. Related Work

Visual object tracking and segmentation. Given an annotated bounding box in the first frame of a test video, visual object tracking (VOT) aims to accurately predict the target’s bounding boxes in the following frames. Similarly, for visual object segmentation (VOS), given an annotated binary mask in the first frame, VOS aims to predict dense target masks in the remaining frames. In the early development of VOT, correlation filter-based approaches [20, 22, 32, 38, 48, 49, 53, 55, 85, 87] are dominant trackers due to their favorable ability in modeling target appearance variation. With the development of deep learning, deep Siamese networks [73, 81] are introduced to VOT. The representative work SiamFC [3] takes template and search images as input for target localization. Based on SiamFC, many improvements have been made, e.g., scale regression [43, 44, 97], online template updating [35, 96, 102], multi-level feature fusion [30], loss design [24] and backbone design [10, 43, 100, 106]. For VOS, matching-based approaches, e.g., STM [62], AOT [99] and STCN [16], achieve promising results on existing VOS benchmarks.

Recent improvements [15, 52] on online memory design further improve the previous SOTA results in VOS. Recent studies including SimTrack [10] and OSTRack [100] show that the ViT backbone [25] with MAE pre-training on ImageNet is effective for object tracking.

Despite the great success of ViT with MAE pre-training on tracking, this static ImageNet-based pre-training still lacks temporal correspondence learning. Moreover, the developments of VOT and VOS show that learning strong temporal matching ability is essential for video tracking tasks. To the best of our best knowledge, we are the first to investigate masked autoencoder self-supervised video pre-training for tracking tasks. Our proposed DropMAE can learn reliable temporal correspondences, which we demonstrate are effective for downstream tracking tasks.

Self-supervised Learning. Self-supervised learning has received significant interest in the past few decades. There are many manually-designed pretext tasks for pre-training, such as image colorization [103], jigsaw puzzle solving [61], future frame prediction [72, 78] and rotation prediction [34]. Contrastive learning approaches [7, 12, 26, 36, 89, 91] are the mainstream self-supervised methods in recent years. However, these methods are sensitive to the type and strength of applied data augmentation, which makes them hard to train. Inspired by masked language modeling [6, 23], masked image modeling (MIM) approaches are proposed for learning unsupervised image [37, 92] or video representations [31, 74], which have been shown to be effective for many downstream tasks including image classification, object detection and video action recognition. However, there is no specifically designed pre-training approaches for temporal matching-based task, such as VOT and VOS. In this work, we propose DropMAE to explore this direction.

3. Method

We propose a self-supervised video pre-training method to learn robust representations for temporal matching-based downstream tasks, including VOT and VOS tasks. We firstly introduce a simple extension of MAE to temporal matching representation learning from video, (TwinMAE). We then illustrate the limitations of the TwinMAE baseline and propose a spatial-attention dropout strategy to facilitate temporal correspondence learning, denoted as DropMAE. The overall pipeline of both DropMAE and TwinMAE is shown in Fig. 3. Finally, we introduce the VOT and VOS methods used for fine-tuning the downstream tasks.

3.1. Temporal Masked Autoencoder (TwinMAE)

The masked autoencoder (MAE) model [37] consists of an encoder and a decoder. Its basic idea is to randomly mask out a large portion (e.g., 75%) of patches in an image and then reconstruct the image pixels. Specifically, the encoder only takes visible patches as input for feature




Figure 3. An illustration of our DropMAE. The proposed adaptive spatial-attention dropout (ASAD) facilitates temporal correspondence learning for temporal matching tasks. TwinMAE follows the same pipeline except that the ASAD module is not used.

learning, and then the decoder is input with both visible and masked patches to produce the image reconstruction. In order to adapt to downstream video matching-based tasks, one naive extension is to directly apply MAE on concatenated video frames, hoping to learn temporal matching representations from video frame pairs, which we denote as TwinMAE.

It should be noted that existing works [31, 74] that extend MAE to video representation learning are mainly designed for the downstream task of video action recognition, where a long video clip (e.g., 16 frames) is used for reconstruction-based pre-training. To keep consistent with our downstream tasks of VOT/VOS, we follow the general training settings used in object tracking [3, 100], where two frames are sampled from one video as input to TwinMAE for pre-training. This adaptation significantly reduces the computational and memory cost compared to existing video pre-training approaches [31, 74], due to the quadratic complexity of ViTs.

Patch embedding. Firstly, we randomly sample 2 frames within a video with a predefined maximal frame gap. For each frame, we follow the vanilla ViT to divide it into non-overlapping patches. The patches extracted from the two frames are then concatenated together to form the overall patch sequence. We then randomly mask out patches in the patch sequence until a predefined mask ratio is reached. Note that we use the same mask ratio (i.e., 75%) with the original MAE, since the information redundancy of two frames should be similar to a single image. The visible patches are embedded by linear projection [25], and the masked patches are embedded using a shared learnable mask token. All the embedded patches are added with positional embeddings [25].

Frame identity embedding. To distinguish between the masked tokens in the same spatial location of the two frames, we use two learnable frame identity embeddings to indicate the two input frames. The corresponding frame identity embedding is added to each embedded patch.

Autoencoder and Training. Following the autoencoding pipeline in the original MAE [37], the encoder only takes visible embedded patches as input, and the decoder is input with all the embedded patches for masked patch reconstruction. We use the same normalized pixel loss from MAE for




Figure 4. The average within-frame and between-frame attention scores obtained by TwinMAE and DropMAE in different decoder layers. The attention score is calculated on 20 randomly sampled K400 validation videos, and is averaged on all heads and locations.

training the whole network architecture.

3.2. Limitation of TwinMAE Baseline

The visualization of the reconstruction for our TwinMAE baseline is shown in Fig. 2. We also quantitatively compare the average within-frame and between-frame attentions during the reconstruction in Fig. 4. Interestingly, we find the TwinMAE reconstruction heavily relies on within-frame patches or spatial cues, which may lead to sub-optimal temporal representations for matching-based video tasks.

When only using within-frame spatial cues, the decoder will perform the reconstruction using only context information in the neighboring patches, and thus the learned encoder representations will embed context information. In contrast, when using between-frame cues, the decoder will learn to perform matching of patches *between* frames so as to recover the corresponding target patch in the other frame. Thus, decoding with between-frame cues will make the encoder learn representations that support temporal matching between frames. Previous works in object tracking [3, 86] also suggest that temporal correspondence learning plays a key role in developing a robust and discriminative tracker. Since TwinMAE relies more on context information, it is still suboptimal for downstream tracking tasks.

3.3. Adaptive Spatial-Attention Dropout

To address issue of TwinMAE discussed in Sec. 3.2, we propose a *Adaptive Spatial-Attention Dropout* (ASAD) to facilitate the temporal correspondence learning in the temporal MAE. Given a query token, our basic idea is to adaptively drop a portion of its within-frame cues in order to facilitate the model to learn more reliable temporal correspondence, i.e., between-frame cues. That is, we restrict the interactions between the query token and tokens in the same frame, and encourage more interactions with tokens in the other frame, through manipulation of the computed spatial-attention in the transformer. Therefore, to minimize the reconstruction loss, the model is facilitated to learn a better temporal matching ability, which is essential in matching-based video tasks.

Before introducing the proposed ASAD, we firstly revisit



Figure 5. Visualization of the temporal matching function f_{tem} on an example frame pair. A large value of $f_{tem}(i)$ indicates that the i -th pixel matches well to a pixel in the other frame.

the multi-head self-attention in ViT [25]. Let $\mathbf{z} \in \mathbb{R}^{N \times D}$ be the input sequence of the two concatenated input frames, N denotes the total patch number in the two frames and D is the feature dimension. The standard multi-head self-attention [25] can be formulated as:

$$[\mathbf{q}, \mathbf{k}, \mathbf{v}] = \mathbf{z} \mathbf{U}_{qkv}, \quad \text{SA}(\mathbf{z}) = \text{softmax}\left(\frac{1}{\sqrt{D_k}} \mathbf{q} \mathbf{k}^T\right), \quad (1)$$

$$\text{MSA}(\mathbf{z}) = [\text{SA}_1(\mathbf{z}); \text{SA}_2(\mathbf{z}); \dots; \text{SA}_k(\mathbf{z})] \mathbf{U}_m, \quad (2)$$

where $\mathbf{U}_{qkv} \in \mathbb{R}^{D \times 3D_k}$ and $\mathbf{U}_m \in \mathbb{R}^{k \cdot D_k \times D}$. Let $A = \frac{\mathbf{q} \mathbf{k}^T}{\sqrt{D_k}} \in \mathbb{R}^{N \times N}$ denote the *attention matrix*. Our ASAD performs spatial-attention dropout on A so as to remove some within-frame interactions.

Temporal matching probability. We first need to consider the best tokens on which to apply ASAD. Intuitively, a query token that has a strong match in the other frame should be a good candidate, since, in the absence of within-frame cues, it can still be reconstructed well using the temporal cues in the other frame. Here, we define a temporal matching function $f_{tem}(\cdot)$ to measure the temporal matching probability of the i -th query token:

$$f_{tem}(i) = \max_{j \in \Omega_t(i)} (\hat{A}_{i,j}), \quad \hat{A} = \text{softmax}_{\text{row}}(A), \quad (3)$$

where the softmax function is applied on each row of A , $f_{tem}(i) \in [0, 1]$, and $\Omega_t(i)$ denotes the *temporal* index set of the i -th query token, which contains all the token indices of the other frame. A larger value of $f_{tem}(i)$ indicates a larger probability that the i -th query token is well-matched in the other frame, and thus a good candidate for ASAD. A visualization of $f_{tem}(\cdot)$ is shown in Fig. 5.

Overall dropout probability measurement. The overall spatial-attention dropout probability at location (i, j) is measured by using both the temporal matching probability and the normalized spatial importance:

$$W_{i,j} = f_{tem}(i) \frac{\hat{A}_{i,j}}{\sum_{j \in \Omega_s(i)} \hat{A}_{i,j}}, \quad (4)$$

where $\Omega_s(i)$ is the *spatial* index set that contains all the other token indices (i.e., excluding the query index itself) in the same frame as the i -th query. When $W_{i,j}$ is large, the i -th query token has a good between-frame match, and meanwhile the j -th within-frame token is important for the i -th

query. In this case, dropping the attention element (i, j) in A facilitates the model to use between-frame (temporally-matched) tokens for token learning or reconstruction. Finally, we set the dropout probability for self-attention and temporal-self-attention to be 0, i.e., $W_{i,i} = W_{i,i+N/2} = 0$.

Note that there are $N(N/2 - 1)$ (i.e., excluding self-attention elements) spatial-attention elements in total, and only these spatial-attention elements are considered for dropout. With a pre-defined dropout ratio P , we globally drop a total of $N_d = PN(N/2 - 1)$ attention elements from A .

Sampling for Dropout. We draw N_d elements from a multinomial distribution based on the dropout probability matrix W . Then we drop the elements in A with the corresponding indices by setting their values to $-\infty$. After applying the softmax function in (1), the corresponding spatial-attention weights are removed. The other operations are the same with the original multi-head self attention mechanisms used in ViT. The PyTorch-like pseudocode is presented in the Supplemental.

Our ASAD method has negligible additional time cost compared with TwinMAE, due to the efficient matrix operation in GPUs. We apply ASAD to each layer in the decoder during the pre-training stage, so as to learn encoder representations that support temporal matching. In the next section, we introduce downstream task fine-tuning based on the well pre-trained ViT model.

3.4. Downstream Temporal Matching Tasks

After obtaining the pre-trained DropMAE model, we fine-tune the well-learned encoder (i.e., the ViT model) on downstream temporal matching tasks, i.e., VOT and VOS.

3.4.1 Video Object Tracking

Recently, the MAE ViT models pre-trained on ImageNet are applied to VOT and show impressive results. We use the state-of-the-art tracker OSTRack [100] as our baseline tracker for fine-tuning. In OSTRack, the cropped template and search images are firstly serialized into sequences and concatenated together. Then the overall sequence is added with the positional embeddings and input to the ViT backbone for joint feature extraction and interaction. Finally, the updated search features are input to a prediction head to predict the target bounding box.

During the fine-tuning stage, we use our pre-trained DropMAE encoder weights to initialize the ViT backbone used in OSTRack. Meanwhile, two frame identity embeddings are respectively added to template and search embeddings, in order to keep consistency with the pre-training stage. We use the same training losses of the original OSTRack. The detailed training parameters are shown in Supplementary for reference.

3.4.2 Video Object Segmentation

For VOS, there are currently no methods based on ViT. Thus, we build a simple VOS baseline with a ViT backbone.

Input serialization. Given a template frame with a binary mask, VOS aims to segment the object-of-interest in each frame of a video. Similar to the pre-training stage, the binary mask map, template and search frames are firstly converted to patch sequences, and then linearly projected and added with positional embeddings. Two frame identity embeddings are added to the template and search embeddings, and the mask embeddings are added to the template embeddings for mask encoding.

Joint feature extraction and interaction. The template and search embeddings are concatenated together and input to the ViT backbone for joint feature extraction and matching. We use the updated search features extracted from the last layer of ViT for mask prediction.

Mask prediction. The existing VOS approaches [15, 16, 62, 99] employ multi-resolution features for mask prediction. However, the updated search features are single-resolution. We follow [47] to upsample the search features to $2\times$ and $4\times$ sizes via two deconvolutional modules. Finally, we use the same decoder used in [16, 62] for mask prediction.

Training loss. We use the commonly-used cross entropy loss [16, 62] to train the whole network architecture.

Online inference. During the online inference, we use the first frame with the mask annotation as the memory frame for online target matching in the search frame.

The pipeline figure and more implementation details can be found in Supplementary.

4. Experiments

4.1. Implementation Details

Datasets. In the pre-training stage, we explore various large-scale video data sources to pre-train our DropMAE model, including Kinetics-400 [41] (K400), Kinetics-600 [8] (K600), Kinetics-700 [9] (K700), Moments in Time [58] (MiT) and WebVid-2M [1]. The detailed performance comparison using different pre-training datasets is shown in the ablation study. For fine-tuning on VOT, we follow the training settings used in our baseline OSTRack [100]. We use the training splits of LaSOT [29], COCO [51], TrackingNet [59] and GOT-10k [40] for training. For the GOT-10k evaluation, we follow the one-shot evaluation and only fine-tune the model on the training split of GOT-10k. For VOS fine-tuning, we use Youtube-VOS [93] and Davis [66] datasets for fine-tuning following the standard convention [16, 62].

Pre-training and fine-tuning. For DropMAE pre-training, the default input is two frames with a spatial size of 224×224 pixels. The frames are randomly sampled from each video within a maximum frame gap of 50. During the pre-training, one epoch is counted when all videos are sampled

Methods	Pre-training Data	Epochs	Pre-train. Time (h)	GOT-10k (VOT)			DAVIS-17 (VOS)		
				AO	SR _{0.5}	SR _{0.75}	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
No Pre-training	-	-	-	62.7	72.8	53.7	69.5	66.9	72.2
Supervised IN1k [75]	IN1K	300	-	69.7	79.0	65.6	78.0	74.8	81.1
Supervised IN21k [68]	IN21K	80	-	70.2	80.7	65.4	78.5	75.4	81.7
CLIP [67]	IN1K	32	-	67.4	76.8	60.0	73.6	70.5	76.7
MOCO-v3 [13]	IN1K	300	-	70.1	80.1	65.3	78.4	75.4	81.5
BeiT [2]	IN1K	800	103.1	67.4	76.8	60.0	76.1	72.7	79.4
MAE [37]	IN1K	1600	84	73.7	83.2	70.8	81.7	78.5	84.9
TwinMAE	K400	400	20.7	72.2	83.2	65.9	79.3	76.4	82.3
TwinMAE	K400	800	41.3	72.9	83.6	68.5	80.7	77.9	83.6
TwinMAE	K400	1600	82.7	74.2	84.9	69.4	81.2	78.1	84.2
DropMAE	K400	400	21.1	73.2	83.9	67.5	81.3	78.5	84.0
DropMAE	K400	800	42.2	74.8	85.4	70.5	82.7	79.7	85.6
DropMAE	K400	1600	84.4	75.8	86.4	72.0	83.1	80.2	86.0
DropMAE	K700	800	92.4	75.9	86.8	72.0	83.0	80.2	85.7

Table 1. Comparison of pre-training methods on downstream VOT and VOS tasks on GOT-10k [40] and DAVIS-17 [66]. All methods adopt the ViT-B/16 model [25] with 224×224 input images for pre-training. The pre-training time is measured on 64 NVIDIA V100 GPUs. The best two results are shown in **red** and **blue**.

Method	Source	GOT-10k [40]			TNL2K [82]		LaSOT _{ext} [28]			LaSOT [29]		
		AO	SR _{0.5}	SR _{0.75}	AUC	P	AUC	P _{Norm}	P	AUC	P _{Norm}	P
SiamFC [3]	ECCVW16	34.8	35.3	9.8	29.5	28.6	23.0	31.1	26.9	33.6	42.0	33.9
MDNet [60]	CVPR16	29.9	30.3	9.9	-	-	27.9	34.9	31.8	39.7	46.0	37.3
ECO [20]	ICCV17	31.6	30.9	11.1	32.6	31.7	22.0	25.2	24.0	32.4	33.8	30.1
SiamPRN++ [43]	CVPR19	51.7	61.6	32.5	41.3	41.2	34.0	41.6	39.6	49.6	56.9	49.1
DiMP [4]	ICCV19	61.1	71.7	49.2	44.7	43.4	39.2	47.6	45.1	56.9	65.0	56.7
SiamR-CNN [77]	CVPR20	64.9	72.8	59.7	52.3	52.8	-	-	-	64.8	72.2	-
LTMU [19]	CVPR20	-	-	-	48.5	47.3	41.4	49.9	47.3	57.2	-	57.2
Ocean [107]	ECCV20	61.1	72.1	47.3	38.4	37.7	-	-	-	56.0	65.1	56.6
TrDiMP [79]	CVPR21	67.1	77.7	58.3	-	-	-	-	-	63.9	-	61.4
TransT [14]	CVPR21	67.1	76.8	60.9	50.7	51.7	-	-	-	64.9	73.8	69.0
AutoMatch [105]	ICCV21	65.2	76.6	54.3	47.2	43.5	37.6	-	43.0	58.3	-	59.9
STARK [95]	ICCV21	68.8	78.1	64.1	-	-	-	-	-	67.1	77.0	-
KeepTrack [57]	ICCV21	-	-	-	-	-	48.2	-	-	67.1	77.2	70.2
MixFormer-L [18]	CVPR22	70.7	80.0	67.8	-	-	-	-	-	70.1	79.9	76.3
SBT [90]	CVPR22	70.4	80.8	64.7	-	-	-	-	-	66.7	-	71.1
UAST [101]	ICML22	63.5	74.1	51.4	-	-	-	-	-	57.1	-	58.7
SwinTrack-384 [50]	NeurIPS22	72.4	80.5	67.8	55.9	57.1	49.1	-	55.6	71.3	-	76.5
AiATrack [33]	ECCV22	69.6	80.0	63.2	-	-	47.7	55.6	55.4	69.0	79.4	73.8
CIA50 [65]	ECCV22	67.9	79.0	60.3	50.9	57.6	-	-	-	66.2	-	69.6
SimTrack-L [10]	ECCV22	69.8	78.8	66.0	55.6	55.7	-	-	-	70.5	79.7	-
OSTrack-384 [100]	ECCV22	73.7	83.2	70.8	55.9	56.7	50.5	61.3	57.6	71.1	81.1	77.6
DropTrack	Ours	75.9	86.8	72.0	56.9	57.9	52.7	63.9	60.2	71.8	81.8	78.1

Table 2. Comparison with state-of-the-art VOT approaches on four large-scale challenging datasets. The best two results are shown in **red** and **blue**. For GOT-10k evaluation, all the methods follow the one-shot protocol, training only on the training set in GOT-10k.

once. For fair comparison, we use the same mask ratio (i.e., 75%) and training hyper-parameters of MAE [37] to pre-train the TwinMAE and DropMAE models. Following [15, 16], we use a bootstrapped cross entropy loss for training. The detailed pre-training and fine-tuning hyper-parameters are in the Supplementary.

4.2. Comparison with Pre-Training Methods

In Table 1, we compare our DropMAE with existing pre-training methods on the downstream tasks of VOT and VOS. DropMAE and TwinMAE are pre-trained using videos (K400, K700), while MAE and other methods are pre-trained on ImageNet 1k or 21k (IN1K, IN21K). The VOT and VOS baselines illustrated in Sec. 3.4 use the

official pre-trained ViT-B/16 models provided by existing pre-training approaches (see Table 1) for fine-tuning. TwinMAE with 800-epoch training performs favorably against MAE on VOT, but achieves inferior results on VOS. There are two main reasons: 1) TwinMAE is not effective enough at learning temporal matches; 2) The number of object classes in K400 is limited, and meanwhile the object classes in DAVIS-17 are included in ImageNet. Thus MAE generalizes well to VOS. Our DropMAE, which is a stronger temporal matching learner, outperforms MAE on both the VOT and VOS tasks with 800-epoch training (i.e., 42.2 hours¹)

¹In our implementation, we perform video decoding and extract the video frames on the fly in each training iteration, which could be further sped up using some tricks like repeated sampling [31].

Method	OTB100 AUC	ITB AUC	TrackingNet AUC P_{Norm}
SiamFC [3]	58.3	44.1	57.1 66.3
Ocean [107]	68.4	47.7	- -
ATOM [21]	68.3	47.2	70.3 77.1
DiMP [4]	53.7	339	74.0 80.1
TransT [14]	69.5	54.7	81.4 86.7
STARK [95]	68.1	57.6	82.0 86.9
OSTrack [100]	-	64.8	83.9 88.5
DropTrack	69.6	65.0	84.1 88.9

Table 3. Comparison with state-of-the-art VOT approaches on OTB100 [88], ITB [46] and TrackingNet [59]. The best two results are shown in red and blue.

by using the K400 dataset. This indicates that our DropMAE is $2\times$ faster than MAE.

5. State-of-the-art Comparison

In this section, we compare our fine-tuned VOT and VOS models, denoted as DropTrack and DropSeg, with state-of-the-art approaches on VOT and VOS benchmarks. We use DropMAE trained on K700 with 800 epochs as the pre-training model for both VOT and VOS fine-tuning.

5.1. Video Object Tracking

To demonstrate the effectiveness of the proposed DropMAE for VOT, we compare our DropTrack with state-of-the-art trackers on 7 challenging tracking benchmarks.

GOT-10k. GOT-10k [40] is a challenging dataset that follows the one-shot evaluation protocol, where the trackers are required to be trained on its training split, and the test object classes have no overlap with the objects in the training split. As shown in Table 2, our DropTrack achieves state-of-the-art results on this dataset, outperforming OSTrack by 2.2% and 3.6% in terms of AO and $SR_{0.5}$. This implies that the temporal correspondence learning in the pre-training is beneficial for the downstream tracking task. Although there exists a domain gap between the pre-training data and the test data (i.e., a large portion of test objects in GOT-10k are animals, vehicles and object parts, whereas K700 only consists of human-centric action videos), the temporal matching ability learned by DropMAE can still be transferred to the downstream tracking task, improving the tracking performance.

LaSOT. LaSOT consists of 280 long test sequences, and our results are presented in Table 2. Our DropTrack sets a new record on this dataset with 71.8% AUC, 81.8% P_{Norm} and 78.1% P, which shows the great potential of our DropTrack in robust long-term visual tracking.

LaSOT_{ext}. LaSOT_{ext} is an extension of LaSOT with more challenging video sequences for testing – similar to GOT-10k, the test split has a large gap with the training split, and sequences with novel object classes (i.e., not appeared in ImageNet) are used for evaluation. Our DropMAE outperforms the other trackers by large margins. This demon-

Method	Source	OL	M	S	DAVIS-2016 [64]			DAVIS-2017 [66]		
					$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
RANet [83]	ICCV19			✓	85.5	85.5	85.4	65.7	63.2	68.2
STM [62]	ICCV19		✓	✓	89.3	88.7	89.9	81.8	79.2	84.3
FRTM [69]	CVPR20	✓			83.5	83.6	83.4	76.7	73.9	79.6
TVOS [104]	CVPR20		✓		-	-	-	72.3	69.9	74.7
LWL [5]	ECCV20	✓	✓		-	-	-	81.6	79.1	84.1
CFBI [98]	ECCV20		✓		89.4	88.3	90.5	81.9	79.1	84.6
UniTrack [84]	NeurIPS21		✓		-	-	-	-	58.4	-
STCN- [16]	NeurIPS21		✓		-	-	-	82.5	79.3	85.7
SSTVOS [27]	CVPR21		✓		-	-	-	82.5	79.9	85.1
SWEM- [52]	CVPR22		✓		89.5	-	-	81.9	-	-
RTS [63]	ECCV22	✓	✓		-	-	-	80.2	77.9	82.6
OSMN [56]	TPAMI18				73.5	74.0	72.9	54.8	52.5	57.1
FAVOS [17]	CVPR18				81.0	82.4	79.5	58.2	54.6	61.8
VideoMatch [39]	ECCV18				-	81.0	-	56.5	-	-
SiamMask [80]	CVPR19				69.8	71.7	67.8	56.4	54.3	58.5
D3S [54]	CVPR20				74.0	75.4	72.6	60.8	57.8	63.8
Siam R-CNN [54]	CVPR20				-	-	-	70.6	66.1	75.0
Unicorn [94]	ECCV22				87.4	86.5	88.2	69.2	65.2	73.2
DropSeg	Ours				92.1	90.9	93.3	83.0	80.2	85.7

Table 4. Comparison with state-of-the-art VOS approaches on the validation sets of DAVIS-2016 [64] and DAVIS-2017 [66]. OL, M and S indicate Online Learning, using Memory mechanism, and using Synthetic videos for pre-training.




Figure 6. Ablation study of the dropout ratio P on the GOT-10k (VOT) and DAVIS-17 (VOS) datasets.

strates that a tracker with DropMAE pre-training generalizes well to unseen objects in generic visual object tracking. **TNL2K.** TNL2K is a large-scale evaluation dataset that consists of 700 test videos with various challenges, such as significant appearance variation and manually added adversarial samples. As illustrated in Table 2, our DropMAE significantly outperforms the other trackers on this dataset. **ITB, TrackingNet and OTB100.** In Table 3, we evaluate our DropTrack on ITB [46], OTB100 [88] and TrackingNet [59], achieving state-of-the-art performance on each one. DropTrack is slightly better than OSTrack on ITB and TrackingNet. We believe the main reason is the fully overlapped training and test object classes in these two datasets, which reduces the effect of pre-training. A competitive tracker on these two datasets can be learned even using supervised ImageNet weights, which has been shown in [100].

On all 7 VOT datasets, our DropTrack outperforms the baseline OSTrack, which demonstrates that our DropMAE pre-training on videos learns better temporal-matching representations than the MAE model trained on ImageNet, resulting in more a robust tracker that generalizes well to both unseen and seen objects.

5.2. Video Object Segmentation

In Table 4, we compare our DropSeg with existing VOS approaches on the DAVIS-16/17 [64, 66].

DAVIS-16. DAVIS-16 is composed of 20 manually anno-

Datasets	No. Videos	No. Actions	VOT			VOS
			AO	SR _{0.5}	SR _{0.75}	$\mathcal{J}\&\mathcal{F}$
K400 [41]	240,000	400	73.2	83.9	67.5	82.7
K600 [8]	390,000	600	74.5	85.5	69.5	82.8
K700 [9]	526,768	700	75.6	86.2	71.4	83.0
MiT [58]	802,244	339	75.1	85.5	70.6	82.8
WebVid [1]	240,000	-	72.8	83.4	67.3	81.5
WebVid [1]	960,000	-	73.4	85.0	69.5	82.9

Table 5. The downstream VOT and VOS performance on GOT-10k and DAVIS-17 obtained by using our DropMAE pre-trained on various video datasets. VOS uses 800 epochs pre-training.

tated test sequences. As shown in Table 4, our one-shot DropSeg approach, without using any online learning and complicated memory mechanisms, achieves the best $\mathcal{J}\&\mathcal{F}$ score of 92.1%, which significantly outperforms the other compared one-shot approaches and is even better than the approaches with complicated pipelines (i.e., OL, M and S). This implies that the pixel-wise correspondence learned during the pre-training is effective for capturing long-range dependencies between various frames in VOS.

DAVIS-17. DAVIS-17 is an extension of DAVIS-16, comprising more challenging videos and supports multi-object segmentation. In Table 4, our DropSeg achieves competitive results of 83.0% $\mathcal{J}\&\mathcal{F}$, 80.2% \mathcal{J} and 85.7% \mathcal{F} , which shows its superiority in handling more challenging videos. More VOS results and comparisons are included in the **supplementary**.

6. Ablation Studies

In this section, we conduct ablation studies to provide more detailed analysis of our method. We use DropMAE with 400-epoch pre-training for the ablation study.

The effect of dropout ratio P . We study the effect of dropout rate P in Fig. 6. A relatively small dropout ratio of $P = 0.1$ works well on both VOT and VOS tasks. Meanwhile, dropping too many spatial cues (e.g., $P=0.2$) degrades the downstream tasks, which is mainly because the spatial cues are also useful for accurate localization and segmentation. $P = 0.1$ is the optimal setting, and thus we adopt it in the following experiments.

Pre-training video sources. Since we are the first to explore masked autoencoder pre-training for temporal matching tasks, it is not clear which video dataset is the optimal choice for pre-training. Here, we use five popular video datasets for pre-training, including K400 [41], K600 [8], K700 [9], MiT [58] and WebVid-2M [1]. For WebVid-2M, we randomly sample 240k and 960k videos for fair comparison and faster validation. The downstream tracking results are reported in Table 5. Performance is not favorable even using 960,000 videos in WebVid for pre-training. This indicates that WebVid is not a good choice for tracking pre-training, which is mainly because it is a video caption dataset that focuses on scene diversity and lacks rich object motion. From the experiments on the K400/600/700 and

Settings	GOT-10k
DropTrack-K400-400E	73.2/67.5
w/ ASAD in Encoder	73.1/68.1
w/ domain specific data	73.4/68.8
w/o frame identity embed	72.9/67.4

Table 6. The tracking performance of AO/SR_{0.75} on GOT10-k reported by variants with different settings.

MiT, tracking benefits from pre-training with from rich action classes (i.e., 700 action classes of K700), from which the model can learn stronger temporal matching ability.

Applying ASAD to the encoder. Here we test applying ASAD to all layers including both encoder and decoder of masked autoencoder. As shown in Table 6, this variant gains improvements (0.6% in SR_{0.75}) over the original baseline. Considering its additional cost and limited performance improvement, we only apply ASAD to the decoder.

Domain specific data. We also add the tracking training data (without using box annotations), including TrackingNet, LaSOT, and GOT-10k, into K400 for pre-training. The downstream tracking performance by using the larger pre-training set is 73.4/68.8, which is better than the baseline. It shows that the domain-specific data is helpful to bridge the domain gap, which can be considered as future work to extend Kinetic datasets with more tracking videos.

Frame identity embedding. During pre-training, the frame identity embedding is used to identify masked patches in the same 2D location of the two frames. From Table 6, we can find that downstream fine-tuning without the frame identity embedding performs worse than with it, since not using it is inconsistent with the pre-training stage.

7. Conclusion

This paper investigated masked autoencoding pre-training for temporal matching-based downstream tasks. Specifically, an adaptive spatial-attention dropout method is proposed to facilitate temporal correspondence learning in self-supervised pre-training on videos. We show that our proposed pre-training method DropMAE can achieve better downstream performance on VOT and VOS than the image-based MAE, while using 50% less pre-training time. In addition, as the first work investigating this problem, we show the guidelines of selecting video sources for pre-training, i.e., selecting videos with rich motion information is more beneficial for temporal matching-based downstream tasks. The experimental results on VOT and VOS show that our DropMAE sets new state-of-the-art results on 8 out of 9 tracking benchmarks. We expect our DropMAE to serve as a strong pre-trained backbone in future work.

8. Acknowledgment

This work was supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. CityU 11215820).

Algorithm 1: ASAD Pseudocode, PyTorch-like

```

# Input: attention matrix A, sequence
length N, drop number  $N_d$ 
W = torch.zeros_like(A) # N-by-N
A = A.detach().softmax(dim=-1) # N-by-N

# get temporal attentions in each row of A
A_tem = temporal_index(A) # N-by-N//2
f_tem = A_tem.max(dim=-1).values # N-by-1

# get spatial attentions in each row of A
A_spa = spatial_index(A) # N-by-N//2
# avoid self-attention dropout
A_spa[0:N//2, 0:N//2].fill_diagonal_(0)
A_spa[N//2:, 0:N//2].fill_diagonal_(0)
A_spa = A_spa / A_spa.sum(dim=-1, keepdim=True)

# calculate overall dropout probability
f_all = f_tem * A_spa # N-by-N//2

# put back to probability matrix W
W[0:N//2, 0:N//2] = f_all[0:N//2, 0:N//2]
W[N//2:, N//2:] = f_all[N//2:, 0:N//2]
# sample  $N_d$  elements based on W

indices = torch.multinomial(W.view(1, -1),  $N_d$ )
return indices

```

In this supplementary material, we provide the additional implementation details, pseudocode, overall pipeline, additional ablation study and qualitative visualization. Sec. I shows the PyTorch-like pseudocode of the proposed adaptive spatial-attention dropout (ASAD). Sec. J details the proposed VOS baseline with the ViT backbone. Sec. K illustrates the implementation details in both pre-training and downstream fine-tuning stages. More ablation studies are conducted in Sec. L. Sec. M shows the qualitative visualization of downstream VOT and VOS results, the video reconstruction results of our DropMAE and the additional quantitative VOS results. We finally discuss the limitation and future work of our DropMAE in Sec. N.

I. Algorithm

We show the pseudocode of our proposed adaptive spatial-attention dropout (ASAD) in Algorithm 1. As can be seen, the implementation is simple and neat, which could be flexibility incorporated into existing approaches like MAE [37]. The implementation mainly consists of three steps: 1) the calculation of temporal matching probability f_{tem} ; 2) the calculation of spatially normalized within-frame attention A_{spa} ; 3) sampling for dropout from a multinomial distribution. Our code and pre-trained models will be publicly available once the paper is published.

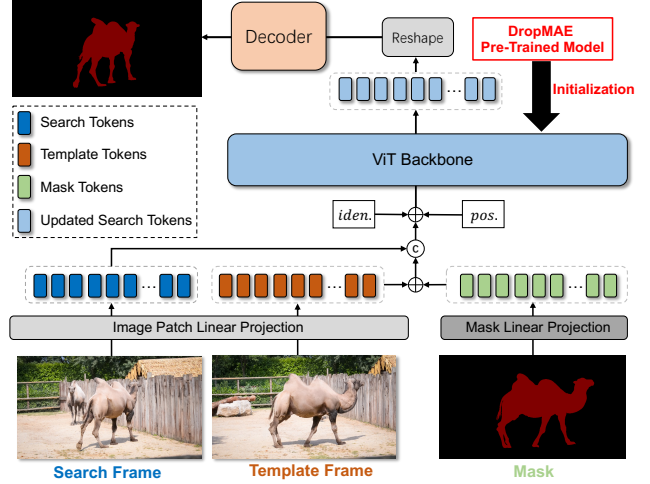


Figure 7. An overall pipeline of the proposed DropSeg for VOS. *iden* and *pos* indicate frame identity embeddings and positional embeddings, respectively. Our DropSeg with the DropMAE pre-trained model sets new state-of-the-art one-shot segmentation results on the DAVIS-16 [64] and DAVIS-17 [66] datasets.

J. VOS Baseline

Since there is no ViT-based VOS approach, we build a simple and effective ViT-based VOS baseline, namely DropSeg, in order to demonstrate the effectiveness of our DropMAE pre-training in VOS.

Architecture. The overall pipeline of our DropSeg is shown in Fig. 7, which mainly consists of the pre-trained ViT backbone and the mask prediction head. Note that the frame identity embeddings are two randomly initialized learnable vectors. We use the standard ViT-B/16 model [25] as the backbone and initialized it with our DropMAE pre-trained model, and the same decoder used in [16, 62] is employed. Since the decoder requires multi-resolution features for mask prediction, we follow [47] to upsample the updated search features to $2\times$ and $4\times$ sizes via two deconvolutional modules.

Fine-tuning. During the fine-tuning stage, given a video, we randomly sample a template frame with the mask annotation, and a search frame of a training video within a predefined maximum frame gap (i.e., 25). Following the convention [62], DAVIS-17 [66] and YouTube-19 [93] datasets are used for training. The detailed training hyper-parameters are in Table 9. We use the same bootstrapped cross entropy loss in [16, 62] for supervision.

K. Implementation Details

In this section, we detail the implementation details of our DropMAE pre-training, and fine-tuning details of downstream VOT and VOS tasks.

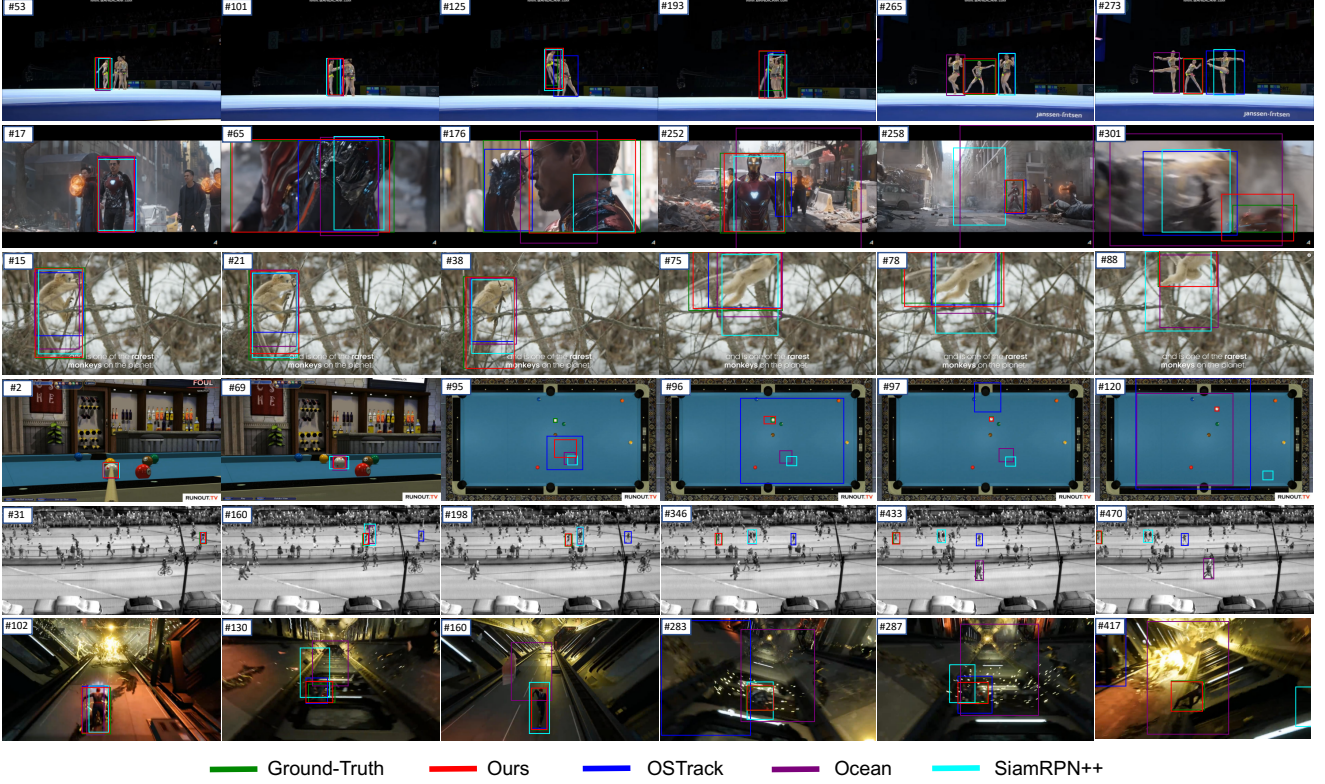


Figure 8. Qualitative tracking results of our DropTrack and state-of-the-art tracking methods, including OTrack [100], Ocean [107] and SiamRPN++ [43]. The six video sequences are collected from TNL2K [82] (from top to down are *SportGirl.video.01_done*, *test_031_IronMan_transform_05_done*, *Monkey.BBC.video.01-Done*, *ZhuoQiu.video.02-Done*, *INF_crow1* and *Zhizhuxia.09-Done*, respectively). The frame number is shown in the *top-left* of each frame.

Config	Value
optimizer	AdamW [42]
base learning rate	1.5e-4
weight decay	0.05
optimizer momentum	$\beta_1, \beta_2 = 0.9, 0.95$ [11]
batch size	4096
learning rate schedule	cosine decay
warmup epochs	40
augmentation	RandomResizedCrop
dropout ratio p	0.1
maximum sampling frame gap	50

Table 7. The pre-training setting for DropMAE.

Config	Value
optimizer	AdamW [42]
learning rate in head	2.5e-4
learning rate in backbone	2.5e-5
weight decay	0.0001
droppath rate	0.1
batch size	128
epoch	300
learning rate decay epoch	240
learning rate decay factor	10
learning rate schedule	steplr
maximum sampling frame gap	200

Table 8. The fine-tuning setting for DropTrack.

K.1. DropMAE Pre-Training

We use the standard ViT-B/16 [25] as our backbone for pre-training. The detailed pre-training hyper-parameters are in Table 7, which mainly follows the training settings used in the original MAE [37]. For the dropout ratio P , we set $P = 0.1$ following the ablation study in Sec. 6 of the main paper. The pre-training is conducted on 64 NVIDIA V100

GPUs. As illustrated in Table 1 of the main paper, the 1600-epoch pre-training takes about 84 hours on the K400 dataset [41]. The whole training time can be further reduced to 58 hours by using 64 NVIDIA A100 GPUs.




Figure 9. Qualitative results of our one-shot approach DropSeg on seven sequences in DAVIS-17 [66], which are respectively *bike-packing*, *bmx-trees*, *camel*, *breakdance*, *car-shadow*, *india* and *soapbox*. The frame number is shown in the *top-left* of each frame, and the ground-truth mask annotation is given in the first frame. Our DropSeg shows favorable segmentation results without using complicated designs such as online fine-tuning and memory mechanisms.




Figure 10. Video frame reconstruction results of DropMAE on K400 validation set. We show the original input frame pairs, masked frame pairs (i.e., with 75% mask ratio) and reconstruction results, sequentially.

Config	Value
optimizer	AdamW [42]
base learning rate	2e-5
weight decay	1e-7
droppath rate	0.1
batch size	32
Iterations	210,000
learning rate decay iteration	125,000
learning rate schedule	steplr
maximum sampling frame gap	25

Table 9. The fine-tuning setting for DropSeg.

K.2. Downstream VOT Fine-Tuning

Following the tracking baseline OTrack [100], we use the template and search sizes of 192×192 and 384×384 pixels, respectively, and fine-tune our DropMAE model with the tracking specific data (see Sec. 4.1 of the main paper) on 4 NVIDIA A100 GPUs. The candidate elimination module proposed in OTrack are also used for a fair comparison. For the full 300-epoch fine-tuning, the detailed hyper-parameters are in Table 8. For fine-tuning on GOT-10k [40], the total training epoch is reduced to 100 and the learning rate decays at 80 epoch. The inference speed of our DropTrack is the same as the baseline OTrack [100], which is 58.1 FPS measured on a single GPU.

K.3. Downstream VOS Fine-Tuning

The detailed fine-tuning setting of DropSeg is shown in Table 9. We use 8 A100 GPUs for fine-tuning, and the whole training takes about 16 hours. The inference speed of DropSeg is about 10 FPS, which is measured on a single A100 GPU.

L. Additional Ablation Studies

In this section, we provide additional ablation studies on parameter selection and effectiveness analysis. We use the model pre-trained on K400 with 400 epochs for the ablation study.

Effect of maximum sampling frame gap. During the pre-training, we randomly sample two frames of a training video with a predefined maximal sampling frame gap g . Here, we study its effect on the downstream VOT task. As shown in Table 10, the VOT task benefits more from the large sampling frame gap, i.e., 50. This is because the stronger temporal matching ability can be learned by using the relatively large sampling frame gap. Since the limited performance improvements from $g = 10$ to $g = 50$, we directly use $g = 50$ for all the pre-training experiments without further searching for the parameters.

Learning static frame representation from K400. To demonstrate the temporal correspondence learning in the

Maximum Sampling Frame Gap	GOT-10k		
	AUC	SR _{0.5}	SR _{0.75}
1	72.2	82.7	65.7
10	72.8	83.4	67.2
50	73.2	83.9	67.5

Table 10. The effect of maximum sampling frame gap on the downstream tracking task.

Methods	GOT-10k		
	AUC	SR _{0.5}	SR _{0.75}
DropMAE	73.2	83.9	67.5
MAE-K400-static	70.4	80.7	65.6

Table 11. The comparison between DropMAE and MAE-K400-static on GOT-10k [40].

pre-training is the key to the success of downstream tracking tasks, we treat K400 [41] as a static image dataset and perform the original MAE pre-training on it. We denote this baseline as MAE-K400-static. Specifically, in each training iteration, one frame image is randomly sampled of a video for masked autoencoding pre-training. To make a fair comparison with our DropMAE, we double the video number in this baseline such that the total sampled frame number in one epoch training is the same as DropMAE. The comparison between MAE-K400-static and DropMAE is shown in Table 11. As can be seen, without temporal correspondence learning, MAE-K400-static is significantly worse than our DropMAE, which further demonstrates the effectiveness of the temporal correspondence learning in the DropMAE pre-training.

Random dropout. The vanilla ViT [25] implements dropout [71] in each multi-head self-attention layer. To see whether this random dropout works in our masked autoencoding pre-training setting, we build a baseline called RandDrop-MAE, which adopts the random dropout in each self-attention layer of the decoder during the pre-training. Different from our adaptive dropout strategy (i.e., ASAD), RandDrop-MAE randomly drops between-frame or within-frame attentions. For a fair comparison, we use the same dropout ratio (i.e., 0.1) for RandDrop-MAE. As shown in Table 13, RandDrop-MAE degrades the performance compared with our DropMAE. We believe this is because the random dropout may excessively drop some attention elements that are essential for reconstruction and thus degrade the learning.

Pre-trained MAE. The downstream VOT and VOS tasks consist of large amounts of objects with diverse classes for evaluation. Considering that K400 is composed of human-action videos, there still exists domain gap between the pre-training and fine-tuning stages. In order to alleviate this gap, we use the original MAE trained on ImageNet as the pre-training weights of our DropMAE, and then we further pre-train our DropMAE on K400 for temporal correspondence learning. From Table 14, we can find that our DropMAE benefits from the pre-trained MAE on both VOT and VOS

Method	#Mem Static BL30K			DAVIS-17			YouTubeVOS-19 val				
				$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}_s	\mathcal{F}_s	\mathcal{J}_u	\mathcal{F}_u
XMEM	~27	✓	✓	87.7	84.0	91.4	85.8	84.8	89.2	80.3	88.8
STCN	27	✓	✓	85.3	82.0	88.6	84.2	82.6	87.0	79.4	87.7
XMEM	~27	✓		86.2	82.9	89.5	85.5	84.3	88.6	80.3	88.6
STCN	27	✓		85.4	82.2	88.6	82.7	81.1	85.4	78.2	85.9
SWEM	27	✓		84.3	81.2	87.4	82.6	82.0	86.1	77.2	85.2
DropSeg⁺	2	✓		86.5	83.5	89.5	83.4	82.9	87.3	77.7	85.6
STCN ⁻	27			82.5	79.3	85.7	-	-	-	-	-
DropSeg	1			83.0	80.2	85.7	-	-	-	-	-

Table 12. ‘Static’ and ‘BL30K’ indicate pre-training with static images and large-scale video BL30K. #Mem is the estimated number of memory frames, which is measured w/ the avg. video length in YouTubeVOS-19.

Methods	GOT-10k		
	AUC	SR _{0.5}	SR _{0.75}
DropMAE	73.2	83.9	67.5
RandDrop-MAE	71.7	82.4	66.2

Table 13. The comparison between DropMAE and RandDropMAE on GOT-10k [40].

Pre-trained MAE	GOT-10k			Davis-17 $\mathcal{J}\&\mathcal{F}$
	AUC	SR _{0.5}	SR _{0.75}	
w/o	73.2	83.9	67.5	81.3
w/	75.2	85.4	71.5	82.6

Table 14. The ablation study on the usage of the pre-trained MAE model for DropMAE pre-training.

tasks, which is mainly because the diverse object classes learned in MAE are beneficial for generic object tracking and segmentation. This also shows the potential that the better downstream performance can be achieved by using the pre-trained MAE and larger video data sources (e.g., K700 [9]).

M. Qualitative Visualization

In this section, we show the qualitative results of our DropTrack and DropSeg on the VOT and VOS tasks, respectively.

M.1. Video Object Tracking

In Fig. 8, we show the qualitative tracking results obtained by our DropTrack and the other 3 compared trackers. The selected sequences contain various challenges including significant appearance variation, background cluster, illumination variation and similar objects. Our DropTrack handles these challenges well due to the robust DropMAE pre-trained model.

M.2. Video Object Segmentation

The qualitative visualization of our DropSeg is shown in Fig. 9. Even without using online fine-tuning or complicated memory mechanisms, our DropSeg can still provide accurate segmentation results in the following frames

by only using the mask annotation in the first frame, which is mainly due to the favorable temporal matching ability learned in the DropMAE pre-training.

M.3. Frame Reconstruction

We show the video frame reconstruction results obtained by our DropMAE in Fig. 10. As can be seen, although less spatial cues are leveraged in the reconstruction, our DropMAE still achieves favorable reconstruction results by exploring temporal cues or between-frame patches.

M.4. Additional Comparison on VOS

In this subsection, we provide the additional comparisons on DAVIS-17 and YouTube-VOS 19 datasets, which are illustrated in Table 12. DropSeg does not use online memory mechanisms, so we have used short videos (e.g., DAVIS-16/17) to focus evaluation on the learned representations. To more fairly compare on the longer YouTube-VOS videos, we employ an improved variant DropSeg⁺, which 1) uses static images for video pre-training (as in [16, 62]); 2) uses the first and previously predicted frames as memory. In Tab. 12, although our DropSeg⁺ only uses 2 memory frames, it is better than STCN w/ static image pre-training on DAVIS-17 & YouTube-19. DropSeg⁺ is comparable to XMEM on DAVIS-17, but worse on YouTube-19, which is mainly due to the lack of complex online memory and BL30K pre-training.

N. Limitation and Future Work

Due to the limited object classes in video datasets (e.g., K400 [41] and K700 [9]), the pre-training video sources still have large domain gap with the downstream VOT and VOS tasks, which results in a sub-optimal pre-trained model. In Table 14, we find that this gap can be alleviated by using the pre-trained MAE model during the DropMAE pre-training. In the future work, we will perform more large-scale DropMAE pre-training (i.e., w/ the MAE pre-trained model) on

larger video sources in order to provide more robust pre-trained models for VOT and VOS communities.

References

- [1] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, 2021. 5, 8
- [2] H. Bao, L. Dong, and F. Wei. Beit: Bert pre-training of image transformers. In *arXiv:2106.08254*, 2021. 6
- [3] L. Bertinetto, J. Valmadre, J.F. Henriques, A. Vedaldi, and P.H.S. Vedaldi. Fully-convolutional siamese networks for object tracking. In *ECCVW*, pages 850–865, 2016. 2, 3, 4, 6, 7
- [4] G. Bhat, M. Danelljan, L. V. Gool, and R. Timofte. Learning discriminative model prediction for tracking. In *ICCV*, pages 6182–6191, 2019. 6, 7
- [5] G. Bhat, F.J. Lawin, M. Danelljan, A. Robinson, M. Felsberg, L.V. Gool, and R. Timofte. Learning what to learn for video object segmentation. In *ECCV*, 2020. 7
- [6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, and Arvind Nee-lakantan. Language models are few-shot learners. In *NeurIPS*, 2020. 1, 3
- [7] Yue Cao, Zhenda Xie, Bin Liu, Youtong Lin, Zheng Zhang, and Han Hu. Parametric instance classification for unsupervised visual feature learning. In *NeurIPS*, 2020. 3
- [8] João Carreira, Eric Noland, Andras Banki-Horvath, Chloe Hillier, and Andrew Zisserman. A short note about kinetics-600. In *arXiv:1808.01340*, 2018. 5, 8
- [9] João Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics700 human action dataset. In *arXiv:1907.06987*, 2019. 5, 8, 13
- [10] B. Chen, P. Li, L. Bai, L. Qiao, Q. Shen, and B. Li. Backbone is all your need: A simplified architecture for visual object tracking. In *ECCV*, 2022. 1, 2, 3, 6
- [11] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pre-training from pixels. In *ICML*, 2020. 10
- [12] T. Chen, S. Kornblith, and M. Norouzi. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 3
- [13] X. Chen, S. Xie, and K. He. An empirical study of training self-supervised vision transformers. In *ICCV*, 2021. 6
- [14] X. Chen, B. Yan, J. Zhu, D. Wang, X. Yang, and H. Lu. Transformer tracking. In *CVPR*, pages 8126–8135, 2021. 6, 7
- [15] H. K. Cheng and A. G. Schwing. Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model. In *ECCV*, 2022. 3, 5, 6
- [16] H. K. Cheng, Y. W. Tai, and C. K. Tang. Rethinking space-time networks with improved memory coverage for efficient video object segmentation. In *NeurIPS*, pages 11781–11794, 2021. 2, 5, 6, 7, 9, 13
- [17] J. Cheng, Y.H. Tsai, W.C. Hung, S. Wang, and M.H. Yang. Fast and accurate online video object segmentation via tracking parts. In *CVPR*, 2018. 7
- [18] Yutao Cui, Cheng Jiang, Limin Wang, and Gangshan Wu. Mixformer: End-to-end tracking with iterative mixed attention. In *CVPR*, 2022. 6
- [19] K. Dai, Y. Zhang, D. Wang, J. Li, H. Lu, and X. Yang. High-performance long-term tracking with meta-updater. In *CVPR*, 2020. 6
- [20] M. Danelljan, G. Bhat, F. S. Khan, and M. Felsberg. Eco: Efficient convolution operators for tracking. In *CVPR*, pages 21–26, 2017. 2, 6
- [21] M. Danelljan, G. Bhat, F. S. Khan, and M. Felsberg. Atom: Accurate tracking by overlap maximization. In *CVPR*, pages 4660–4669, 2019. 7
- [22] M. Danelljan, A. Robinson, F. S. Khan, and M. Felsberg. Beyond correlation filters: learning continuous convolution operators for visual tracking. In *ECCV*, pages 472–488, 2016. 2
- [23] Acob Devlin, Ming-Wei Chang, Kenton Lee, , and Kristina Toutanova. Language models are few-shot learners. In *NAACL*, 2019. 1, 3
- [24] X. Dong and J. Shen. Triplet loss in siamese network for object tracking. In *ECCV*, 2018. 2
- [25] A. Dosovitskiy, L. Beyer, and A. Kolesnikov. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 3, 4, 6, 9, 10, 12
- [26] Alexey Dosovitskiy, Jost Tobias, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *NeurIPS*, pages 766–774, 2014. 3
- [27] Brendan Duk, Abdalla Ahmed, Christian Wolf, Parham Aarabi, and Graham W. Taylor. Sstvos: Sparse spatiotemporal transformers for video object segmentation. In *CVPR*, 2021. 7
- [28] H. Fan, H. Bai, L. Lin, F. Yang, P. Chu, G. Deng, S. Yu, M. Huang, J. Liu, and Y. Xu. Lasot: A high-quality large-scale single object tracking benchmark. In *IJCV*, 2021. 6
- [29] H. Fan, L. Lin, and F. Yang. Lasot: A high-quality benchmark for large-scale single object tracking. In *CVPR*, pages 5374–5383, 2019. 5, 6
- [30] H. Fan and H. Ling. Siamese cascaded region proposal networks for real-time visual tracking. In *ICCV*, 2019. 2
- [31] C. Feichtenhofer, H. Fan, Y. Li, and K. He. Masked autoencoders as spatiotemporal learners. In *arXiv:2205.09113*, 2022. 3, 6
- [32] H. Galoogahi, A. Fagg, and S. Lucey. Learning background-aware correlation filters for visual tracking. In *ICCV*, 2017. 2
- [33] Shenyuan Gao, Chunlun Zhou, Chao Ma, Xinggang Wang, and Junsong Yuan. Attention in attention for transformer visual tracking. In *ECCV*, 2022. 6
- [34] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Un-supervised representation learning by predicting image rotations. In *arXiv:1803.07728*, 2018. 3
- [35] Q. Guo, W. Feng, C. Zhou, R. Huang, L. Wan, and S. Wang. Learning dynamic siamese network for visual object tracking. In *ICCV*, pages 1763–1771, 2017. 2
- [36] K. He and H. Fan and Y. Wu. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020. 3

- [37] K. He, X. Chen, and S. Xie. Masked autoencoders are scalable vision learners. In *CVPR*, pages 16000–16009, 2022. [1](#), [3](#), [6](#), [9](#), [10](#)
- [38] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista. High-speed tracking with kernelized correlation filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(3):583–596, 2015. [2](#)
- [39] Y.T. Hu, J.B. Huang, and A.G. Schwing. Videomatch: Matching based video object segmentation. In *ECCV*, 2018. [7](#)
- [40] L. Huang, X. Zhao, and K. Huang. Got-10k: A large high-diversity benchmark for generic object tracking in the wild. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. [1](#), [5](#), [6](#), [7](#), [12](#), [13](#)
- [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. In *arXiv:1705.06950*, 2017. [5](#), [8](#), [10](#), [12](#), [13](#)
- [42] Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. In *arXiv:1412.6980*, 2014. [10](#), [12](#)
- [43] B. Li, W. Wu, Q. Wang, F. Zhang, J. Xing, and J. Yan. Siamrpn++: Evolution of siamese visual tracking with very deep networks. In *CVPR*, 2019. [2](#), [6](#), [10](#)
- [44] B. Li, W. Wu, Z. Zhu, and J. Yan. High performance visual tracking with siamese region proposal network. In *Proceedings of the CVPR*, pages 8971–8980, 2018. [2](#)
- [45] N. Li, S. Liu, and Y. Liu. Neural speech synthesis with transformer network. In *AAAI*, pages 6706–6713, 2019. [1](#)
- [46] X. Li, Q. Liu, W. Pei, Q. Shen, Y. Wang, H. Lu, and M.H. Yang. An informative tracking benchmark. In *arXiv:2112.06467*, 2021. [7](#)
- [47] Yanghao Li, Hanzi Mao, Ross Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. In *ECCV*, 2022. [5](#), [9](#)
- [48] Y. Liang, Q. Wu, Y. Liu, and H. Wang. Robust correlation filter tracking with shepherded instance-aware proposals. In *ACM MM*, 2018. [2](#)
- [49] Y. Liang, Q. Wu, Y. Liu, Y. Yan, and H. Wang. Deep correlation filter tracking with shepherded instance-aware proposals. In *IEEE Transactions on Intelligent Transportation Systems*, 2021. [2](#)
- [50] L. Lin, H. Fan, Y. Xu, and H. Ling. Swintrack: A simple and strong baseline for transformer tracking. In *arXiv:2112.00995*, 2021. [6](#)
- [51] T. Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. [5](#)
- [52] Z. Lin, T. Yang, M. Li, Z. Wang, C. Yuan, W. Jiang, and W. Liu. Swem: Towards real-time video object segmentation with sequential weighted expectation-maximization. In *CVPR*, pages 1362–1372, 2022. [3](#), [7](#)
- [53] Y. Liu, Y. Liang, Q. Wu, L. Zhang, and H. Wang. A new framework for multiple deep correlation filters based object tracking. In *ICASSP*, 2022. [2](#)
- [54] A. Lukezic, J. Matas, and M. Kristan. D3s: a discriminative single shot segmentation tracker. In *CVPR*, 2020. [7](#)
- [55] A. Lukezic and T. Vojir. Discriminative correlation filter with channel and spatial reliability. In *CVPR*, 2017. [2](#)
- [56] K.K. Maninis, S. Caelles, Y. Chen, J. Pont-Tuset, L. Leal-Taix'e, D. Cremers, and L. Van Gool. Video object segmentation without temporal information. In *TPAMI*, 2018. [7](#)
- [57] C. Mayer, M. Danelljan, D.P. Paudel, and L. Van Gool. Learning target candidate association to keep track of what not to track. In *ICCV*, 2021. [6](#)
- [58] M. Monfort, A. Andonian, B. Zhou, K. Ramakrishnan, S. A. Bargal, T. Yan, and A. Oliva. Moments in time dataset: one million videos for event understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(2):502–508, 2019. [5](#), [8](#)
- [59] M. Muller, A. Bibi, and Giancola S. Trackingnet: A large-scale dataset and benchmark for object tracking in the wild. In *ECCV*, pages 300–317, 2018. [5](#), [7](#)
- [60] H. Nam and B. Han. Learning multi-domain convolutional neural networks for visual tracking. In *CVPR*, pages 4293–4302, 2016. [6](#)
- [61] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016. [3](#)
- [62] S. W. Oh, J. Y. Lee, N. Xu, and S. J. Kim. Video object segmentation using space-time memory networks. In *ICCV*, pages 9226–9235, 2019. [2](#), [5](#), [7](#), [9](#), [13](#)
- [63] Matthieu Paul, Martin Danelljan, Christoph Mayera, , and Luc Van Gool. Robust visual tracking by segmentation. In *ECCV*, 2022. [7](#)
- [64] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. Van Gool, M. Gross, and A. Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, 2016. [7](#), [9](#)
- [65] Zhixiong Pi, Weitao Wan, Chong Sun, Changxin Gao, Nong Sang, and Chen Li. Hierarchical feature embedding for visual tracking. In *ECCV*, 2022. [6](#)
- [66] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool. The 2017 davis challenge on video object segmentation. In *arXiv:1704.00675*, 2017. [5](#), [6](#), [7](#), [9](#), [11](#)
- [67] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, and I. Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. [6](#)
- [68] T. Ridnik, E. Ben-Baruch, A. Noy, and L. Zelnik-Manor. Imagenet-21k pretraining for the masses. In *arXiv:2104.10972*, 2021. [6](#)
- [69] A. Robinson, F.J. Lawin, M. Danelljan, F.S. Khan, and M. Felsberg. Learning fast and robust target models for video object segmentation. In *CVPR*, 2020. [7](#)
- [70] O. Russakovsky, J. Deng, and et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015. [1](#)
- [71] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. In *The journal of machine learning research*, 2014. [12](#)

- [72] N. Srivastava, E. Mansimov, and R. Salakhudinov. Unsupervised learning of video representations using lstms. In *ICML*, 2015. 3
- [73] R. Tao, E. Gavves, and A. W.M. Smeulders. Siamese instance search for tracking. In *CVPR*, pages 1420–1429, 2016. 2
- [74] Z. Tong, Y. Song, J. Wang, and L. Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. In *arXiv:2203.12602*, 2022. 3
- [75] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou. Training data-efficient image transformers distillation through attention. In *ICML*, 2021. 6
- [76] T.D. Truong, C.N. Duong, and A.H. Pham. The right to talk: An audio-visual transformer approach. In *CVPR*, pages 1105–1114, 2021. 1
- [77] P. Voigtlaender, J. Luiten, P.H.S. Torr, and B. Leibe. Siam r-cnn: Visual tracking by re-detection. In *CVPR*, pages 6578–6588, 2020. 6
- [78] C. Vondrick, H. Pirsiavash, and A. Torralba. Generating videos with scene dynamics. In *NeurIPS*, 2016. 3
- [79] N. Wang, W. Zhou, J. Wang, and H. Li. Transformer meets tracker: Exploiting temporal context for robust visual tracking. In *ICCV*, 2021. 6
- [80] Q. Wang, L. Zhang, and L. Bertinetto. Fast online object tracking and segmentation: A unifying approach. In *CVPR*, 2019. 7
- [81] X. Wang, C. Li, B. Luo, and J. Tang. Sint++: Robust visual tracking via adversarial positive instance generation. In *CVPR*, pages 4864–4873, 2018. 2
- [82] X. Wang, X. Shu, Z. Zhang, B. Jiang, Y. Wang, Y. Tian, and F. Wu. Towards more flexible and accurate object tracking with natural language: Algorithms and benchmark. In *CVPR*, 2021. 6, 10
- [83] Z. Wang, J. Xu, L. Liu, F. Zhu, and L. Shao. Ranet: Ranking attention network for fast video object segmentation. In *ICCV*, 2019. 7
- [84] Z. Wang, H. Zhao, Y.L. Li, S. Wang, P. Torr, and L. Bertinetto. Do different tracking tasks require different appearance models? In *NeurIPS*, 2021. 7
- [85] Q. Wu and A.B. Chan. Meta-graph adaptation for visual object tracking. In *ICME*, 2021. 2
- [86] Q. Wu, J. Wan, and A. B. Chan. Progressive unsupervised learning for visual object tracking. In *CVPR*, pages 2993–3002, 2021. 2, 4
- [87] Q. Wu, Y. Yan, Y. Liang, Y. Liu, and H. Wang. Dsnet: Deep and shallow feature learning for efficient visual tracking. In *ACCV*, pages 119–134, 2018. 2
- [88] Y. Wu, J. Lim, and M.-H. Yang. Object tracking benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(9):1834–1848, 2015. 7
- [89] Z. Wu, Y. Xiong, and S. Yu. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, pages 3733–3742, 2018. 3
- [90] F. Xie, C. Wang, G. Wang, Y. Cao, W. Yang, and W. Zeng. Correlation-aware deep tracking. In *CVPR*, 2022. 6
- [91] Zhenda Xie, Yutong Lin, Zheng Zhang, Yue Cao, Stephen Lin, and Han Hu. Propagate yourself: Exploring pixel-level consistency for unsupervised visual representation learning. In *CVPR*, pages 16684–16693, 2021. 3
- [92] Z. Xie, Z. Zhang, Y. Cao, Y. Lin, J. Bao, and Z. Yao. Sim-mim: A simple framework for masked image modeling. In *CVPR*, pages 9653–9663, 2022. 1, 3
- [93] N. Xu, L. Yang, Y. Fan, D. Yue, Y. Liang, J. Yang, and T. Huang. Youtube-vos: A large-scale video object segmentation benchmark. In *arXiv:1809.03327*, 2018. 5, 9
- [94] Bin Yan, Yi Jiang, Peize Sun, Dong Wang, Zehuan Yuan, Ping Luo, and Huchuan Lu. Towards grand unification of object tracking. In *ECCV*, 2022. 7
- [95] B. Yan, H. Peng, J. Fu, D. Wang, and H. Lu. Learning spatio-temporal transformer for visual tracking. In *ICCV*, pages 10448–10457, 2021. 6, 7
- [96] T. Yang and A. B. Chan. Learning dynamic memory networks for object tracking. In *ECCV*, pages 152–167, 2018. 2
- [97] T. Yang, P. Xu, and R. Hu. Roam: Recurrently optimizing tracking model. In *CVPR*, pages 6718–6727, 2020. 2
- [98] Z. Yang, Y. Wei, and Y. Yang. Collaborative video object segmentation by foreground-background integration. In *ECCV*, 2020. 7
- [99] Z. Yang, Y. Wei, and Y. Yang. Associating objects with transformers for video object segmentation. In *NeurIPS*, pages 2491–2502, 2021. 2, 5
- [100] B. Ye, H. Chang, B. Ma, and S. Shan. Joint feature learning and relation modeling for tracking: A one-stream framework. In *ECCV*, pages 341–357, 2022. 1, 2, 3, 5, 6, 7, 10, 12
- [101] Dawei Zhang, Yanwei Fu, and Zhonglong Zheng. Uast: Uncertainty-aware siamese tracking. In *ICML*, 2022. 6
- [102] L. Zhang, A. Gonzalez-Garcia, J. Weijer, M. Danelljan, and F. Khan. Learning the model update for siamese trackers. In *ICCV*, 2019. 2
- [103] Richard Zhang, Phillip Isola, and Alexei A Efros. Image colorization. In *ECCV*, 2016. 3
- [104] Y. Zhang, Z. Wu, H. Peng, and S. Lin. A transductive approach for video object segmentation. In *CVPR*, 2020. 7
- [105] Z. Zhang, Y. Liu, X. Wang, B. Li, and W. Hu. Learn to match: Automatic matching network design for visual tracking. In *ICCV*, pages 19339–13348, 2021. 6
- [106] Z. Zhang and H. Peng. Deeper and wider siamese networks for real-time visual tracking. In *CVPR*, 2017. 2
- [107] Z. Zhang, H. Peng, J. Fu, B. Li, and W. Hu. Ocean: Object-aware anchor-free tracking. In *ECCV*, 2020. 6, 7, 10