

TIM: A Time Interval Machine for Audio-Visual Action Recognition

Jacob Chalk^{1*}

Jaesung Huh^{2*}

Evangelos Kazakos³

Andrew Zisserman²

Dima Damen¹

¹University of Bristol

²VGG, University of Oxford

³Czech Technical University in Prague

Abstract

Diverse actions give rise to rich audio-visual signals in long videos. Recent works showcase that the two modalities of audio and video exhibit different temporal extents of events and distinct labels. We address the interplay between the two modalities in long videos by explicitly modelling the temporal extents of audio and visual events. We propose the Time Interval Machine (TIM) where a modality-specific time interval poses as a query to a transformer encoder that ingests a long video input. The encoder then attends to the specified interval, as well as the surrounding context in both modalities, in order to recognise the ongoing action.

We test TIM on three long audio-visual video datasets: EPIC-KITCHENS, Perception Test, and AVE, reporting state-of-the-art (SOTA) for recognition. On EPIC-KITCHENS, we beat previous SOTA that utilises LLMs and significantly larger pre-training by 2.9% top-1 action recognition accuracy. Additionally, we show that TIM can be adapted for action detection, using dense multi-scale interval queries, outperforming SOTA on EPIC-KITCHENS-100 for most metrics, and showing strong performance on the Perception Test. Our ablations show the critical role of integrating the two modalities and modelling their time intervals in achieving this performance. Code and models at: <https://github.com/JacobChalk/TIM>.

1. Introduction

Long videos exhibit a quick succession of auditory and visual events. The latest attempts to annotate events in these modalities separately [15, 30], showcase that both the temporal extents and class labels differ between the two. However, these events still remain correlated – identifying temporally close events in both modalities can improve recognition of actions in both visual and audio.

Furthermore, most methods to date typically only utilise the exact temporal extent of an action; a precise, trimmed clip of the action is fed into a convolutional [5, 8, 39] or

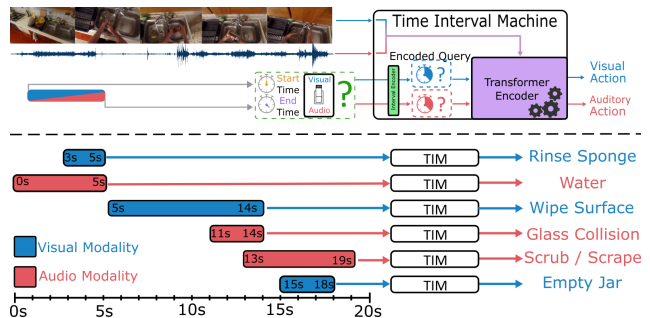


Figure 1. Time Interval Machine (TIM): **Top:** Given a visual and auditory stream input, the ongoing action in a particular *time interval* is determined by a query specifying the start and end time of the interval, along with the modality of interest. **Bottom:** TIM can query for visual (e.g. ‘Rinse Sponge’) and auditory (e.g. ‘Water’) action classes, as well as distinguish between overlapping actions within the same modality (‘Glass Collision’ and ‘Scrub / Scrape’).

transformer-based [1, 11, 22] backbone, which predicts the action taking place. Even when the surrounding context is utilised to improve action recognition [18, 31, 42], again, this context is supplied in the form of exact clips of neighbouring actions, rather than the untrimmed long input video.

In this paper, we propose an approach that encodes multiple events that occur in both visual and auditory streams of a long video input. We achieve this by elevating *time intervals* to first-class citizens, utilising them to specify a query within an accompanying modality. We term this mechanism a **Time Interval Machine (TIM)**. It is able to receive a long video input, and output the actions that occur *within the queried intervals of the queried modalities*.

Consider the example in Figure 1. The input contains the sound of water running while a sponge is being rinsed, which is then used to wipe a surface. These distinct events may vary significantly in duration and may be more prominent in the audio or visual modality. Despite differences between these events, there are likely many correlations between them and the surrounding context, which may be beneficial to recognising a given event (e.g. the sound of water is relevant to rinsing a sponge, providing useful information for recognising the visual action). TIM is able

*Equal technical contribution.

to exploit this by accessing the context within both modalities, including the background when no events occur. It can then distinguish between different, potentially overlapping, events within the same input by querying the time interval of a particular event within a given modality.

We test TIM on three challenging audio-visual **recognition** datasets consisting of long videos: EPIC-KITCHENS [7], which recently offered distinct audio annotations through EPIC-SOUNDS [15], the Perception Test [30], and AVE [36]. We show that TIM can effectively learn both visual and auditory classes within a long input, outperforming the current SOTA top-1 accuracy on EPIC-KITCHENS by 2.9% and 1.4% on EPIC-SOUNDS, despite competing methods for the former using far larger pre-training datasets, large language models or higher resolution input. We also outperform models pre-trained with public datasets on AVE by 0.6% and improve over a strong baseline on the Perception Test in visual and audio action recognition by 9.9% and 3.2% respectively.

Additionally, we adapt TIM to action **detection**, through fixed multi-scale dense querying with an added *interval regression* loss. We report strong detection results on EPIC-KITCHENS and the Perception Test, outperforming Action Former [49] by 1.6 avg mAP and 4.3 respectively.

Our contributions are summarised as: (i) we propose the TIM query mechanism for attending to modality-specific intervals in long videos. (ii) we efficiently train TIM to encode/query multiple audio-visual actions using time intervals. (iii) we showcase the value of TIM for both visual and auditory action recognition, and adapt it for detection with an added interval regression loss. (iv) we achieve new SOTA in both video and multi-modal recognition on multiple datasets.

2. Related Works

Audio-visual action recognition. A number of works have employed audio and visual modalities for action recognition [10, 17, 18, 26, 41, 44]. Some introduce new architectures to effectively fuse modalities [17, 18, 26, 44]; others propose unique training techniques to solve problems occurring while training multi-modal models, such as Gradient Blending [41], to tackle overfitting at different speeds for each modality, or contrastive learning for cross-modal discrimination [25]. However, these works use the same set of semantic and temporal labels for both modalities. Recent works have shown that both the temporal intervals and semantics of events differ between modalities [15, 30]. [37] temporally annotates visual and auditory events independently, although they share the same set of labels. In this work, we train with distinct labels for each modality to leverage discriminative audio and visual actions.

Leveraging temporal context. Several works have considered incorporating temporal context [18, 27, 42, 43, 48], a

direction orthogonal to employing multiple modalities and particularly useful in untrimmed videos. An auto-regressive LSTM-based encoder-decoder is proposed in [27] for action sequence classification, effectively leveraging past action context to predict the current action. The Temporal Query Network [48] uses learnable query vectors that correspond to specific attributes of a long video, allowing the model to attend to the aspects of the video and its surrounding context to produce a response for each attribute. [42] proposes to enhance the representation of the action by aggregating temporal context from neighbouring action clips using a Long-Term Feature Bank along with an attention mechanism. [43] crafts a more sophisticated memory bank by storing keys and values of all the intermediate layers of a transformer to aggregate the past context. Lastly, [18] exploits multi-modal temporal context from surrounding actions using vision, audio, and language.

[18, 42, 43] are the closest to our approach, in that the common goal is to enrich the representation of the action of interest using surrounding context from the untrimmed video, rather than neighbouring clips. Nevertheless, [42, 43] are single modality models, recognising visual actions solely. [18] assumes the temporal extents of all actions are known, including for the test set, which is restrictive.

Queries in visual models. Learning visual queries with Transformer architectures has gained recent attention [4, 14, 16, 23, 48]. Commonly, approaches employ a set of learnable vectors that are used to inquire about the presence of a concept in the input. For example, in [4, 23] the learnable queries correspond to different objects, whereas in [14] they are used for multi-task learning and each learnable query corresponds to a different task. [16] has incorporated learnable queries for adapting a pre-trained model while keeping the rest of its parameters frozen. Closest to our motivation is [48], where the queries correspond to events and their attributes for fine-grained action recognition in videos. The authors note that the queries also have the role of temporally localising the events in untrimmed videos.

Different from [48] and other works, our queries are primarily temporal with no semantic interpretation and are applied to multiple modalities. Importantly, since time is continuous, we cannot use a predefined set of queries. Instead, we employ an MLP architecture to encode time, in a form akin to a universal clock. We present our approach next.

3. Time Interval Machine

In this section, we describe the **Time Interval Machine** (TIM), a multi-modal transformer encoder architecture where all inputs, both features and queries, are encoded with their associated *time intervals*. A time interval incorporates the duration and position of each audio and visual feature and is also used to *query* the network for any action occurring within the given time interval.

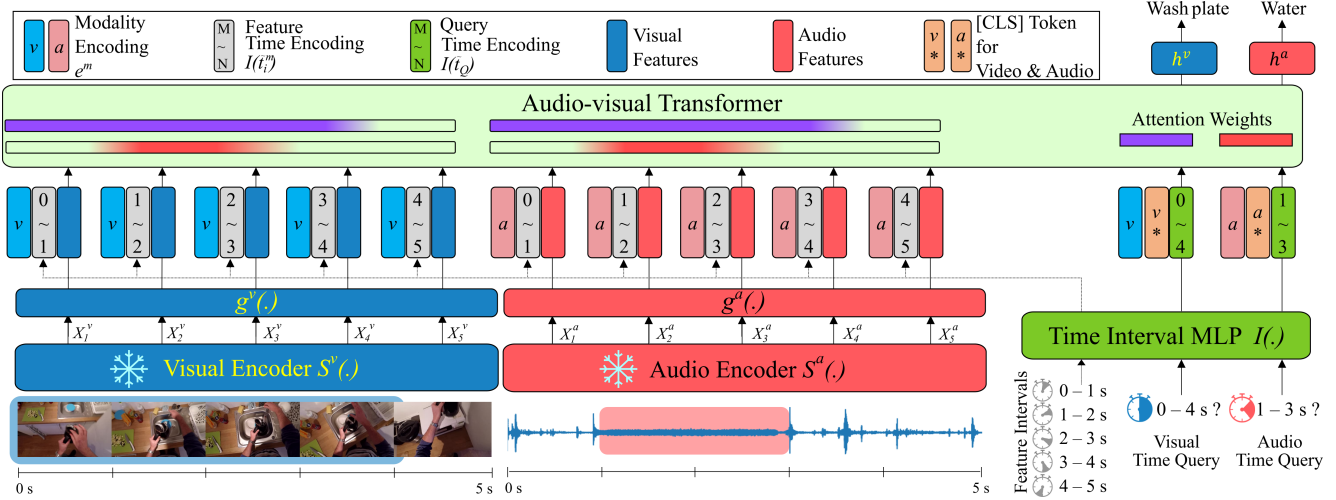


Figure 2. **Overview of the Time Interval Machine (TIM).** The model ingests a sequence of audio and visual features from a video, with each feature time-stamped by the temporal interval it spans, and encoded with its modality. To infer the action occurring over a temporal interval (a visual or audio event) a query is formed specifying the interval and modality of interest.

The architecture of TIM is illustrated in Figure 2. It ingests a large video input, represented as a sequence of audio and visual features, and outputs the ongoing auditory or visual action label for the provided query time intervals.

3.1. Model architecture

Input. The input to TIM is a long crop of the untrimmed video, represented by extracted features. When considering two modality inputs, such as video and audio, each modality is embedded separately as follows: for each modality m , let $\mathbf{X}^m = [X_1^m, \dots, X_{N^m}^m]$ be N^m temporally-ordered feature representations of the input video, obtained from a pre-trained feature extractor $S^m(\cdot)$. We feed the features through modality-specific embedding layers $g^m(\cdot)$, projecting them to a lower, common dimension D across all modalities. The embedded features¹ are then tagged with modality encodings and time interval encodings, forming the input to the transformer encoder. We now detail how we encode the time intervals.

Encoding Time Intervals. In this work, we introduce a new type of learnt query network, the **Time Interval MLP**, which produces a single D -dimensional vector representing a given time interval. This network is used within TIM to encode the time intervals of the input features and the time interval we wish to query, and later classify. Figure 3 illustrates the concept of this network.

The Time Interval MLP $I(\cdot) : \mathbb{R}^2 \rightarrow \mathbb{R}^D$ receives a time interval, represented by start and end time, and produces a single D -dimensional encoding. Note that this is distinct from encoding the start and end times separately. Specifically, let t_s and t_e be the start and end time of an interval of interest, normalised by the length of the long video input.

¹Note that the number of features can differ between modalities




Figure 3. Illustration of the Time Interval MLP $I(\cdot)$. It inputs the **two** dimensional vector, start and end times of an interval, and produces a single vector, which can be concatenated along the channel dimension to either input features or [CLS] tokens. The figure shows three time interval inputs and three corresponding outputs. Note that in practice, time intervals are ingested simultaneously.

$I(\cdot)$ receives the interval $\tilde{t} = [t_s, t_e]$ as input, and outputs a D -dimensional vector encoding of that interval. This vector encodes both the relative position of the time interval within the input, as well as its duration. This vector then acts as a *query* for the model concerning the action taking place within the interval. Furthermore, each feature $\{X_i^m\}$ spans a certain time interval within the input. Thus, it is important to also encode the time intervals of the features.

In summary, the Time Interval MLP acts as an *universal clock*, which encodes the temporal extent of features, from any modality, within the input. Note that it is critical that the same Time Interval MLP is used for encoding all time intervals of the input features and queries across both modalities to accurately encode universal time. It is also important to note that Time interval MLP can cover *continuous* time intervals, whereas traditional positional encoding only covers a fixed set of positions of the input features. The Time Interval MLP is trained end-to-end along with the transformer.

Transformer Feature Inputs. Let $\tilde{\mathbf{t}}^m = [\tilde{t}_1^m, \dots, \tilde{t}_{N^m}^m]$

be the corresponding time intervals of the video’s features \mathbf{X}^m from the modality m . We inject the encoded time interval $I(\tilde{t}^m)$ into the embedded features via channel-wise concatenation. A learnable modality-specific encoding $e^m \in \mathbb{R}^{2D}$ is then summed to the temporally-encoded features to discriminate between each modality. In summary, the feature inputs \mathbf{E}^m for TIM are computed by,

$$E_i^m = [g^m(X_i^m), I(\tilde{t}_i^m)] + e^m \quad \forall i \in [1, \dots, N^m] \quad (1)$$

where $[\cdot, \cdot]$ indicates concatenation.

Transformer Query Inputs. To query for an action within an interval of interest, we adopt a standard approach of appending a learnable classification token to the input sequence, CLS^m . If \tilde{t}_Q is an interval of interest, we concatenate the time interval representation $I(\tilde{t}_Q)$ to this classification token along the channel dimension, which acts as a query for the network in order to predict the corresponding action happening within \tilde{t}_Q . We also add the modality-specific encoding e^m to each classification token, as a flag to distinguish between which modality we are querying. The encoded $[\text{CLS}]^m$ tokens can be more formally defined as:

$$[\text{CLS}]^m = [\text{CLS}^m, I(\tilde{t}_Q)] + e^m \quad (2)$$

During training, we add a classification token for *each* action within the input video, resulting in multiple $[\text{CLS}]$ tokens across both modalities.

Transformer Encoder. We use a transformer encoder to perform self-attention on the input sequence to aggregate relevant temporal context and cross-modal relations.

We form the transformer input sequence with the encoded feature inputs \mathbf{E}^m and **one or more** classification tokens $[\text{CLS}]^m$, representing each time interval query, and feed these into the encoder. Note that we recognise all actions from any modality simultaneously by appending multiple CLS^m tokens to the input. The transformer output representation of $[\text{CLS}]^m$, namely Z_{CLS}^m , is then passed to the corresponding linear classifier to predict the action labels.

Importantly, we use an attention mask to prevent queries from attending to one another, and we similarly prevent input features from attending to queries. This ensures each query is recognised without the privileged knowledge of any other query, or action boundary, during inference.

3.2. Training and Testing in TIM

To train TIM, we consider all long segments of W seconds and stride H_w across the entirety of the untrimmed videos. We randomly select batches from these. For each window, we query all annotated audio and visual actions that overlap with the window by more than $\delta = 0.2$ seconds.

All queries in the window are encoded and concatenated to separate CLS tokens. To classify queries, let $h_{\text{CLS}}^m(\cdot)$ be a linear classifier for modality m , and let $\hat{y}_{\text{CLS}}^m = h^m(Z_{\text{CLS}}^m)$

be the predicted action of the output representation Z_{CLS}^m . We train TIM by using a cross-entropy classification loss $CE(\cdot)$ on the ground truth y_{CLS}^m by:

$$L^m = \frac{1}{N_Q} \sum^{N_Q} CE(\hat{y}_{\text{CLS}}^m, y_{\text{CLS}}^m) \quad (3)$$

where N_Q is the number of queries within the batch.

Temporal Distance Loss. In addition to the standard classification loss, we introduce a Temporal Distance (TD) loss as an auxiliary loss for training TIM. Inspired by [21], where relative patch positions in token embeddings are learnt using self-supervision, we similarly train the network to take two transformer outputs and predict the elapsed time between their corresponding time intervals.

Let the $\mathbf{Z}_{1:\sum_m N^m}$ be the transformer outputs of the features from all modalities. We randomly sample a set of feature pairs $\mathbb{B} \subset \mathbf{Z}_{1:\sum_m N^m}$ from these outputs, concatenate along the channel dimension and feed them to the temporal distance regression head $h_{\tilde{t}}(\cdot) : \mathbb{R}^{4D} \rightarrow \mathbb{R}^1$ to predict the time interval difference between each pair. Note that feature pairs can be sampled both *within* and *across* modalities. In our case, we sample across modalities by pairing one visual feature with another audio feature. This helps the model to learn the temporal relations between modalities.

Formally, a TD loss L^{td} is computed as:

$$L^{td} = \sum_{\{Z_i, Z_j\} \in \mathbb{B}} |h_{\tilde{t}}(Z_i, Z_j) - d_{ij}| \quad (4)$$

where d_{ij} is the temporal distance between intervals \tilde{t}_i, \tilde{t}_j .

Training objective and regime. For our final training loss, we sum the losses across modalities along with the TD loss:

$$L^{total} = \left(\sum_{m \in \mathbb{M}} \lambda^m L^m \right) + \lambda^{td} L^{td} \quad (5)$$

where \mathbb{M} is a set of modalities, λ^m controls the strength of each modality’s loss and λ^{td} is a hyperparameter that controls the strength of the TD loss.

Test-Time Augmentation. We use test-time augmentation, as this generally increases prediction robustness and performance [29, 32]. In TIM, we use a sliding window across the untrimmed video, thus feeding the same interval query with varying contexts. We then aggregate predictions of the same interval query across windows to make the final prediction.

3.3. Adapting for Detection

While primarily designed for recognition, we can adapt TIM for detection. The backbone remains largely unchanged from recognition, but there are two main differences. First, we construct dense multi-scale interval queries spanning the entirety of the video input at each scale. These

are used as interval queries in both training and detection inference. The multi-scale intervals allow for detecting both long and short actions. Second, we introduce an additional interval regression head, which regresses the query interval to the action’s exact temporal duration.

During training, we deem any query in the multi-scale pyramid that overlaps with a ground truth action by more than some IoU threshold as a positive query. In addition to classifying the query, we train a DIOU regression loss [52] to predict the exact interval of the action. Both classification and interval regression losses are trained jointly. We provide full details in the ArXiv appendix.

4. Experiments

This section describes the datasets used to evaluate our model, implementation details and results along with a comparison to the state-of-the-art methods.

4.1. Dataset

EPIC-KITCHENS-100 [7] is a large-scale video dataset, including 700 egocentric videos recording actions in kitchens. It consists of 89,977 segments of fine-grained actions. Inspired by prior works [11, 34, 35], we directly predict the action out of 3806 classes present in the train and validation set to avoid predicting invalid actions.

EPIC-SOUNDS [15] offers audio annotations capturing temporal extents and class labels within the audio stream of EPIC-KITCHENS-100. The annotations contain 78,366 labelled audio events. We combine the visual annotations from EPIC-KITCHENS with the audio annotations from EPIC-SOUNDS to train our audio-visual model. TIM can recognise actions from both datasets using a single model.

AVE [36] contains 4,143 videos covering a range of real-life scenes and labelled with 27 categories, such as church bell, male speaking and dog barking. Each video is equally divided into 10 segments, each 1 second in length. We evaluate TIM on the supervised audio-visual event localisation task. Given a 1 second segment, we recognise the ongoing action out of the 27 categories *plus* a background class.

Perception Test [30] is a recent multimodal video benchmark of 11,620 videos with an average length of 23 seconds, and provides both temporal action and sound annotations. There are 73,503 visual annotations spanning 63 classes, versus 137,128 sound annotations over 16 classes.

4.2. Implementation Details

Architectural Details. Visual and audio embedding layers g_m consist of a single 512-D feed-forward layer, followed by a GELU [13] activation and layer normalisation [2] are used to project the features to a common space. The Time Interval MLP I consists of three linear layers with 512-D hidden dimension, followed by ReLU activations, with

layer normalisation after the output of the last linear layer. We include 512-D learnable [CLS] tokens: $[\text{CLS}]_{action}^m$ for each query in each modality which become 1024-D after concatenation with the encoded time interval. They are then summed with 1024-D modality encodings; e^m .

The audio-visual transformer contains four encoder layers, each with 8 attention heads, GELU activations, and 1024-D keys, queries and values. A dropout rate of $p = 0.1$ is applied within the encoder layers. We also apply channel-wise dropout with $p = 0.5$ directly to the raw input features, as well as to the encoded transformer input. The temporal distance head consists of two linear layers with a hidden dimension of 1024 and a third which outputs a single number corresponding to the elapsed time between each time interval. We include the architectural ablations on encoder layers and the temporal distance head in the ArXiv appendix.

Training / Validation Details. We train each model for 100 epochs, using AdamW [24] with a batch size of 64 and a weight decay of $1e-4$. A linear learning rate warm-up is applied for the first two epochs, starting from $1e-6$ to a target learning rate, and we use a cosine learning rate scheduler. We set the TD loss weight λ^{td} to 0.3. We pad the queries for each window in the batch to the maximum number of queries in a single window in each dataset. We provide implementation details per dataset in the ArXiv appendix.

4.3. Results

We compare TIM with SOTA models for each dataset.

EPIC-KITCHENS / EPIC-SOUNDS Results. We train a single model on both visual and audio labels of the EPIC-KITCHENS videos, reporting results on both datasets.

For the visual features, we concatenate Omnivore [11] and VideoMAE-L [38] features along the channel dimension, forming 2048-D features. For the audio features, we use Auditory SlowFast [19], which generalises well across diverse audio domains [40]. For both modalities, we extract 1 second features every 0.2s. For training, we extract additional augmented feature sets - with RandAugment [6] for visual and SpecAugment [28] for audio features.

Table 1 compares TIM with the SOTA models on EPIC-KITCHENS-100. We outperform M&M Mix [45] by 5.1% on verb, 0.9% on noun and 3.9% on action. Compared to our model, both MTV and M&M Mix are trained with an additional private dataset [33] which contains 194K hours of 70 million videos while we only use the open-source visual backbone pre-trained with public datasets. We also outperform LaViLa [51] and AVION [50] which leverage pre-trained LLMs to learn video representations.

We note that we outperform all prior works, often without additional techniques that boost performance. For example, we use short-sided cropped 224×224 images while [45] uses 420×420 , which enlarges the spatial resolution of objects in the egocentric video, enabling better noun recog-

| Model | x_p | LLM | Verb | Noun | Action |
|----------------------------|-------|-----|-------------|-------------|-------------|
| <i>Visual-only models</i> | | | | | |
| MFormer-HR [29] | 336p | ✗ | 67.0 | 58.5 | 44.5 |
| MoViNet-A6 [20] | 320p | ✗ | 72.2 | 57.3 | 47.7 |
| MeMViT [43] | 224p | ✗ | 71.4 | 60.3 | 48.4 |
| Omnivore [11] | 224p | ✗ | 69.5 | 61.7 | 49.9 |
| MTV [47] | 280p | ✗ | 69.9 | 63.9 | 50.5 |
| LaViLa (TSF-L) [51] | 224p | ✓ | 72.0 | 62.9 | 51.0 |
| AVION (ViT-L) [50] | 224p | ✓ | 73.0 | 65.4 | 54.4 |
| TIM (ours) | 224p | ✗ | 76.2 | 66.4 | 56.4 |
| <i>Audio-visual models</i> | | | | | |
| TBN [17] | 224p | ✗ | 66.0 | 47.2 | 36.7 |
| MBT [26] | 224p | ✗ | 64.8 | 58.0 | 43.4 |
| MTCN [18] | 336p | ✗ | 70.7 | 62.1 | 49.6 |
| M&M [45] | 420p | ✗ | 72.0 | 66.3 | 53.6 |
| TIM (ours) | 224p | ✗ | 77.1 | 67.2 | 57.5 |

Table 1. Comparisons to state-of-the-art *recognition* models on the EPIC-KITCHENS validation set. We report the top-1 accuracy for verb, noun and action (%). LLM: large language model is used during pre-training. x_p : input resolution of $x \times x$.

| Model | SSAST [12] | ASF [19] | DiffSED [3] | TIM (A) | TIM (A+V) |
|------------------|------------|----------|-------------|---------|-------------|
| Top-1 acc | 53.5 | 53.8 | 56.9 | 55.7 | 58.3 |

Table 2. Comparisons to state-of-the-art *sound recognition* models on EPIC-SOUNDS. We report the top-1 accuracy (%) on Val. The performance of SSAST and ASF are from [15].

| Model | V | A | Average Precision (AP) | | | | | | |
|-----------------------------|---|---|------------------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | | | Task | @0.1 | @0.2 | @0.3 | @0.4 | @0.5 | Avg. |
| G-TAD [46] | ✓ | ✗ | Verb | 12.1 | 11.0 | 9.4 | 8.1 | 6.5 | 9.4 |
| | | | Noun | 11.0 | 10.0 | 8.6 | 7.0 | 5.4 | 8.4 |
| ActionFormer [49] | ✓ | ✗ | Verb | 26.6 | 25.4 | 24.2 | 22.3 | 19.1 | 23.5 |
| | | | Noun | 25.2 | 24.1 | 22.7 | 20.5 | 17.0 | 21.9 |
| ActionFormer - Our Features | ✓ | ✗ | Verb | 29.6 | 28.8 | 26.9 | 24.4 | 21.6 | 26.3 |
| | | | Noun | 34.3 | 32.6 | 30.2 | 27.4 | 22.6 | 29.4 |
| TIM | ✓ | ✓ | Verb | 32.9 | 31.6 | 29.6 | 27.0 | 22.2 | 28.6 |
| | | | Noun | 36.4 | 34.8 | 32.1 | 28.7 | 22.7 | 31.0 |

Table 3. Comparisons to state-of-the-art *detection* models on the EPIC-KITCHENS validation set. We report the average precision at IOU thresholds [0.1, 0.2, 0.3, 0.4, 0.5] as well as their average across all thresholds on verb, noun.

recognition. We expect a further performance boost when implementing any of: higher resolution feature extractors, additional large-scale pre-training and the introduction of a LLM. We leave this as an avenue for future work.

Table 2 compares TIM against prior results on EPIC-SOUNDS, where TIM outperforms SOTA by 1.4%.

For detection, we show that TIM can produce competitive results when compared to models primarily designed for this task in Table 3. TIM adapted for detection outperforms ActionFormer [49] by 2.3 mAP on verb and 1.6 mAP on noun using the same set of features.

AVE Results. As this dataset contains joint audio-visual labels, we train TIM by duplicating the query, i.e. using a [CLS] for each modality, and combine their logits during training and inference. We use the pre-trained publicly available models from [36] for a fair comparison with other

| Model | PSP | CPSP | CSSNet | TIM | | | |
|------------------|------|------|-------------|------|------|------|------|
| | [53] | [54] | [9]† | V | A | AV | AV* |
| Top-1 acc | 77.8 | 78.6 | 80.5 | 62.8 | 65.5 | 79.2 | 79.8 |

Table 4. Top-1 event classification accuracy (%) on AVE Test set. †: no official code or public model provided to replicate results. We show the models trained only with publicly available datasets. *: results with Omni+ASF features.

| <i>Perception Test Action</i> | | | | |
|-------------------------------|---------|----------------|---------|-------------|
| Model | MLP (V) | MTCN [18](A+V) | TIM (V) | TIM (A+V) |
| Top-1 acc | 43.7 | 51.2 | 56.1 | 61.1 |
| <i>Perception Test Sound</i> | | | | |
| Model | MLP (A) | MTCN [18](A+V) | TIM (A) | TIM (A+V) |
| Top-1 acc | 50.6 | 52.9 | 54.8 | 56.1 |

Table 5. Comparisons to trained recognition baselines on the Perception Test validation split. We show both action and sound recognition and the benefit of including audio-visual in TIM for both challenges. V : visual and A : audio input features. MLP is the result by training an MLP classifier with the features directly.

| Model | Average Precision (AP) | | | | | |
|-------------------------------|------------------------|-------------|-------------|-------------|-------------|-------------|
| | @0.1 | @0.2 | @0.3 | @0.4 | @0.5 | Avg. |
| <i>Perception Test Action</i> | | | | | | |
| ActionFormer [49] | 27.8 | 27.6 | 25.2 | 23.0 | 20.0 | 24.5 |
| TIM | 33.5 | 32.2 | 29.8 | 26.4 | 22.0 | 28.8 |
| <i>Perception Test Sound</i> | | | | | | |
| ActionFormer [49] | 34.7 | 31.3 | 27.5 | 22.7 | 17.7 | 26.8 |
| TIM | 37.5 | 33.1 | 27.9 | 22.8 | 17.2 | 27.7 |

Table 6. Comparisons to strong *detection* models on the Perception Test validation set for action and sound localisation. We report the average precision at IOU thresholds [0.1, 0.2, 0.3, 0.4, 0.5] as well as the average across all thresholds.

works. We also apply AVGA [36] to spatial visual features from VGG-19 before feeding them to the transformer.

Table 4 shows our results on the AVE dataset. Combining audio and video significantly improves the performance on TIM. The results from [9] perform best but could not be replicated. We also report TIM using the Omnivore visual features and Auditory Slowfast features used for EPIC-KITCHENS, which achieves a 0.6% boost in performance.

Perception Test Results. We use the same backbone for the Omnivore features and Auditory Slowfast features and train a single model using both the visual and audio labels. Table 5 compares the results on the newly introduced Perception Test. We train an MLP classifier, with two linear layers and a ReLU activation, on the features directly as a baseline. We also evaluate on an audio-visual model that does use context with MTCN. Compared to these methods TIM clearly shows significant improvements. Results are improved over MTCN by 9.9% and 3.2% on visual and audio recognition tasks respectively. We also provide detection results in Table 6. TIM improves over ActionFormer [49] by 3.3 average mAP on visual actions and by 0.9 average mAP on sound, when using the same features.

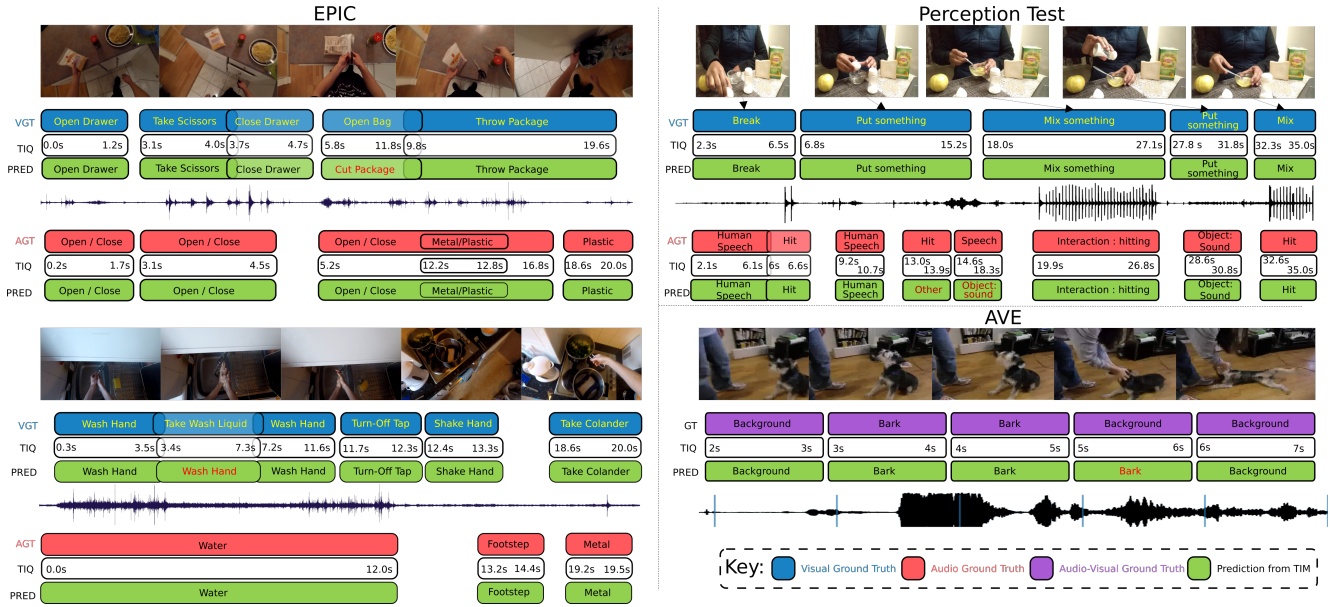


Figure 4. Qualitative results for all datasets. **PRED**: Prediction by TIM, **TIQ**: Time Interval Queries, **V/AGT**: Visual/Audio Ground Truth.

Cross-Modality in TIM. When referring to our previous results, we see that including the additional modality provides a performance boost in all cases, highlighting TIM’s ability to utilise and distinguish between different modalities. For example, on EPIC-KITCHENS-100, including audio improves visual action accuracy by 0.9%. For EPIC-SOUNDS, the visual modality further improves accuracy by 2.6%. In the Perception Test, including the audio modality improves visual recognition by 5.0%, and visual increases sound recognition by 1.3%. Finally, for AVE, we see a significant improvement, where an audio-visual model increases accuracy by 13.7% from audio-only.

Qualitative Results. We present qualitative results in Figure 4. We see that in EPIC-KITCHENS, TIM can competently recognise actions across the two modalities including overlapping queries. Furthermore, we see consecutive actions are correctly recognised with varying interval lengths, such as the ‘open / close’ audio actions between 0.2s and 4.5s. For AVE, TIM is able to distinguish between the background and a ‘barking’ audio-visual event based on the time interval query. For the Perception Test, we see that TIM can distinguish between heavily overlapping actions across both modalities, such as between ‘break’, ‘human speech’, ‘hit’ and ‘put something’. However, there are also failure cases, such as in EPIC-KITCHENS when the action ‘take washing up liquid’ is recognised as ‘wash hand’, as the model is likely confused by the context predominantly associated with the highly overlapping ‘wash hand’ actions.

4.4. Analysing Time Intervals

We showcase the importance of effectively encoding time intervals, as well as how they differ from alternative strate-




Figure 5. TSNE plot for time encodings $I(\cdot)$ on all datasets. In each plot, we use colour maps to indicate encodings of the time interval’s duration (left), start time (middle) and end time (right).

gies. We perform this analysis on the EPIC-KITCHENS-100 and EPIC-SOUNDS recognition tasks.

Time Encoding Representation. To show TIM encodes time intervals across all datasets in Figure 5. We use three colour maps on the same TSNE projection to show three properties of the encoded interval: duration, start time and end time. Interestingly, the 1D time encoding perfectly captures all three attributes and across the datasets. While the encodings differ per dataset, as these differ in the positions and durations of actions, we see clear similarities in the learnt time encoding projections. For example, the duration is perfectly captured across the x-axis of the TSNE plot with lower values indicating longer time intervals.



Figure 6. Attention heatmaps of the second encoder layer for two random 30s clips in EPIC-KITCHENS. **x-axis**: input feature time intervals; **y-axis** query time intervals of varying position and duration. The attention magnitude relates to the query CLS token.




Figure 7. The impact of shifting the time interval query on visual performance (top) and audio performance (bottom), both on short actions (< 2 sec) and long actions (> 2 sec) and overall validation set (**all**) for EPIC-KITCHENS-100 and EPIC-SOUNDS.

Interval Query Attention. We plot two attention heatmaps in Figure 6 for 5 separate queries with varying positions and scales in EPIC-KITCHENS-100. We extract the attention weights from the second transformer encoder layer, as this appears to be the most relevant to the interval query. The learnt attention clearly applies to the feature time intervals contained within the query. We note the similarity between the attention in the two randomly selected windows.

Shifting Intervals. To show how TIM effectively encodes the time interval of actions, we shift the time interval queries from their correct action interval by -1.5s to 1.5s, assessing the impact of these adjustments on the performance.

Figure 7 shows the result. We see the performance gradually drops, both in visual and audio, as the query interval moves away from the correct action interval. The drop is also symmetric showcasing no bias. Unsurprisingly, the performance drops significantly when shifting short actions both in video (-57.9%) and audio (-35.2%), while it is less

| | EPIC-KITCHENS | | | EPIC-SOUNDS |
|-------------------------|---------------|-------------|-------------|---------------|
| Encoding | Verb | Noun | Action | Audio Actions |
| Learned | 43.8 | 44.3 | 29.6 | 23.7 |
| Sinusoidal | 43.8 | 44.6 | 30.0 | 13.4 |
| Centre | 74.3 | 65.8 | 55.6 | 56.4 |
| Separate-add | 76.0 | 66.2 | 56.4 | 57.7 |
| Interval-add | 76.3 | 66.5 | 56.9 | 58.8 |
| Separate-cat | 76.8 | 67.4 | 57.1 | 58.4 |
| Interval-cat (proposed) | 77.1 | 67.2 | 57.5 | 58.3 |

Table 7. Ablating the choice of encoding time intervals.

extreme in long actions (-14.5% and -11.2%). We assess the impact of scaling the time interval in the ArXiv appendix.

Time Interval Encodings. The Time Interval MLP encodes the interval of the query. Here, we compare this to traditional positional encodings, both sinusoidal and learned. We also experiment on five different variations of the Time Interval MLP, namely: (i) Centre – we only encode the centre timestamp of the interval; (ii) Separate-Add/Cat – we encode the intervals’ start and end time separately and add the encoded output vectors together, or concatenate along the channel dimension; and (iii) Interval-Add/Cat – we encode the start and end time within the same vector and add, or concatenate, the encoded output to the input sequence.

We show the results in Table 7. In all cases, the final encoding is of the same dimension for comparable results. Performance is significantly worse with sinusoidal or learned positional encoding, as they are unable to capture the complexity of overlapping actions. There is also a drop when encoding only the centre of the time interval.

Separate-Add/Cat are alternative ways to encode the intervals (and hence include duration information) resulting in comparable performance to the interval counterparts. Our proposed approach to encoding the interval into the MLP shows the best performance for visual and while maintaining strong auditory performance.

5. Conclusions

In this paper, we propose to utilise the action’s time interval as a query to an audio-visual transformer which learns to recognise the action from its interval and the unaltered surrounding context. We jointly train the model on modality-specific time intervals and label sets, allowing the Time Interval Machine (TIM) to recognise multiple events across both visual and auditory modalities.

TIM is sensitive to the interval’s position and duration. This allows the model, as is, to produce competitive results on action detection through multi-scale dense querying.

Acknowledgements. This work uses public datasets. It is supported by EPSRC Doctoral Training Program, EPSRC UMPIRE EP/T004991/1 and EPSRC Programme Grant VisualAI EP/T028572/1; and by the use of the EPSRC funded Tier 2 facility JADE-II.

References

- [1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6836–6846, 2021. 1
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 5
- [3] Swapnil Bhosale, Sauradip Nag, Diptesh Kanojia, Jiankang Deng, and Xiatian Zhu. Diffused: Sound event detection with denoising diffusion. *arXiv preprint arXiv:2308.07293*, 2023. 6
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Proc. ECCV*, 2020. 2
- [5] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 1
- [6] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703, 2020. 5
- [7] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Jian Ma, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Rescaling egocentric vision: Collection pipeline and challenges for epic-kitchens-100. *International Journal of Computer Vision (IJCV)*, 2021. 2, 5
- [8] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. SlowFast Networks for Video Recognition. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2019. 1
- [9] Fan Feng, Yue Ming, Nannan Hu, Hui Yu, and Yuanan Liu. Cst-net: A consistent segment selection network for audio-visual event localization. *IEEE Transactions on Multimedia*, 2023. 6
- [10] Ruohan Gao, Tae-Hyun Oh, Kristen Grauman, and Lorenzo Torresani. Listen to look: Action recognition by previewing audio. In *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [11] Rohit Girdhar, Mannat Singh, Nikhila Ravi, Laurens van der Maaten, Armand Joulin, and Ishan Misra. Omnivore: A Single Model for Many Visual Modalities. In *CVPR*, 2022. 1, 5, 6
- [12] Yuan Gong, Cheng-I Lai, Yu-An Chung, and James Glass. Ssast: Self-supervised audio spectrogram transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10699–10709, 2022. 6
- [13] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. 5
- [14] Roei Herzig, Ofir Abramovich, Elad Ben-Avraham, Assaf Arbelle, Leonid Karlinsky, Ariel Shamir, Trevor Darrell, and Amir Globerson. Promptonomyvit: Multi-task prompt learning improves video transformers using synthetic scene data. *arXiv preprint arXiv:2212.04821*, 2022. 2
- [15] Jaesung Huh, Jacob Chalk, Evangelos Kazakos, Dima Damen, and Andrew Zisserman. EPIC-SOUNDS: A Large-Scale Dataset of Actions that Sound. In *IEEE International Conference on Acoustics, Speech, & Signal Processing (ICASSP)*, 2023. 1, 2, 5, 6
- [16] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *European Conference on Computer Vision (ECCV)*, 2022. 2
- [17] Evangelos Kazakos, Arsha Nagrani, Andrew Zisserman, and Dima Damen. Epic-fusion: Audio-visual temporal binding for egocentric action recognition. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2019. 2, 6
- [18] Evangelos Kazakos, Jaesung Huh, Arsha Nagrani, Andrew Zisserman, and Dima Damen. With a little help from my temporal context: Multimodal egocentric action recognition. In *Proc. BMVC*, 2021. 1, 2, 6
- [19] Evangelos Kazakos, Arsha Nagrani, Andrew Zisserman, and Dima Damen. Slow-fast auditory streams for audio recognition. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021. 5, 6
- [20] Dan Kondratyuk, Liangzhe Yuan, Yandong Li, Li Zhang, Mingxing Tan, Matthew Brown, and Boqing Gong. Movinets: Mobile video networks for efficient video recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16020–16030, 2021. 6
- [21] Yahui Liu, Enver Sangineto, Wei Bi, Nicu Sebe, Bruno Lepri, and Marco Nadai. Efficient training of visual transformers with small datasets. *Advances in Neural Information Processing Systems*, 34, 2021. 4
- [22] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3202–3211, 2022. 1
- [23] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [24] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 5
- [25] Pedro Morgado, Nuno Vasconcelos, and Ishan Misra. Audio-visual instance discrimination with cross-modal agreement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12475–12486, 2021. 2
- [26] Arsha Nagrani, Shan Yang, Anurag Arnab, Cordelia Schmid, and Chen Sun. Attention bottlenecks for multimodal fusion. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 2, 6
- [27] Yan Bin Ng and Basura Fernando. Human action sequence classification. *CoRR*, abs/1910.02602, 2019. 2

- [28] Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D Cubuk, and Quoc V Le. Specaugment: A simple data augmentation method for automatic speech recognition. *arXiv preprint arXiv:1904.08779*, 2019. [5](#)
- [29] Mandela Patrick, Dylan Campbell, Yuki M. Asano, Ishan Misra Florian Metzger, Christoph Feichtenhofer, Andrea Vedaldi, and João F. Henriques. Keeping your eye on the ball: Trajectory attention in video transformers. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. [4](#), [6](#)
- [30] Viorica Pătrăucean, Lucas Smaira, Ankush Gupta, Adria Recasens Contente, Larisa Markeeva, Dylan Banarse, Skanda Koppula, Joseph Heyward, Mateusz Malinowski, Yi Yang, Carl Doersch, Tatiana Matejovicova, Yury Sulsky, Antoine Miech, Alex Frechette, Hanna Klimczak, Raphael Koster, Junlin Zhang, Stephanie Winkler, Yusuf Aytar, Simon Osindero, Dima Damen, Andrew Zisserman, and João Carreira. Perception test: A diagnostic benchmark for multimodal video models, 2023. [1](#), [2](#), [5](#)
- [31] Fadime Sener, Dipika Singhania, and Angela Yao. Temporal aggregate representations for long-range video understanding. In *Proceedings of European Conference on Computer Vision (ECCV)*, 2020. [1](#)
- [32] Divya Shanmugam, Davis Blalock, Guha Balakrishnan, and John Gutttag. Better aggregation in test-time augmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2021. [4](#)
- [33] Jonathan C Stroud, Zhichao Lu, Chen Sun, Jia Deng, Rahul Sukthankar, Cordelia Schmid, and David A Ross. Learning video representations from textual web supervision. *arXiv preprint arXiv:2007.14937*, 2020. [5](#)
- [34] Swathikiran Sudhakaran, Adrian Bulat, Juan-Manuel Perez-Rua, Alex Falcon, Sergio Escalera, Oswald Lanz, Brais Martinez, and Georgios Tzimiropoulos. Saic_cambridge-hupba-fbk submission to the epic-kitchens-100 action recognition challenge 2021. *arXiv preprint arXiv:2110.02902*, 2021. [5](#)
- [35] Tsung-Ming Tai, Oswald Lanz, Giuseppe Fiameni, Yi-Kwan Wong, Sze-Sen Poon, Cheng-Kuang Lee, Ka-Chun Cheung, and Simon See. Nvidia-unibz submission for epic-kitchens-100 action anticipation challenge 2022. *arXiv preprint arXiv:2206.10869*, 2022. [5](#)
- [36] Yapeng Tian, Jing Shi, Bochen Li, Zhiyao Duan, and Chenliang Xu. Audio-visual event localization in unconstrained videos. In *ECCV*, 2018. [2](#), [5](#), [6](#)
- [37] Yapeng Tian, Dingzeyu Li, and Chenliang Xu. Unified multisensory perception: Weakly-supervised audio-visual video parsing. In *Proc. ECCV*, 2020. [2](#)
- [38] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. VideoMAE: Masked autoencoders are data-efficient learners for self-supervised video pre-training. In *Advances in Neural Information Processing Systems*, 2022. [5](#)
- [39] 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*, pages 20–36. Springer, 2016. [1](#)
- [40] Luyu Wang, Pauline Luc, Yan Wu, Adria Recasens, Lucas Smaira, Andrew Brock, Andrew Jaegle, Jean-Baptiste Alayrac, Sander Dieleman, Joao Carreira, et al. Towards learning universal audio representations. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4593–4597. IEEE, 2022. [5](#)
- [41] Weiyao Wang, Du Tran, and Matt Feiszli. What makes training multi-modal classification networks hard? In *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [2](#)
- [42] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. Long-term feature banks for detailed video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 284–293, 2019. [1](#), [2](#)
- [43] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13587–13597, 2022. [2](#), [6](#)
- [44] Fanyi Xiao, Yong Jae Lee, Kristen Grauman, Jitendra Malik, and Christoph Feichtenhofer. Audiovisual slowfast networks for video recognition. *arXiv preprint arXiv:2001.08740*, 2020. [2](#)
- [45] Xuehan Xiong, Anurag Arnab, Arsha Nagrani, and Cordelia Schmid. M&m mix: A multimodal multiview transformer ensemble. *arXiv preprint arXiv:2206.09852*, 2022. [5](#), [6](#)
- [46] Mengmeng Xu, Chen Zhao, David S. Rojas, Ali Thabet, and Bernard Ghanem. G-tad: Sub-graph localization for temporal action detection, 2020. [6](#)
- [47] Shen Yan, Xuehan Xiong, Anurag Arnab, Zhichao Lu, Mi Zhang, Chen Sun, and Cordelia Schmid. Multiview transformers for video recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3333–3343, 2022. [6](#)
- [48] Chuhan Zhang, Ankush Gupta, and Andrew Zisserman. Temporal query networks for fine-grained video understanding. In *Proc. CVPR*, 2021. [2](#)
- [49] Chen-Lin Zhang, Jianxin Wu, and Yin Li. Actionformer: Localizing moments of actions with transformers. In *European Conference on Computer Vision*, pages 492–510, 2022. [2](#), [6](#)
- [50] Yue Zhao and Philipp Krähenbühl. Training a large video model on a single machine in a day. *arXiv preprint arXiv:2309.16669*, 2023. [5](#), [6](#)
- [51] Yue Zhao, Ishan Misra, Philipp Krähenbühl, and Rohit Girdhar. Learning video representations from large language models. In *CVPR*, 2023. [5](#), [6](#)
- [52] Zhaohui Zheng, Ping Wang, Wei Liu, Jinze Li, Rongguang Ye, and Dongwei Ren. Distance-iou loss: Faster and better learning for bounding box regression. In *The AAAI Conference on Artificial Intelligence (AAAI)*, pages 12993–13000, 2020. [5](#)
- [53] Jinxing Zhou, Liang Zheng, Yiran Zhong, Shijie Hao, and Meng Wang. Positive sample propagation along the audio-visual event line. In *Proceedings of the IEEE/CVF Con-*

ference on Computer Vision and Pattern Recognition, pages 8436–8444, 2021. 6

- [54] Jinxing Zhou, Dan Guo, and Meng Wang. Contrastive positive sample propagation along the audio-visual event line. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 6