

The Wisdom of Crowds: Temporal Progressive Attention for Early Action Prediction

Alexandros Stergiou^{1,2,*} Dima Damen³

¹Vrije University of Brussels, Belgium ²imec, Belgium ³University of Bristol, UK

Abstract

Early action prediction deals with inferring the ongoing action from partially-observed videos, typically at the outset of the video. We propose a bottleneck-based attention model that captures the evolution of the action, through progressive sampling over fine-to-coarse scales. Our proposed **Temporal Progressive (TemPr)** model is composed of multiple attention towers, one for each scale. The predicted action label is based on the collective agreement considering confidences of these towers. Extensive experiments over four video datasets showcase state-of-the-art performance on the task of Early Action Prediction across a range of encoder architectures. We demonstrate the effectiveness and consistency of TemPr through detailed ablations.[†]

1. Introduction

Early action prediction (EAP) is the task of inferring the action label corresponding to a given video, from only partially observing the start of that video. Interest in EAP has increased in recent years due to both the ever-growing number of videos recorded and the requirement of processing them with minimal latency. Motivated by the advances in action recognition [6, 57], where the entire video is used to recognize the action label, recent EAP methods [3, 15, 34, 45, 60] distill the knowledge from these recognition models to learn from the observed segments. Despite promising results, the information that can be extracted from partial and full videos is inevitably different. We instead focus on modeling the observed partial video better.

Several neurophysiological studies [11, 29] have suggested that humans understand actions in a predictive and not reactive manner. This has resulted in the *direct matching hypothesis* [18, 46] where, actions are believed to be perceived through common patterns. Encountering any of these patterns prompts the expectation of specific action(s), even before the action is completed. Although the early pre-

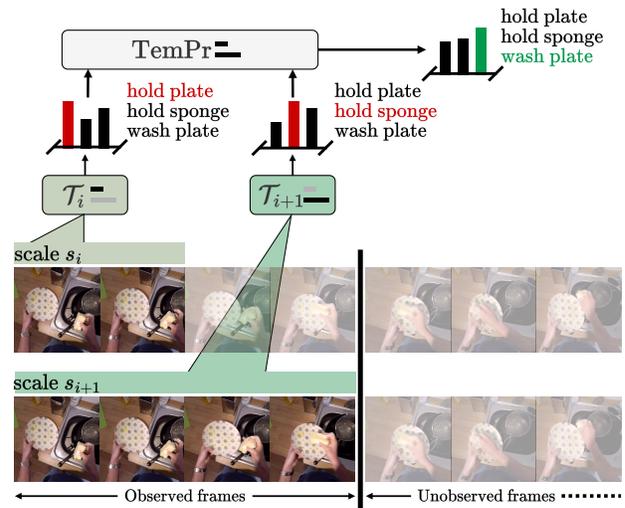


Figure 1. **Early action prediction with TemPr** involves the use of multiple scales for extracting features over partially observed videos. Encoded spatio-temporal features are attended by distinct transformer towers (\mathcal{T}) at each scale. We visualize two scales, where the fine scale \mathcal{T}_i predicts ‘hold plate’, and the coarse scale \mathcal{T}_{i+1} predicts ‘hold sponge’. Informative cues from both scales are combined for early prediction of the action ‘wash plate’.

diction of actions is an inherent part of human cognition, the task remains challenging for computational modeling.

Motivated by the direct matching hypothesis, we propose a Temporally Progressive (TemPr) approach to modeling partially observed videos. Inspired by multi-scale representations in images [7, 69] and video [27, 62], we represent the observed video by a set of sub-sequences of temporally increasing lengths as in Figure 1, which we refer to as scales. TemPr uses distinct transformer towers over each video scale. These utilize a shared latent-bottleneck for cross-attention [28, 37], followed by a stack of self-attention blocks to concurrently encode and aggregate the input. From tower outputs, a shared classifier produces label predictions for each scale. Labels are aggregated based on their collective similarity and individual confidences.

In summary, our contributions are as follows: (i) We

*Work carried out while A. Stergiou was at University of Bristol

[†]Code is available at: <https://tinyurl.com/temprog>

propose a progressive fine-to-coarse temporal sampling approach for EAP. (ii) We use transformer towers over sampled scales to capture discriminative representations and adaptively aggregate tower predictions, based on their confidence and collective agreement. (iii) We evaluate the effectiveness of our approach over four video datasets: UCF-101 [53], EPIC-KITCHENS [8], NTU-RGB [51] and Something-Something (sub-21 & v2) [21], consistently outperforming prior work.

2. Related Work

The task of EAP is related to but distinctly different from the tasks of action recognition and action anticipation. EAP predicts the *ongoing* action label, partially observed. In contrast, recognition assumes the *completed* action has been fully observed, while anticipation forecasts potential *upcoming* actions, seconds before the action starts. We first review prior EAP approaches, before relating our method to those used for other video understanding tasks.

Early action prediction: Most of the early attempts have focused on the probabilistic modeling of partially observed videos [4, 36, 38, 39, 47]. For example, Ryoo *et al.* [47] used a bag-of-words approach to model feature distributions over multiple partially observed videos. Later approaches aimed to overcome errors where large appearance variations occur, by either sparse coded feature bases [4] or through a scoring function [31, 33], combining prior knowledge and the sequential order of frames. Lan *et al.* [36] studied the representation of movements within the partially observed video, using a hierarchical structure.

More recent methods [3, 15, 25, 26, 32, 60, 64, 67, 71] have used learned-features. Specifically, knowledge distillation [24, 44] has been used to transfer class knowledge from the complete videos to the corresponding partial videos. This was achieved using Long Short-Term Memory (LSTM) models [26, 43, 60] and teacher-student frameworks [3, 15, 60]. Other methods are based on recurrent architectures with additional memory cells [32] for matching similar characteristics between the full and partial videos. Xu *et al.* [67] proposed a conditional generative adversarial network to generate feature representations for the entire video, from the partially observed video. Approaches have also focused on the propagation of residual features [71] or exploration with graph convolutions through relation reasoning [63, 64]. Foo *et al.* [16] proposed specializing features during training into instance-specific and general features. Instance-specific features are learned from a subset of videos focusing on subtle cues, while general features are learned from the entire dataset.

In contrast, we hypothesize that it is more beneficial to represent the partial video progressively. Our method is based on sampling at varying-length scales from the observed video to understand the temporal progression of ac-

tions. We show that aggregating these predictors can lead to notable improvements in accuracy. To our knowledge, we are the first to study EAP in this progressive manner.

Multi-scale representations for other video understanding tasks. The usage of scales, i.e. sequences of varying lengths or sampling at differing rates, is common in other video understanding task. For action recognition, video scales have been primarily used as a sampling method for either relational reasoning [14, 50, 73] or to select the most salient scale(s) as input to the network [42, 65, 72]. Xu *et al.* [66] proposed the Long Short-Term Transformer, an encoder-decoder for relating current actions with their long-term context. In action anticipation, methods utilize different scales to combine features from video snippets and anticipate one or more upcoming actions [17, 20]. Different from these tasks, and based on the fact that informative parts of partially observed videos do not have fixed lengths, we propose to utilize progressive video scales, which capture fine-to-coarse representations making them more suitable for partially observed videos.

Attention for video tasks. Attention-based video methods [59, 61] have initially been used as part of spatio-temporal CNNs [6, 57]. The recent introduction of Vision Transformer [10] has inspired subsequent works on action recognition by either focusing on how spatio-temporal information can be processed [1, 2] or architectural optimizations for spatio-temporal data [12, 40, 48, 68, 70]. Motivated by the recent advances of transformers for action recognition, we combine multiple transformer towers in TemPr.

3. Our Approach

In this section, we overview our TemPr model (shown in Figure 2). We first introduce our prime contribution of progressive scales for sampling from the observed video in Section 3.2. Each scale corresponds to an attention tower, which captures the progression of the action, and predicts the ongoing action, as explained in Section 3.3. Multiple scales/towers are then combined for a final prediction by an aggregation function, detailed in Section 3.4.

3.1. EAP: Problem Definition

We follow the standard definition of the EAP task from recent works [3, 64, 67, 71]. We denote the full video with T frames as $\mathbf{v}_{\{1, \dots, T\}}$. We define the *observation ratio* $0 < \rho < 1$ as the proportion of frames observed. EAP assumes $0 < \rho$, i.e. at least one frame of the video depicting the action has been observed, and $\rho < 1$, i.e. part of the video remains unobserved. Accordingly, $T_\rho = \lceil \rho \cdot T \rceil$ is the number of observed frames. In EAP, the prediction of the ongoing action label y conveyed in the full video $\mathbf{v}_{\{1, \dots, T\}}$ is attempted from only the observed T_ρ frames.

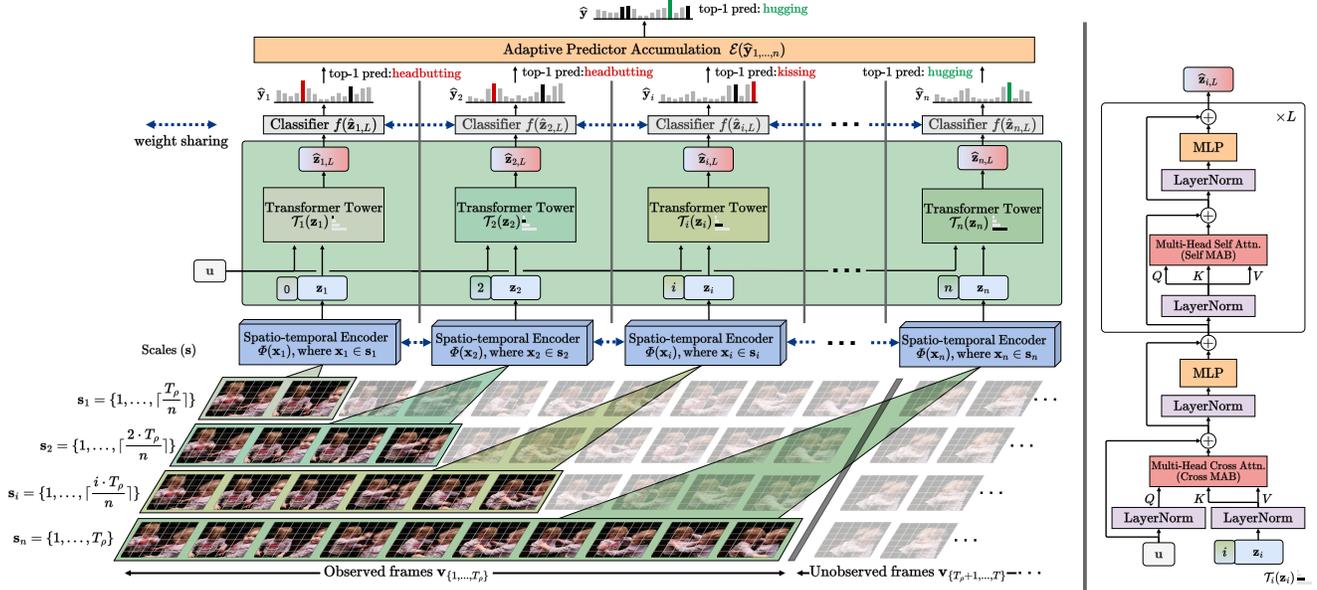


Figure 2. (Left) **TemPr architecture**. Features are extracted over each input \mathbf{x}_i sampled from video scale \mathbf{s}_i , and combined with scale and spatio-temporal positional encodings. The encoded features \mathbf{z}_i are passed to attention towers \mathcal{T}_i which output tensors $\hat{\mathbf{z}}_{i,L}$ in the latent space. Shared-weight classifier $f(\cdot)$ is applied to every tower output to make per-scale predictions. These predictions are aggregated by aggregation function $\mathcal{E}(\cdot)$, for early action prediction over the observed frames. (Right) **Attention Tower**. Each utilizes pre-norm and a shared latent array \mathbf{u} for the cross-attention block (Cross MAB). This is followed by a stack of L self-attention blocks (Self MAB).

3.2. Progressive Video Scales

Given the partial observation of the action, we speculate that the sampling strategy is critical for capturing distinctive representations of the ongoing action. This is different from the sampling typically utilized in action recognition, where the video is uniformly split into equally-sized segments [58]. Equal-sized segments, in partially observed videos, can miss the discriminative action pattern when this pattern spans across segments. We thus propose to sample at multiple scales within the observed video, which we refer to as *progressive* sampling.

Given the partially observed video of T_p frames, we examine the ongoing action over n scales $\mathbf{s}_{\{1,\dots,n\}}$. Each scale \mathbf{s}_i has a larger temporal extent to sample from than \mathbf{s}_{i-1} . We represent each scale \mathbf{s}_i as:

$$\mathbf{s}_i = \{1, \dots, T_{\mathbf{s}_i}\}; \quad T_{\mathbf{s}_i} = \lceil \frac{i}{n} \cdot T_p \rceil \quad \forall i \in \mathbf{N} = \{1, \dots, n\} \quad (1)$$

Over each scale, we sample F frames randomly to capture a progressive fine-to-coarse representation. Considering the variable input length per scale, sampling a fixed number of frames F , is required to standardize the encoder inputs.

3.3. Temporal Progressive Attention Towers

We use a shared encoder $\Phi(\cdot)$ to extract features from the sampled frames, over the progressive scales. Corresponding to each scale \mathbf{s}_i , we define input volume \mathbf{x}_i of

size $3 \times F \times H \times W$, with F temporally ordered frames, H height and W width. We thus define $\mathbf{z}_i = \Phi(\mathbf{x}_i)$ to be the per-scale, multi-dimensional spatio-temporal encoded feature volume, of size $C \times t \times h \times w$. Given the scales' spatio-temporal features $\mathbf{z}_1, \dots, \mathbf{z}_n$, we reshape these to $C \times (thw)$, and concatenate Fourier Positional Embeddings (PE) of size $n \times (thw)$ to encode each scale and space-time position. Features \mathbf{z}_i form the input to attention tower \mathcal{T}_i .

We attend each scale's features using tower \mathcal{T}_i , so that $\hat{\mathbf{z}}_i = \mathcal{T}_i(\mathbf{z}_i)$, where $\hat{\mathbf{z}}_i$ is the feature volume after attending input volume \mathbf{z}_i over the transformer blocks. Motivated by the recent architectural approaches for dealing with the quadratic scaling of complexity in transformers [28, 37], each tower uses two attention components consisting of one cross-attention bottleneck block and a stack of self-attention blocks as shown in Figure 2 (right). Towers are indexed by $i \in \mathbf{N}$ and attention blocks, per tower, are indexed by $j \in \{0, \dots, L\}$. We describe these components next.

Cross Multi-Head Attention Block (Cross MAB), employs a latent array \mathbf{u} of $C \times d$ size ($d \ll thw$). This latent array alongside \mathbf{z}_i are used to create the asymmetric query-key-value (QKV) attention function in which $\mathbf{Q} \in \mathbb{R}^{C \times d}$, $\mathbf{K} \in \mathbb{R}^{C \times (thw)}$, $\mathbf{V} \in \mathbb{R}^{C \times (thw)}$. The Cross MAB block consists of Multi-Head Cross Attention (MCA), Layer Normalization (LN), and Multilayer Perceptron (MLP) modules:

$$\begin{aligned} \hat{\mathbf{z}}_{i,0} &= MLP(LN(\mathbf{h}_{i,0})) + \mathbf{h}_{i,0}, \text{ where} \\ \mathbf{h}_{i,0} &= MCA(LN(\mathbf{u}), LN(\mathbf{z}_i)) + \mathbf{u} \quad \forall i \in \mathbf{N} \end{aligned} \quad (2)$$

in which, the MCA computes the dot-product asymmetric attention of tensors \mathbf{u} and \mathbf{z}_i .

By exploiting the Cross MAB [28] bottleneck, the transformer towers are significantly more efficient than a deep stack of self-attention blocks. The use of a parameterizable size latent vector can benefit the creation of performance-balanced models, minimizing feature redundancies.

Stacked Self-Attention Blocks (Self MAB), correspond to a stack of L transformer blocks [10], symmetrically attending to tensors $\hat{\mathbf{z}}_{i,j} \forall j \in \{0, \dots, L-1\}$. Including Multi-Head Self Attention (MSA), the block is denoted as:

$$\begin{aligned} \hat{\mathbf{z}}_{i,j} &= MLP(LN(\mathbf{h}_{i,j})) + \mathbf{h}_{i,j}, \text{ where} \\ \mathbf{h}_{i,j} &= MSA(LN(\hat{\mathbf{z}}_{i,j-1})) + \hat{\mathbf{z}}_{i,j-1} \forall i \in \mathbf{N}, j \in \{1, \dots, L\} \end{aligned} \quad (3)$$

Attention tower predictors. Towers additionally include a linear classifier $\hat{\mathbf{y}}_i = f(\hat{\mathbf{z}}_{i,L})$ that maps the output $\hat{\mathbf{z}}_{i,L}$ to $\hat{\mathbf{y}}_i$ class predictions. As features $\hat{\mathbf{z}}_{i,L}$ are bound to scale s_i , towers cannot relate features across scales, which limits their modeling capabilities. We thus share classifier weights across scales to establish a joint feature space.

Predictions from the n attention towers are thus obtained. We describe our proposed aggregation approach next.

3.4. Aggregation Function for EAP

We wish to accumulate class predictions from the individual fine-to-coarse scales into an overall EAP for the observed T_p frames.

We introduce an aggregation function $\mathcal{E}(\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_n)$ for accumulating tower predictions. The function is formulated based on the agreement between predictions and the individual towers' confidence in the produced prediction.

Predictor agreement. We trust that predictions with a high degree of resemblance, in terms of their class probability distribution, can reduce the uncertainty of individual predictors. We utilize Exponential Inverse Coefficient Weighting (eICW) [54] for the weighted aggregation of probabilities $\hat{\mathbf{y}}_i$ per scale, based on their similarity to the mean probability distribution $\bar{\mathbf{y}}$:

$$\mathcal{E}_{eICW}(\hat{\mathbf{y}}_i, \bar{\mathbf{y}}) = \frac{e^{DSC(\hat{\mathbf{y}}_i, \bar{\mathbf{y}})^{-1}}}{\sum_{k \in \mathbf{N}} e^{DSC(\hat{\mathbf{y}}_k, \bar{\mathbf{y}})^{-1}}} \cdot \hat{\mathbf{y}}_i \quad (4)$$

in which $DSC(\cdot)$ is the Dice-Sørensen coefficient [9] between class probabilities $\hat{\mathbf{y}}_i$ and mean probabilities $\bar{\mathbf{y}}$.

Predictor confidence. Aggregation is performed based on the sharpness of the probability distribution. We calculate the exponential maximum (i.e. softmax) across all predictions. Predictions with high class probability for a single or a small set of classes are weighted higher:

$$\mathcal{E}_{eM}(\hat{\mathbf{y}}_i) = \frac{e^{\hat{\mathbf{y}}_i}}{\sum_{k \in \mathbf{N}} e^{\hat{\mathbf{y}}_k}} \cdot \hat{\mathbf{y}}_i \quad (5)$$

A combination of the two strategies is used for the final adaptive predictor aggregation function $\mathcal{E}(\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_n)$. As in [54], we use a parameter $0 \leq \beta \leq 1$, which we learn during training, to determine the proportion of each method:

$$\mathcal{E}(\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_n) = \sum_{i \in \mathbf{N}} \beta \cdot \mathcal{E}_{eICW}(\hat{\mathbf{y}}_i, \bar{\mathbf{y}}) + (1 - \beta) \cdot \mathcal{E}_{eM}(\hat{\mathbf{y}}_i) \quad (6)$$

We refer to this aggregation function as our proposed *adaptive* aggregation function for attention tower predictions.

During training, we use the adaptive probability distribution from $\mathcal{E}(\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_n)$ to calculate the divergence from the target one-hot categorical distribution for class vector \mathbf{y} . In inference, the arg max class is used as the EAP label.

In summary, our proposed method combines progressive scales of the observed video, individual attention towers with shared classifier weights, and an aggregation function that backpropagates through all individual attention towers. We evaluate our method next.

4. Experiments

The datasets used, alongside implementation and training scheme details, are explained in Section 4.1. We include state-of-the-art model comparisons in Section 4.2 followed by ablation studies in Section 4.3.

4.1. Datasets and Implementation Details

Datasets We report our method's performance over a diverse set of video datasets previously used for EAP. *UCF-101* [53] consists of 101 classes and 13K videos depicting various types of actions such as human-object interactions, human-human interactions, playing musical instruments, and sports. *Something-Something* (SSv1/SSsub21/SSv2) [21] is a collection of 100K (SSv1) & 220K (SSv2) videos of 174 fine-grained human-object action and interaction categories. The v1 of the dataset also includes a 21-action categories subset (SSsub21) of 11K videos used previously by [63, 64] for EAP. We report on this subset, for direct comparisons and v2 for large-scale benchmarking. *EPIC-KITCHENS-100* (EK-100) [8] contains unscripted egocentric actions and activities across 45 kitchen environments. Labels are composed of 97 verb classes, 300 noun classes, and 4025 action classes of combined nouns and verbs. We also use the RGB-only version of *NTU RGB+D* [51], as in [34, 38], containing 60 action classes and 57K videos of daily human actions.

Previous EAP works [3, 4, 32, 34, 49, 60, 63, 64] have evaluated their performance over smaller datasets ($< 100K$ videos) that are only partially indicative of the approaches' generalizability. We thus set new EAP baselines by evaluating on two large-scale datasets: the temporally challenging SSv2 as well as EK-100.

Table 1. **Top-1 accuracies (%) of action prediction methods on UCF-101 over different observation ratios (ρ)**. Methods are grouped w.r.t. the backbone used. We report **TemPr** results on 5 backbones. The best results per ρ are in **bold** and second best are underlined.

Method	Backbone	dim	Observation ratios (ρ)								
			0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
RGN-KF [71]	Inception [55]	2D	83.3	85.2	87.8	90.6	91.5	92.3	92.0	93.0	92.9
GGNN [64]			82.4	85.6	89.0	-	91.3	-	92.4	-	93.0
TS (2xL) [60]			83.3	87.1	88.9	89.8	90.9	91.0	91.3	91.2	91.3
AAPNet [35]	C3D [56]	3D	59.9	80.4	86.8	86.5	86.9	88.3	88.3	89.9	90.9
MSSC [4]	ResNet-18	2D [23]	34.1	53.8	58.3	57.6	62.6	61.9	63.5	64.3	62.7
MTSSVM [33]			40.1	72.8	80.0	82.2	82.4	83.2	83.4	83.6	83.7
DeepSCN [34]			45.0	77.7	83.0	85.4	85.8	86.7	87.1	87.4	87.5
mem-LSTM [32]			51.0	81.0	85.7	85.8	88.4	88.6	89.1	89.4	89.7
MSRNN [26]			68.0	87.2	88.2	88.8	89.2	89.7	89.9	90.3	90.4
GGNN [64]			75.9	81.7	87.8	-	88.7	-	89.4	-	90.2
TemPr $\underline{\underline{=}}$ (ours)			3D [22]	84.3	90.2	90.4	90.9	91.2	91.8	92.1	92.3
AA-GAN [19]	ResNet-50	2D [23]	-	84.2	-	-	85.6	-	-	-	-
GGNN [64]			84.1	88.5	89.8	-	90.9	-	91.4	-	91.8
TS+JVS+JCC+JFIP [15]			-	85.8	-	-	-	-	-	-	-
TemPr $\underline{\underline{=}}$ (ours)	3D [22]	84.8	90.5	91.2	91.8	91.9	92.2	92.3	92.4	92.6	
DBDNet [43]	ResNeXt101 [22]	3D	82.7	86.6	88.3	89.7	90.6	91.2	91.7	91.9	92.0
IGGNN [63]			80.2	-	89.8	-	92.9	-	94.1	-	94.4
ERA [16]			89.1	-	92.4	-	94.3	-	<u>95.4</u>	-	95.7
TemPr $\underline{\underline{=}}$ (ours)			85.7	91.4	92.1	92.7	93.5	<u>93.9</u>	94.4	94.6	94.9
TemPr $\underline{\underline{=}}$ (ours)	X3D _M [13]	3D	87.9	<u>93.4</u>	<u>94.5</u>	<u>94.8</u>	<u>95.1</u>	95.2	95.6	<u>96.4</u>	96.3
TemPr $\underline{\underline{=}}$ (ours)	MoViNet-A4 [30]	3D	<u>88.6</u>	93.5	94.9	94.9	95.4	95.2	95.3	96.6	<u>96.2</u>
TemPr $\underline{\underline{=}}$	MoViNet-A4	3D	87.3	93.1	94.9	94.6	95.2	94.9	94.6	95.1	95.0
TemPr $\underline{\underline{=}}$			85.6	92.9	93.6	94.5	94.4	94.2	94.2	94.6	94.8
TemPr $\underline{\underline{=}}$			85.2	92.1	92.5	92.9	93.3	93.7	93.5	93.8	93.7

Model settings. We evaluate our model over four scales $n = \{1, 2, 3, 4\}$. We use the concise visual notation: $-$, $\underline{\underline{=}}$, $\underline{\underline{=}}$, $\underline{\underline{=}}$, $\underline{\underline{=}}$ to refer to these 4 configurations. Except during ablations, we follow model configurations similar to [28, 37] for each attention tower ($L = 8$, $d = 256$, $H_C = 4$, $H_S = 8$)[‡]. We sample $F = 16$ frames for each scale[§].

Overall, we employ four encoder architectures. MoViNet-A4 [30] is used for UCF-101, SSsub21 and NTU-RGB in Section 4.2 due to its efficiency and high accuracy on action recognition. A 3D ResNet-18 with TemPr $\underline{\underline{=}}$ is used to compare against models with the same feature encoder in Section 4.2 and for the ablation studies in Section 4.3. We additionally experiment with the widely used encoder networks, SlowFast-R50 [14] for EK-100 and (video) Swin-B [40] on SSv2. All convolutional encoders are pre-trained on Kinetics-700 [52] and then trained on each dataset over the full videos. Swin-B is initialized with the official weights pre-trained on Kinetics-600 [5].

[‡] L : number of self-attention layers, d : size of the latent bottleneck, H_C and H_S : numbers of cross and self-attention heads respectively.

[§]We use adaptive average pooling for down-scaling encoder output features \mathbf{z}_i across datasets to a fixed size of $t=16$, $h=4$, and $w=4$

Training scheme. For UCF-101, EK-100, and NTU-RGB, we process the videos by scaling the height to 384px and taking a center crop to size 384×384 px followed by a random crop of 224×224 px. Because of SSsub21’s low frame resolution, we scale the input frames to 100×176 px. We initialize β with 0.5 and train for 60 epochs with $1e^{-2}$ base learning rate for TemPr and $1e^{-3}$ for β . Both learning rates are reduced on epochs $\{14, 32, 44\}$ by $1e^{-1}$. We use batch sizes of 32 for UCF-101, EK-100, NTU-RGB & SSv2 and 64 for SSsub21 with AdamW & $1e^{-5}$ weight decay.

4.2. Comparative Results

UCF-101. For a fair comparison to prior methods, we structure our results based on the feature encoder. In the top half of Table 1, we demonstrate that our TemPr $\underline{\underline{=}}$ model consistently outperforms all other methods with the same ResNet-18 encoder [4, 26, 32–34, 64], for every observation ratio. Across our tests, the largest improvements are observed in small ratios in which, we achieve +8.4% improvement for $\rho = 0.1$, +3.0% for $\rho = 0.2$ and +2.2% for $\rho = 0.3$, compared to the previous top-performing models.

We also outperform prior works [15, 19, 43, 63, 64] on the

Table 2. **Top-1 accuracy (%) of EAP** over different observation ratios (ρ).

(a) NTU-RGB.							(b) SSsub21.						(c) SSv2.					
Method	Observation ratios (ρ)						Method	Observation ratios (ρ)						Method	Obs. ratios (ρ)			
	0.1	0.2	0.3	0.5	0.7	0.9		0.1	0.2	0.3	0.5	0.7	0.9		0.1	0.3	0.5	0.7
RankLSTM [41]	11.5	16.5	25.7	48.0	61.0	66.1	mem-LSTM [32]	14.9	17.2	18.1	20.4	23.2	24.5	Baseline (Inference)	6.9	17.6	28.9	36.0
DeepSCN [34]	16.8	21.5	30.6	48.8	58.2	60.0	MS-LSTM [49]	16.9	16.6	16.8	16.7	16.9	17.1	Baseline (Fine-tuned)	14.4	23.5	31.1	39.6
MSRNN [26]	15.2	20.3	29.5	51.6	63.9	68.9	MSRNN [49]	20.1	20.5	21.1	22.5	24.0	27.1	TemPr $\underline{\text{L}}$ (ours)	20.5	28.6	41.2	47.1
TS (2xL) [60]	27.8	35.8	46.3	67.4	77.6	81.5	GGN [64]	21.2	21.5	23.3	27.4	30.2	30.5					
TemPr $\underline{\text{L}}$ (ours)	29.3	38.7	50.2	70.1	78.8	84.2	IGGN [63]	22.6	-	25.0	28.3	32.2	34.1					
							TemPr $\underline{\text{L}}$ (ours)	28.4	34.8	37.9	41.3	45.8	48.6					

(d) EK-100.																				
Method	Verb						Noun						Action							
	Observation ratios (ρ)																			
	0.1	0.2	0.3	0.5	0.7	0.9	0.1	0.2	0.3	0.5	0.7	0.9	0.1	0.2	0.3	0.5	0.7	0.9		
Baseline (Inference)	17.3	19.7	27.0	48.7	60.5	64.2	19.5	21.7	25.3	38.5	46.7	49.1	5.4	7.6	11.1	24.3	34.1	37.6		
Baseline (Fine-tuned)	20.6	21.8	29.4	49.8	61.3	64.3	21.3	24.2	27.6	39.4	47.3	49.1	6.9	9.1	12.8	25.5	34.9	37.5		
TemPr $\underline{\text{L}}$ (ours)	21.4	22.5	34.6	54.2	63.8	67.0	22.8	25.5	32.3	43.4	49.2	53.5	7.4	9.8	15.4	28.9	37.3	40.8		

same backbone for every ρ . Our method does not outperform [16] on the 48M parameters ResNeXt101 backbone. However, using the more efficient MoViNet-A4 or X3D_M networks with 5M and 4M parameters respectively, we outperform [16] in all but $\rho = 0.1$. We get best performance of TemPr $\underline{\text{L}}$ when using MoViNet-A4; e.g. at $\rho = 0.3$ we outperform all prior work by 2.5%. For $\rho = 0.1$, we speculate that methods like [16] benefit from specializing to subtle differences when only a handful of frames are observed. The final three rows of Table 1 present results across for $n = 1, 2$ and 3. Results steadily increase, across observation ratios as more scales are incorporated in TemPr. Further results are available in §S1 in Supplementary Material.

NTU-RGB. Results on NTU-RGB are presented in Table 2a. Compared to the state-of-the-art, our TemPr $\underline{\text{L}}$ consistently outperforms other models across the six observation ratios used. We observe the largest improvement in accuracy over [60] at $\rho = 0.3$ with 3.9%. For smaller observation ratios, accuracy increases by 1.5% and 2.9% for $\rho = 0.1$ and $\rho = 0.2$, respectively.

Something-Something (sub21). Table 2b demonstrates the SSsub21 class-averaged accuracy, across observation ratios as in [63, 64]. Our proposed TemPr $\underline{\text{L}}$ surpasses state-of-the-art models [63, 64] with a significant improvement over all observation ratios. Compared to the previous top-performing model per observation ratio, accuracy increases include 5.8% at $\rho = 0.1$, 13.3% at $\rho = 0.2$, 13.6% at $\rho = 0.7$, and 14.5% at $\rho = 0.9$.

Something-Something (SSv2). Table 2c shows results on SSv2 per observation ratio with video Swin-B, which achieves 66.3% when evaluated on full videos (i.e. $\rho = 1.0$)[‡]. We note the significant drop in performance when evaluated on partially-observed videos. Even when $\rho =$

0.7, the model can only achieve 36.0% top-1 accuracy. The improvement remains modest when the classifier is fine-tuned. On average, TemPr $\underline{\text{L}}$ outperforms the inference-only model by 12.0% and the fine-tuned model by 7.2%. Improvements are also evident across ρ . This not only demonstrates the benefits of our proposed TemPr model for EAP, but also the distinction between the tasks of action classification and EAP, and thus the need for EAP-specific models. **EPIC-KITCHENS-100 (EK-100).** We also investigate EAP on EK-100. We believe that a challenging part of EK-100 is the inclusion of fine-grained verb labels. For example, the class ‘hold’ is easily confused with partially-observed videos of classes ‘put’, ‘throw’, ‘insert’ or ‘stack’. These classes start with objects being held before the action is initiated. We are the first to use EK-100 as a benchmark for EAP. As in SSv2, we report inference-only and classifier fine-tuned models alongside TemPr $\underline{\text{L}}$.

Table 2d demonstrates the performance per observation ratio. TemPr $\underline{\text{L}}$ outperforms the baselines and showcases that EK-100 is more challenging than all other benchmarks when focusing on action performance - 28.9% for $\rho = 0.5$ compared to 95.4%, 70.1% and 41.2% for UCF-101, NTU-RGB, and SSv2. We note that EAP is higher for noun classes in smaller ρ while classifying verbs becomes easier for larger ρ . This highlights that actions, which require correct prediction of the verb and the noun, are challenging to be predicted in cases where very few frames are observed.

4.3. Ablation Studies and Qualitative Results

In this section we conduct ablation studies on UCF-101 reporting accuracy over different observation ratios. Unless specified, we use the ResNet-18 backbone. Computations and memory use are reported solely for TemPr, without the encoder, to demonstrate the differences clearer.

Video scales strategy. Different strategies can be used for

[‡]We note that the difference from the reported 69.6% accuracy in [40] is due to our use of 16 frames instead of the reported 32 frames as input.

Table 3. **Ablation studies on UCF-101** with TemPr $\underline{\text{L}}$ across obs. ratios. We use \spadesuit to denote softmax during training and \clubsuit for $\theta = \frac{1}{2n}$.

(a) Video Scales Strategy.					(b) Aggregation function.			(c) Weight sharing over attention towers and classifiers.			(d) Latent array (\mathbf{u}) sharing.					
Scale strategy	Observation ratios (ρ)				Aggregation	ρ		Weight sharing	ρ			\mathbf{u} shared	Mem. (GB)	ρ		
	0.2	0.4	0.6	0.8		0.2	0.4	MAB	$f(\cdot)$	0.2	0.4	0.6	\times	4.0	0.2	0.4
full \equiv	86.4	88.3	88.8	89.0	avg	89.5	90.1	\checkmark	\times	73.4	76.2	79.0	\times	3.0	90.2	91.0
equal \sphericalangle	83.7	84.6	86.3	87.1	softmax	87.8	89.4	\times	\times	84.7	85.8	87.3	\checkmark		90.2	90.9
random \equiv	88.8	89.7	90.2	90.6	top \spadesuit	84.6	87.5	\checkmark	\checkmark	89.2	90.0	90.7				
decreasing ∇	90.0	90.9	91.6	92.6	gate ($\theta=0.1$)	85.4	88.5	\times	\checkmark	90.2	90.9	91.8				
increasing $\underline{\text{L}}$	90.2	90.9	91.8	92.3	ICW	89.7	90.1									
Table 4. Video Scales Strategies on SSsub21 with TemPr $\underline{\text{L}}$.					weighted	88.5	89.0									
Scale strategy	Obs. ratios (ρ)				weighted (θ) \clubsuit	83.4	85.8									
	0.2	0.3	0.5	0.7	adaptive ($\mathcal{E}(\cdot)$)	90.2	90.9									
full \equiv	32.6	36.4	39.3	42.9	(e) CMAB replacements.											
equal \sphericalangle	29.8	34.5	37.2	41.8	MAB	ρ		Par. (M)	GfLOPs							
random \equiv	33.4	37.1	40.6	44.3		0.2	0.4									
decreasing ∇	35.2	38.3	40.7	45.2	Self	83.2	84.5	84.6	8.59							
increasing $\underline{\text{L}}$	34.8	37.9	41.3	45.8	Cross	90.2	90.9	23.0	1.47							

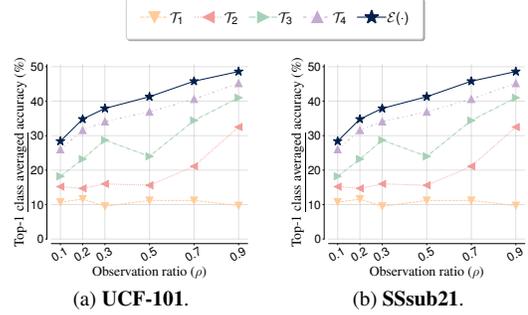


Figure 3. **Top-1 accuracy of each TemPr $\underline{\text{L}}$ tower T_i per ρ .**

selecting video scales. We compare our proposed temporal progressive sampling (Section 3.2) to other common strategies and potential baselines in Table 3a. In all settings, we keep $n = 4$ scales. The *full* strategy \equiv uses n scales of fixed length matching the entire observation video. In *equal* \sphericalangle , scales/segments have equal lengths as in [58]. The *random* strategy \equiv uses scales of random length. Finally, the *increasing* $\underline{\text{L}}$ and *decreasing* ∇ strategies utilize our proposed progressive approach, sampling the fine scale from either the start or the end of the observed video. Accuracy is consistently lower when scales are of the same length, either matching the observed video (*full*) or equally-sized (*equal*). This is in contrast to the success of this sampling approach for action recognition [58], further emphasizing the distinction between the two tasks. The use of progressive (*increasing* or *decreasing*) video scales exhibits an average +3.6% accuracy increase across ρ , compared to other sampling approaches. We note that no model component depends on the order of the scales, thus the performance over increasing or decreasing scales is expected to be similar.

In Table 4, we compare sampling strategies on SSsub21, as this dataset is more challenging temporally. We use TemPr $\underline{\text{L}}$ with MoViNet-A4. Similar to Table 3a, progressive (*increasing* $\underline{\text{L}}$ or *decreasing* ∇) scales is a better-suited strategy, with an average +2.4% accuracy increase over ρ . This emphasizes the need for fine-to-coarse sampling, independent of where the fine sample is taken from.

Prediction aggregation. Table 3b presents comparisons over different aggregation functions. In the case that the predictor with the highest confidence is chosen (top \spadesuit), we use softmax during training to ensure that gradients are propagated across the entire network. The largest drop in

performance is observed when using individual predictions (softmax, top, gate). Methods that are instead based on using all predictors uniformly by averaging them, or by weighting them with Inverse Covariance Weighting (ICW) improve the final predictions. A further +0.7% accuracy over ICW is observed by our adaptive approach with the combination of predictor agreement and confidences.

Weight sharing combinations. We consider the two model components that can share their weights across scales. The first is the multihead-attention blocks (MAB) and the second is their classifier layer. Table 3c shows that using individual classifier weights for each tower decreases performance. Classifier weight sharing improves performance.

Latent array (\mathbf{u}). Table 3d shows the effect on both performance and memory when sharing the Cross MAB latent array \mathbf{u} across attention towers. With marginal difference in accuracy, sharing \mathbf{u} increases efficiency with a significant reduction in memory. Thus, we share \mathbf{u} in all experiments.

CMAB replacements. We include ablations on the effect of cross/self-MAB in accuracy, compute and memory on Table 3e. We note that self-MAB-only towers significantly increase memory and computation costs.

Scale per Observation Ratio. We additionally plot the performance of individual predictors for both UCF-101 and SSsub21 in Figure 3 with respect to different observation ratios. As shown, datasets such as Something-Something that are less appearance-based can benefit more from the proposed aggregated progressive scales. Class accuracies across scales are presented in §S1. Overall, towers of smaller scales (T_1 $\underline{\text{L}}$ and T_2 $\underline{\text{L}}$) performed more favorably for classes that are distinguishable from the only first

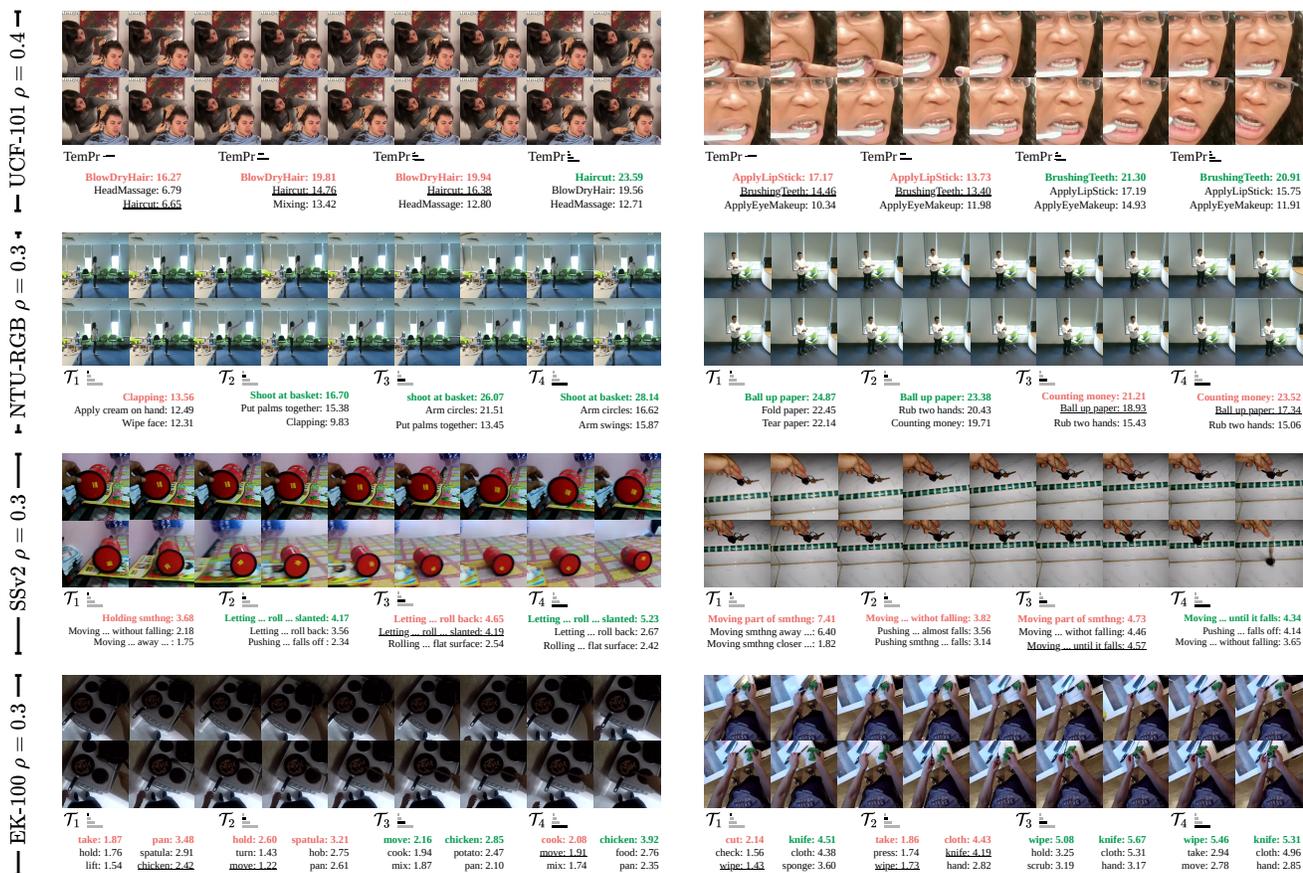


Figure 4. Examples from UCF-101, NTU-RGB, SSV2 and EK-100. Top 3 action label confidences are reported for either TemPr model or over individual tower predictors (\mathcal{T}_i). We show the 16 frames sampled per video. Green/red highlight correct/incorrect top 1 predictions, and we underline true label when in top-3. We show verb and noun predictions for EK-100. See additional examples in §S6.

few frames. In contrast, towers of larger scales (\mathcal{T}_3 and \mathcal{T}_4) were better suited for classes that the action become distinguishable with a larger part of the video observed.

Qualitative results. The first row of Figure 4 demonstrates UCF-101 instances where predictions differ across TemPr $\rho = 0.4$, $\rho = 0.3$, $\rho = 0.3$, $\rho = 0.3$. The increase in the number of scales allows the network to capture features that are more descriptive of the target action e.g. the two *BrushingTeeth* instances. In the first example, the subtle motion of *Hair Cutting* is only confidently predicted when the finest scale is incorporated in TemPr (comparing $\rho = 0.4$ to $\rho = 0.3$). In the following three rows of Figure 4, predictions from individual towers \mathcal{T}_1 , \mathcal{T}_2 , \mathcal{T}_3 and \mathcal{T}_4 are shown across NTU-RGB, SSV2, and EK-100. In the second row, fine scales benefit subtle motion e.g. in the *Ball up paper*. In the third row, coarse scales assist prediction as the end of the sequence changes the prediction to the correct class, e.g. *Moving something until it falls* in SSV2. In the fourth row, coarser scales are required to distinguish *taking cloth* from *wiping knife* in EK-100.

5. Conclusions

We have proposed to utilize progressive scales from partially observed videos for early action prediction. Based on these scales, we introduce a temporal progressive (TemPr) model consisting of bottleneck-based attention towers, in order to capture the progression of an action over multiple fine-to-coarse scales. We aggregate scale predictors considering the similarity in their probability distributions as well as their confidence. Extensive experiments over five encoders and four video datasets demonstrate the merits of TemPr. Additionally, we are the first to investigate the unique difficulties of EAP for large-scale datasets - evaluating EAP on SSV2 and EK-100. We hope that our approach of progressive, rather than single continual, scales can pave a new path for subsequent methods.

Acknowledgments. We use publicly available datasets. Research is funded by the United Nation’s End Violence Fund (iCOP 2.0) and EPSRC UMPIRE (EP/T004991/1). We utilized Bristol’s HPC Blue Crystal 4 facility.

References

- [1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. ViViT: A video vision transformer. In *International Conference on Computer Vision (ICCV)*, 2021. 2
- [2] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *International Conference on Machine Learning (ICML)*, 2021. 2
- [3] Yijun Cai, Haoxin Li, Jian-Fang Hu, and Wei-Shi Zheng. Action knowledge transfer for action prediction with partial videos. In *Association for the Advancement of Artificial Intelligence (AAAI)*, 2019. 1, 2, 4
- [4] Yu Cao, Daniel Barrett, Andrei Barbu, Siddharth Narayanaswamy, Haonan Yu, Aaron Michaux, Yuewei Lin, Sven Dickinson, Jeffrey Mark Siskind, and Song Wang. Recognize human activities from partially observed videos. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013. 2, 4, 5
- [5] Joao Carreira, Eric Noland, Andras Banki-Horvath, Chloe Hillier, and Andrew Zisserman. A short note about kinetics-600. *arXiv preprint arXiv:1808.01340*, 2018. 5
- [6] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? A new model and the Kinetics dataset. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2
- [7] Yunpeng Chen, Haoqi Fan, Bing Xu, Zhicheng Yan, Yanis Kalantidis, Marcus Rohrbach, Shuicheng Yan, and Jiashi Feng. Drop an octave: Reducing spatial redundancy in convolutional neural networks with octave convolution. In *International Conference on Computer Vision (ICCV)*, 2019. 1
- [8] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Evangelos Kazakos, Jian Ma, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Rescaling egocentric vision: Collection, pipeline and challenges for EPIC-KITCHENS-100. *International Journal of Computer Vision (IJCV)*, 130:33–55, 2022. 2, 4
- [9] Lee R Dice. Measures of the amount of ecologic association between species. *Ecology*, 26(3):297–302, 1945. 4
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations (ICLR)*, 2020. 2, 4
- [11] Luciano Fadiga, Leonardo Fogassi, Giovanni Pavesi, and Giacomo Rizzolatti. Motor facilitation during action observation: a magnetic stimulation study. *Journal of neurophysiology*, 73(6):2608–2611, 1995. 1
- [12] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *International Conference on Computer Vision (ICCV)*, 2021. 2
- [13] Christoph Feichtenhofer. X3D: Expanding architectures for efficient video recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 5
- [14] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. SlowFast networks for video recognition. In *International Conference on Computer Vision (ICCV)*, 2019. 2, 5
- [15] Basura Fernando and Samitha Herath. Anticipating human actions by correlating past with the future with jaccard similarity measures. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 5
- [16] Lin Geng Foo, Tianjiao Li, Hossein Rahmani, Qihong Ke, and Jun Liu. ERA: Expert retrieval and assembly for early action prediction. In *European Conference on Computer Vision (ECCV)*, 2022. 2, 5, 6
- [17] Antonino Furnari and Giovanni Maria Farinella. Rolling-unrolling LSTMs for action anticipation from first-person video. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):4021–4036, 2020. 2
- [18] Vittorio Gallese, Luciano Fadiga, Leonardo Fogassi, and Giacomo Rizzolatti. Action recognition in the premotor cortex. *Brain*, 119(2):593–609, 1996. 1
- [19] Harshala Gammulle, Simon Denman, Sridha Sridharan, and Clinton Fookes. Predicting the future: A jointly learnt model for action anticipation. In *International Conference on Computer Vision (ICCV)*, 2019. 5
- [20] Rohit Girdhar and Kristen Grauman. Anticipative video transformer. In *International Conference on Computer Vision (ICCV)*, 2021. 2
- [21] 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 *International Conference on Computer Vision (ICCV)*, 2017. 2, 4
- [22] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3D CNNs retrace the history of 2D CNNs and ImageNet? In *Computer Vision and Pattern Recognition (CVPR)*, 2018. 5
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition (CVPR)*, pages 770–778. IEEE, 2016. 5
- [24] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In *Conference on Neural Information Processing Systems Workshops (NIPSW)*, 2015. 2
- [25] Jingyi Hou, Xinxiao Wu, Ruiqi Wang, Jiebo Luo, and Yunde Jia. Confidence-guided self refinement for action prediction in untrimmed videos. *IEEE Transactions on Image Processing*, 29:6017–6031, 2020. 2
- [26] Jian-Fang Hu, Wei-Shi Zheng, Lianyang Ma, Gang Wang, Jianhuang Lai, and Jianguo Zhang. Early action prediction by soft regression. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(11):2568–2583, 2018. 2, 5, 6
- [27] Noureldien Hussein, Efstratios Gavves, and Arnold WM Smeulders. Timeception for complex action recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1

- [28] Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General perception with iterative attention. In *International Conference on Machine Learning (ICML)*, 2021. 1, 3, 4, 5
- [29] Evelyn Kohler, Christian Keysers, M Alessandra Umiltà, Leonardo Fogassi, Vittorio Gallese, and Giacomo Rizzolatti. Hearing sounds, understanding actions: action representation in mirror neurons. *Science*, 297(5582):846–848, 2002. 1
- [30] Dan Kondratyuk, Liangzhe Yuan, Yandong Li, Li Zhang, Mingxing Tan, Matthew Brown, and Boqing Gong. MoViNets: Mobile video networks for efficient video recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 5
- [31] Yu Kong and Yun Fu. Max-margin action prediction machine. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(9):1844–1858, 2015. 2
- [32] Yu Kong, Shangqian Gao, Bin Sun, and Yun Fu. Action prediction from videos via memorizing hard-to-predict samples. In *Association for the Advancement of Artificial Intelligence (AAAI)*, 2018. 2, 4, 5, 6
- [33] Yu Kong, Dmitry Kit, and Yun Fu. A discriminative model with multiple temporal scales for action prediction. In *European Conference on Computer Vision (ECCV)*, 2014. 2, 5
- [34] Yu Kong, Zhiqiang Tao, and Yun Fu. Deep sequential context networks for action prediction. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 4, 5, 6
- [35] Yu Kong, Zhiqiang Tao, and Yun Fu. Adversarial action prediction networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(3):539–553, 2018. 5
- [36] Tian Lan, Tsung-Chuan Chen, and Silvio Savarese. A hierarchical representation for future action prediction. In *European Conference on Computer Vision (ECCV)*, 2014. 2
- [37] Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosiorek, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *International Conference on Machine Learning (ICML)*, 2019. 1, 3, 5
- [38] Kang Li and Yun Fu. Prediction of human activity by discovering temporal sequence patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(8):1644–1657, 2014. 2, 4
- [39] Kang Li, Jie Hu, and Yun Fu. Modeling complex temporal composition of actionlets for activity prediction. In *European conference on computer vision (ECCV)*, 2012. 2
- [40] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video Swin transformer. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 5, 6
- [41] Shugao Ma, Leonid Sigal, and Stan Sclaroff. Learning activity progression in LSTMs for activity detection and early detection. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [42] Yue Meng, Chung-Ching Lin, Rameswar Panda, Prasanna Sattigeri, Leonid Karlinsky, Aude Oliva, Kate Saenko, and Rogerio Feris. AR-Net: Adaptive frame resolution for efficient action recognition. In *European Conference on Computer Vision (ECCV)*, 2020. 2
- [43] Guoliang Pang, Xionghui Wang, Jianfang Hu, Qing Zhang, and Wei-Shi Zheng. DBDNet: Learning bi-directional dynamics for early action prediction. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2019. 2, 5
- [44] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [45] Jie Qin, Li Liu, Ling Shao, Bingbing Ni, Chen Chen, Fumin Shen, and Yunhong Wang. Binary coding for partial action analysis with limited observation ratios. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1
- [46] Giacomo Rizzolatti, Luciano Fadiga, Vittorio Gallese, and Leonardo Fogassi. Premotor cortex and the recognition of motor actions. *Cognitive brain research*, 3(2):131–141, 1996. 1
- [47] Michael S Ryoo. Human activity prediction: Early recognition of ongoing activities from streaming videos. In *International Conference on Computer Vision (ICCV)*, 2011. 2
- [48] Michael S Ryoo, AJ Piergiovanni, Anurag Arnab, Mostafa Dehghani, and Anelia Angelova. TokenLearner: What can 8 learned tokens do for images and videos? In *Conference on Neural Information Processing Systems (NeurIPS)*, 2021. 2
- [49] Mohammad Sadegh Aliakbarian, Fatemeh Sadat Saleh, Mathieu Salzmann, Basura Fernando, Lars Petersson, and Lars Andersson. Encouraging LSTMs to anticipate actions very early. In *International Conference on Computer Vision (ICCV)*, 2017. 4, 6
- [50] Pierre Sermanet, Corey Lynch, Jasmine Hsu, and Sergey Levine. Time-contrastive networks: Self-supervised learning from multi-view observation. In *Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2017. 2
- [51] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. NTU RGB+D: A large scale dataset for 3D human activity analysis. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 4
- [52] Lucas Smaira, João Carreira, Eric Noland, Ellen Clancy, Amy Wu, and Andrew Zisserman. A short note on the kinetics-700-2020 human action dataset. *arXiv preprint arXiv:2010.10864*, 2020. 5
- [53] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 2, 4
- [54] Alexandros Stergiou and Ronald Poppe. Adapool: Exponential adaptive pooling for information-retaining downsampling. *arXiv preprint arXiv:2111.00772*, 2021. 4
- [55] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. 5
- [56] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with

- 3D convolutional networks. In *International Conference on Computer Vision (ICCV)*, 2015. 5
- [57] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2
- [58] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *European Conference on Computer Vision (ECCV)*, 2016. 3, 7
- [59] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [60] Xionghui Wang, Jian-Fang Hu, Jian-Huang Lai, Jianguo Zhang, and Wei-Shi Zheng. Progressive teacher-student learning for early action prediction. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 4, 5, 6
- [61] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. Long-term feature banks for detailed video understanding. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [62] Chao-Yuan Wu, Yanghao Li, Karttikeya Mangalam, Haoqi Fan, Bo Xiong, Jitendra Malik, and Christoph Feichtenhofer. MeMViT: Memory-augmented multiscale vision transformer for efficient long-term video recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1
- [63] Xinxiao Wu, Ruiqi Wang, Jingyi Hou, Hanxi Lin, and Jiebo Luo. Spatial-temporal relation reasoning for action prediction in videos. *International Journal of Computer Vision*, 129(5):1484–1505, 2021. 2, 4, 5, 6
- [64] Xinxiao Wu, Jianwei Zhao, and Ruiqi Wang. Anticipating future relations via graph growing for action prediction. In *Association for the Advancement of Artificial Intelligence (AAAI)*, 2021. 2, 4, 5, 6
- [65] Zuxuan Wu, Caiming Xiong, Chih-Yao Ma, Richard Socher, and Larry S Davis. Adaframe: Adaptive frame selection for fast video recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [66] Mingze Xu, Yuanjun Xiong, Hao Chen, Xinyu Li, Wei Xia, Zhuowen Tu, and Stefano Soatto. Long short-term transformer for online action detection. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 2
- [67] Wanru Xu, Jian Yu, Zhenjiang Miao, Lili Wan, and Qiang Ji. Prediction-cgan: Human action prediction with conditional generative adversarial networks. In *International Conference on Multimedia (ACMMM)*, 2019. 2
- [68] Shen Yan, Xuehan Xiong, Anurag Arnab, Zhichao Lu, Mi Zhang, Chen Sun, and Cordelia Schmid. Multiview transformers for video recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [69] Pengchuan Zhang, Xiyang Dai, Jianwei Yang, Bin Xiao, Lu Yuan, Lei Zhang, and Jianfeng Gao. Multi-scale vision long-former: A new vision transformer for high-resolution image encoding. In *International Conference on Computer Vision (ICCV)*, 2021. 1
- [70] Yanyi Zhang, Xinyu Li, Chunhui Liu, Bing Shuai, Yi Zhu, Biagio Brattoli, Hao Chen, Ivan Marsic, and Joseph Tighe. Vidtr: Video transformer without convolutions. In *International Conference on Computer Vision (ICCV)*, 2021. 2
- [71] He Zhao and Richard P Wildes. Spatiotemporal feature residual propagation for action prediction. In *International Conference on Computer Vision (ICCV)*, 2019. 2, 5
- [72] Yin-Dong Zheng, Zhaoyang Liu, Tong Lu, and Limin Wang. Dynamic sampling networks for efficient action recognition in videos. *IEEE Transactions on Image Processing*, 29:7970–7983, 2020. 2
- [73] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. In *European Conference on Computer Vision (ECCV)*, 2018. 2