

# Learning Streaming Video Representation via Multitask Training

Yibin Yan<sup>1,2\*</sup>, Jilan Xu<sup>3,4\*</sup>, Shangzhe Di<sup>1</sup>, Yikun Liu<sup>1</sup>, Yudi Shi<sup>1</sup>,  
 Qirui Chen<sup>1</sup>, Zeqian Li<sup>1</sup>, Yifei Huang<sup>4</sup>, Weidi Xie<sup>1</sup>

<sup>1</sup>School of Artificial Intelligence, Shanghai Jiao Tong University

<sup>2</sup>Shanghai Innovation Institute    <sup>3</sup>Fudan University    <sup>4</sup>Shanghai AI Laboratory



Figure 1. **StreamFormer** learns streaming video representations of various granularities through multitask training, making it applicable for diverse downstream tasks such as Online Action Detection, Online Video Instance Segmentation and Video Question Answering.

## Abstract

Understanding continuous video streams plays a fundamental role in real-time applications, including embodied AI and autonomous driving. Unlike offline video processing, streaming video understanding requires the ability to process video streams frame by frame, preserve historical information, and make low-latency decisions. To address these challenges, our main contributions are three-fold. (i) we develop a novel streaming video backbone, termed as **StreamFormer**, by incorporating causal temporal attention into a pre-trained vision transformer. This enables efficient streaming video processing while maintaining image representation capability. (ii) to train **StreamFormer**, we propose to unify diverse spatial-temporal video understanding tasks within a multitask visual-language alignment framework. Hence, **StreamFormer** learns global semantics, temporal dynamics, and fine-grained spatial relationships simultaneously. (iii) we conduct extensive experiments for online action detection, online video instance segmentation, and

video question answering. **StreamFormer** achieves competitive performance while maintaining efficiency, demonstrating its potential for real-time applications.

## 1. Introduction

The growing demand for real-time applications, such as robotics [8, 28, 79], autonomous driving [22, 50, 73], and interactive systems [39, 45, 66], has driven a shift in video understanding research from traditional offline analysis [7, 35, 53, 102] to online streaming paradigms [17, 25, 46, 77, 121]. While offline video understanding focuses on analyzing pre-recorded videos with full temporal context, streaming video understanding addresses the challenge of dynamically processing continuous video streams in real time. As illustrated in Figure 1, models in this setting are required to sequentially process incoming frames, deliver timely predictions, and operate without access to future frames or the ability to frequently revisit distant past frames [10, 110, 129, 133]. This transition reflects the growing need for adaptive systems that prioritize latency, efficiency, and scalability in real-world scenarios.

\*Equal Contribution.

<sup>1</sup>Project Page: <https://go2heart.github.io/streamformer>

Comparing to traditional offline approaches, streaming video understanding poses several unique challenges: *First*, efficient frame-by-frame processing is crucial, as incoming visual signals must be handled promptly and continuously. Unlike existing offline video backbones [3, 35, 102], which typically process clips as their basic unit, streaming models must operate seamlessly at a per-frame basis. *Second*, long-term context preservation is essential, requiring models to effectively capture temporal dependencies across frames under strict memory and computational constraints. *Third*, low-latency decision-making is indispensable, ensuring predictions or actions to be promptly delivered for real-time scenarios. Addressing these challenges necessitates specialized backbone architectures specifically designed to balance efficiency, memory management, and accuracy.

Recent works adopt pre-trained image encoders to extract frame features [99, 124], which are subsequently processed auto-regressively by large language models (LLMs). Although this strategy aligns well with streaming video requirements, the reliance on aggregating coarse historical information within LLMs often leads to the loss of important fine-grained spatial and temporal details. To address these limitations, we propose training a streaming video backbone capable of supporting diverse spatiotemporal tasks, ranging from global scene understanding to fine-grained spatiotemporal event recognition. The proposed model, termed as **StreamFormer**, adapts Vision Transformer architectures pre-trained on images (*e.g.*, SigLIP [119]) by introducing causal attention along the temporal dimension at the spatial-temporal feature extraction stage. This approach allows the backbone to explicitly model temporal dependencies at early vision encoding stages, resulting in richer and more temporally coherent feature representations.

While training the streaming video backbone, we leverage various datasets that are manually annotated by the vision community [2, 11, 12, 37, 38, 107, 113], unifying various spatial-temporal video understanding tasks into a coherent multi-task learning framework. In particular, we combine three complementary learning objectives: (i) video-level classification for global semantics; (ii) per-frame supervision for temporal understanding, and (iii) pixel-wise segmentation for fine-grained spatial comprehension. By leveraging shared textual embeddings from SigLIP, our framework supports both closed-set and open-set tasks. This text-visual alignment mechanism facilitates knowledge transfer across tasks with varying granularities, enabling compatibility with diverse label spaces. In contrast to existing video backbone training paradigms [61, 102, 127] that predominantly rely on global contrastive learning using massive video-text datasets [6, 70, 103], our joint optimization of diverse objectives explicitly ensures capturing detailed temporal motion cues and spatial precision within streaming contexts.

In summary, our contributions are threefold: (i) we propose **StreamFormer**, which effectively converts pre-trained image backbones into streaming video models via temporal causal attention and spatial low-rank adaptations. Our approach achieves efficient streaming video processing without sacrificing spatial representational quality; (ii) we introduce a hierarchical multitask learning paradigm, systematically unifying global-level, temporal-level, and spatial-level learning objectives. This paradigm enables our model to simultaneously comprehend global semantics, temporal dynamics, and fine-grained spatial relationships within streaming video; (iii) we empirically demonstrate the effectiveness and efficiency of our method on diverse downstream tasks, including online action detection, online video instance segmentation, and video question answering. Our approach achieves promising performance across these tasks, effectively balancing accuracy and computational efficiency in streaming scenarios.

## 2. Related Work

**Video Backbone.** Early video understanding methods employ 3D convolutional neural networks [12, 34, 91] or factorized convolutions [92, 93, 97] to learn spatial-temporal features. Recent works have adopted Vision Transformers [27] to capture long-range dependencies in videos by incorporating global spatiotemporal self-attention [3], local self-attention [68], local-global self-attention [54], divided space-time attention [7], and multi-scale attention mechanisms [33, 57]. Despite the progress in offline video tasks, the use of bi-directional attention in visual backbones makes them less suitable for real-time video tasks. In NLP, LLMs typically adopt a decoder-only architecture [4, 62, 89, 90], where causal attention enables efficient autoregressive generation during inference. Currently, autoregressive modeling in vision primarily focuses on training backbones for image or video generation [5, 75, 85, 87, 94]. For visual understanding tasks, iGPT [18] learns visual representations by predicting next pixels, while AIM [30] explores the scaling laws for autoregressive image pre-training. Toto [81] trains a causal video transformer to predict the next visual patch. Unlike these works, we target designing StreamFormer with causal temporal attention to address various video understanding tasks.

**Video Language Pre-training.** Building on the success of image-text pretraining [41, 78, 119], recent video-language models [1, 102, 104, 108, 111, 127, 130] learn generalizable representations by training on massive video-text pairs [6, 103, 118]. UMT [53] and InternVideo2 [104] combine feature distilling, masked video modeling and video-text contrastive learning, to learn unimodal and multimodal representations. VideoPrism [127] integrates global-local distillation and token shuffling to enhance video represen-

tation learning. Despite strong performance, these methods typically require extensive training on million-scale video-text pairs. In contrast, we explore leveraging the rich, fine-grained human annotations available in existing video datasets for video language training by adopting a multitask learning paradigm, achieving higher data efficiency.

**Streaming Video Understanding.** Streaming video understanding tasks necessitate the processing of video inputs on a frame-by-frame basis and demand timely reactions [40, 45]. Existing streaming video approaches mainly target at online action detection [16, 96, 109, 110, 128], online video instance segmentation [105, 116, 123] and dense video captioning [133]. Leveraging the autoregressive paradigm in LLMs, recent approaches extract frame-wise features using pre-trained image encoders and pass them to LLMs for online question answering [25, 46, 121]. However, aggregating coarse frame features often results in the loss of fine-grained spatial and temporal details. In contrast, StreamFormer addresses causal temporal interaction in the video backbone to preserve spatial-temporal details.

**Multitask Learning.** Multitask models provide different outputs for a given input, simultaneously supporting various tasks [13, 15, 29, 49, 71, 83, 117]. Motivated by the success of unifying diverse tasks as sequence-to-sequence modeling in LLMs [9, 23, 80], recent works either train a unified multimodal and multitask transformer model [19, 69, 86, 98, 120, 134], or leverage LLM as a bridge to combine task-specific heads [100, 106, 112]. Inspired by these works, we train our streaming video backbone in a unified multitask visual-language alignment framework.

### 3. Method

We begin by formulating the problem in Sec. 3.1. Next, Sec. 3.2 describes the streaming video encoding backbone, followed by the language encoding model in Sec. 3.3. Finally, in Sec. 3.4, we elaborate on training the streaming video backbone to address diverse video understanding tasks within a unified framework, as illustrated in Figure 2.

#### 3.1. Problem Formulation

This paper considers the task of streaming video understanding, which aims to continuously process video frames by extracting spatiotemporal features in real-time, while retaining historical context and supporting diverse video understanding tasks. To achieve this, we propose to train a unified streaming video backbone capable of capturing global semantics, temporal dependencies, and fine-grained spatial details, all within a single framework.

Specifically, given a streaming video, represented as  $\mathcal{V}^T = [I_1, I_2, \dots, I_T]$ , where  $T$  denotes the total number of frames in the video. At any time step  $t$  ( $1 \leq t \leq T$ ), the

visual features of different granularities can be computed as:

$$\{\mathbf{v}^t, \mathbf{f}^t, \mathbf{F}^t\} = f_{\text{video}}(\mathcal{V}^t), \quad (1)$$

where  $\mathbf{v}^t \in \mathbb{R}^d$ ,  $\mathbf{f}^t \in \mathbb{R}^{t \times d}$ , and  $\mathbf{F}^t \in \mathbb{R}^{t \times h \times w \times d}$  denote the global, temporal, and spatial video features, respectively. The computation details are provided in Sec. 3.2.

To unify the diverse spatiotemporal tasks in a streaming setting, we reformulate them as visual-language alignment. Given a set of  $N$  tasks, denoted as  $\mathcal{Q} = \{q_1, q_2, \dots, q_N\}$ , the output for each task  $q_i$  (e.g., action recognition, temporal localization, or object segmentation) can be viewed as an alignment score  $\mathcal{A}_i$  in a shared visual-language space:

$$\mathcal{A}_i = g_{\text{tasks}}(f_{\text{video}}(\mathcal{V}^t), f_{\text{text}}(y), q_i), \quad (2)$$

where  $g_{\text{tasks}}(\cdot)$  produces the video-text alignment score  $\mathcal{A}_i$  given a task  $q_i$ , and  $f_{\text{text}}(y)$  is the text representation for a given action prompt or free-form video caption  $y$  (detailed in Sec. 3.3). In this way, training multiple tasks can be unified as optimizing the alignment score  $\mathcal{A}$ . As far as we know, this is the first work to develop a streaming video backbone that supports multiple spatiotemporal tasks under a unified framework, making it scalable and feasible for tackling a wide range of video understanding tasks.

**Discussion.** Reformulating various video understanding tasks as visual-language alignment offers several key advantages: (i) it eliminates the need for task-specific architectures by unifying all tasks as interactions between visual and textual features, enabling a streamlined and consistent framework; (ii) it facilitates multitask learning, allowing the model to generalize across task boundaries by leveraging shared spatiotemporal representations. This improves data efficiency and reduces training overhead; (iii) the framework is inherently extensible, supporting the incorporation of new tasks by simply defining additional alignment mechanisms, without requiring modifications to the underlying video backbone. By leveraging this unified framework, our method provides a flexible and efficient solution for streaming video understanding across varying complexity and granularity, including video semantics, temporal dependencies, and fine-grained spatial reasoning.

#### 3.2. Video Encoding

For the input video  $\mathcal{V}^T \in \mathbb{R}^{T \times H \times W \times 3}$ , it is first transformed into patch embeddings, with spatial and temporal position embeddings added separately, and the resulting video representation  $z \in \mathbb{R}^{T \times h \times w \times d}$  is the input to the video backbone. We adopt a series of Transformer blocks that employ divided space-time attention [7], with a key modification to make the temporal attention causal.

Specifically, while computing the visual feature for the current frame, the causal temporal attention operates on the




Figure 2. **Overall framework of StreamFormer.** Our StreamFormer is trained under a unified visual-language alignment framework, enabling simultaneous understanding of global semantics, temporal dynamics, and fine-grained spatial relationships. Each level utilizes features of different granularities: (i) last frame for the global level, (ii) per frame for the temporal level, (iii) and per frame per patch feature for the spatial level.

**same spatial positions** across **all historical frames** observed so far. The output for the  $l^{\text{th}}$  layer is denoted as:

$$z^{l+1} = f_{\text{FFN}}(f_{\text{space}}(f_{\text{time}}(z^l))) \quad (3)$$

where  $f_{\text{time}}(\cdot)$ ,  $f_{\text{space}}(\cdot)$  and  $f_{\text{FFN}}(\cdot)$  refer to causal temporal attention, spatial attention, and feed-forward network, respectively. Next, we illustrate  $f_{\text{time}}(\cdot)$  and  $f_{\text{space}}(\cdot)$  in detail.

**Temporal Causal Attention.** In contrast to the conventional bi-directional temporal attention used in video understanding backbones [3, 7], causal temporal attention models the temporal correlations between frames in a streaming setting, where frames are processed on-the-fly as they arrive, and each frame is restricted to only access features from preceding frames, *i.e.*, future frames remain inaccessible. This constraint is enforced by applying a temporal causal mask  $M \in \mathbb{R}^{T \times T}$  within the attention mechanism. At the  $l^{\text{th}}$  layer, the causal temporal attention is formulated as:

$$x^l = f_{\text{time}}(z^l) = \text{softmax} \left( \frac{QK^T}{\sqrt{d}} + M \right) V + z^l, \quad (4)$$

$$M_{ij} = \begin{cases} 0, & \text{if } i \geq j, \\ -\infty, & \text{otherwise,} \end{cases}$$

where  $Q, K$  and  $V$  are linear transformations of  $z$ , and  $\text{softmax}$  is operated along the temporal dimension. This design enables efficient processing of videos in a streaming setup, while also preserving temporal dependencies.

**Spatial Attention.** The spatial attention layer takes the output of the causal temporal attention layer and models the spatial relationship within each video frame. Following prior works [3, 7, 56, 74], we initialize the spatial attention layers from a pre-trained image Vision Transformer (*e.g.* SigLIP [119]). To preserve its strong feature representation while enabling video-specific adaptation, we employ

low-rank adaptation (LoRA) [43] strategy, which only introduces a small number of trainable parameters to modulate the pre-trained spatial weights effectively. The spatial attention is represented by:

$$\hat{x}^l = \text{softmax} \left( \frac{Q'K'^T}{\sqrt{d}} \right) V' + x^l. \quad (5)$$

$$Q' = W_Q^T x^l + W_{QB}^T (W_{QA}^T x^l)$$

where the learnable LoRA parameters  $W_{QA} \in \mathbb{R}^{d \times r}$  and  $W_{QB} \in \mathbb{R}^{r \times d}$  ( $r < d$ );  $K'$  and  $V'$  are transformed accordingly.  $\text{softmax}$  is operated along the spatial dimension.

**Outputs.** StreamFormer consists of 12 Transformer blocks, and outputs features of various resolutions: (1) spatiotemporal video features  $\mathbf{F} \in \mathbb{R}^{T \times h \times w \times d}$  by taking the last layer's hidden states, (2) temporal feature  $\mathbf{f} \in \mathbb{R}^{T \times d}$  by applying an attention pooling layer to  $\mathbf{F}$  over the spatial dimension following SigLIP [119], (3) global video feature  $\mathbf{v} \in \mathbb{R}^d$  by taking the last frame feature. All visual features are projected into a shared visual-text representation space using a visual projector  $\phi_{\text{visual}}(\cdot)$ . To avoid notation abuse, we reuse  $\mathbf{v}$ ,  $\mathbf{f}$  and  $\mathbf{F}$  to represent projected features.

### 3.3. Language Encoding

The groundtruth annotations for all considered training tasks encompass two forms: (1) **closed-set category**: an action or object label in the video (*e.g.*, kicking a football), and (2) **open-ended caption**: a natural language description of the video (*e.g.*, a man opens the window).

As in the former case, we construct prompt templates to convert the category labels into natural language descriptions, *e.g.*, a video clip of *kicking a football*. Details on the prompt templates can be found in the supplementary material. In this way, we can reformulate all groundtruth annotations into free-form languages ( $y$ ), and compute the text embeddings ( $\mathbf{t} \in \mathbb{R}^d$ ), *i.e.*,  $\mathbf{t} = f_{\text{text}}(y)$ , where  $f_{\text{text}}(\cdot)$  denotes



the pre-trained text encoder from SigLIP. This visual-text alignment facilitates knowledge transfer across tasks with varying granularities, enabling compatibility with diverse label spaces.

### 3.4. Spatiotemporal Multitasking

As for training StreamFormer, we unify a wide range of spatiotemporal tasks under a cohesive multitask learning framework. Unlike existing training paradigms that primarily use global contrastive learning on large-scale video-text pairs, we exploit the publicly available datasets from the vision community and focus on learning fine-grained spatiotemporal representations that capture both temporal dynamics and spatial precision. Specifically, we optimize the backbone with three complementary objectives: (i) capturing global semantics through video-level global supervision, (ii) modeling temporal dynamics with frame-level supervision, and (iii) enabling spatial precision via pixel-wise supervision.

**Global-level training.** The objective here is to learn the global semantics of a video clip, encompassing both the primary action and the scene. To achieve this, the global-level training consists of two sub-tasks: action recognition (AR) and video-text retrieval (VTR).

In action recognition, the model assigns videos to a closed set of action labels, while in video-text retrieval, the model aligns videos with free-form natural language descriptions. The language encoding procedure has been detailed in Sec. 3.3. To preserve causal temporal modeling, we use the feature of the last frame as the video representation. This design ensures that the model accumulates the temporal information from previous frames to construct a comprehensive global video representation. For consistency with SigLIP’s training process, we use sigmoid loss [119]:

$$\mathcal{L}_{AR} = -\frac{1}{|\mathcal{B}|} \sum_{i=1}^{|\mathcal{B}|} \sum_{c=1}^{|\mathcal{C}|} \log \frac{1}{1 + e^{y_{ic}(-\tau \mathbf{v}_i \cdot \mathbf{t}_c + b)}} \quad (6)$$

where  $y_{ic}$  is the label for video  $\mathbf{v}_i$  and text  $\mathbf{t}_c$ , which equals 1 if aligned and -1 otherwise. A learnable temperature  $\tau$  and bias  $b$  are used to alleviate initial imbalance.

For the video-text retrieval loss  $\mathcal{L}_{VTR}$ , we simply replace the prompted action label with a free-form language description, and calculate the sigmoid loss with negative pairs sampled from the mini-batch [119].

**Temporal-level training.** The objective here is to develop fine-grained discriminative capabilities along the temporal dimension, enabling the model to perform tasks such as per-frame action understanding and perceiving events that occur across frames. To achieve this, we consider two sub-tasks: temporal action localization (TAL) and temporal video grounding (TVG).

In temporal action localization, the model assigns each frame to either one of the pre-defined actions or background, while in temporal video grounding, the model enables to localize video segments corresponding to free-form natural language queries. Both sub-tasks leverage the temporal structure of videos, encouraging the model to capture temporal dependencies and align them with semantic textual representations. The language encoding procedure has been detailed in Sec. 3.3. Formally, the temporal-level loss takes temporal feature  $\mathbf{f}_i = [\mathbf{f}_i^1, \mathbf{f}_i^2, \dots, \mathbf{f}_i^T] \in \mathbb{R}^{T \times d}$ , where  $\mathbf{f}_i^t \in \mathbb{R}^d$  is the frame representation at timestep  $t$ . For temporal action localization with action classes  $\mathcal{C}$ :

$$\mathcal{L}_{TAL} = -\frac{1}{|\mathcal{B}|T} \sum_{i=1}^{|\mathcal{B}|} \sum_{t=1}^T \sum_{c=1}^{|\mathcal{C}|} \log \frac{1}{1 + e^{y_{ic}^t(-\tau \mathbf{f}_i^t \cdot \mathbf{t}_c + b)}}, \quad (7)$$

where  $y_{ic}^t = 1$  if frame  $t$  contains action class  $c$  and  $-1$  otherwise.

For temporal video grounding, we also formulate it as a contrastive learning paradigm. Specifically, when a certain frame in the video falls within the interval corresponding to its paired text query, we consider the pair as a positive sample. Negative samples include: (1) the current frame paired with other language samples in the batch; (2) background frames paired with all language samples in the batch. We calculate  $\mathcal{L}_{TVG}$  following Eq. 7, by replacing the action label with a free-form query.

**Spatial-level training.** Moving to the finest granularity, our spatial-level pre-training tackles pixel-wise comprehension through two sub-tasks: **video object segmentation (VOS)** for closed-set categories and **referring video object segmentation (RVOS)** for open-set language queries. VOS aims to predict a class label (including the background) for each pixel in the video, while RVOS needs pixel-level grounding given a natural language query.

To maintain the consistency of multitask learning, we directly compute dot product between each patch embedding with the corresponding textual embedding to get a coarse-grained segmentation map with logits. After up-sampling the patch-level logits to a higher resolution (e.g.  $224 \times 224$ ), a cross-entropy loss is utilized to optimize our predictions to the annotated masks. Formally, the spatial-level loss takes the spatiotemporal feature  $\mathbf{F}_i = [\mathbf{F}_i^1, \mathbf{F}_i^2, \dots, \mathbf{F}_i^{T \times H \times W}] \in \mathbb{R}^{(T \times H \times W) \times d}$ , where  $\mathbf{F}_i^{st} \in \mathbb{R}^d$  refers to the pixel-level representation at time  $t$  and spatial location  $s$ . For closed-set video object segmentation with classes  $\mathcal{C}$ , the loss is defined as:

$$\mathcal{L}_{VOS} = -\frac{1}{|\mathcal{B}|T \times H \times W} \sum_{i=1}^{|\mathcal{B}|} \sum_{st=1}^{T \times H \times W} \sum_{c=1}^{|\mathcal{C}|} \log \frac{1}{1 + e^{y_{ic}^{st}(-\tau \mathbf{F}_i^{st} \cdot \mathbf{t}_c + b)}}, \quad (8)$$

where  $y_{ic}^{st} = 1$  if the pixel at spatiotemporal location  $st$  belongs to class  $c$  and  $-1$  otherwise.

For referring video object segmentation, we extend the contrastive learning formulation in temporal video grounding to the spatiotemporal domain. Specifically, if a pixel in the video belongs to the object corresponding to the paired natural language query, we consider this pixel and the text query as a positive sample pair. Negative pairs include: (1) the current foreground pixel and other text queries within the batch, and (2) background pixels and all text queries within the batch. Again, we compute the loss  $\mathcal{L}_{RVOS}$  using the same formulation as Eq. 8, only changing the object label to the free-form language query.

**Training objective.** To jointly train the StreamFormer with global-level, temporal-level, and spatial-level tasks, we define the overall objective as:

$$\mathcal{L} = \underbrace{(\mathcal{L}_{AR} + \mathcal{L}_{VTR})}_{\text{global-level}} + \underbrace{(\mathcal{L}_{TAL} + \mathcal{L}_{TVG})}_{\text{temporal-level}} + \underbrace{(\mathcal{L}_{VOS} + \mathcal{L}_{RVOS})}_{\text{spatial-level}}$$

At training time, we alternately sample data from each task for a forward pass, and use gradient accumulation to perform backpropagation and parameter update collectively after iterating through all tasks, avoiding potential bias towards a specific task. Joint optimization of these objectives allows our model to learn unified video representations capable of capturing global semantics, temporal dynamics and spatial details, all within a single backbone.

## 4. Experiments

In this section, we first present datasets used for multitask training (Sec. 4.1), followed by the implementation details (Sec. 4.2). We present results on different downstream tasks for evaluation (Sec. 4.3) and ablation studies (Sec. 4.4).

### 4.1. Pre-training Data

We leverage various publicly available video datasets for training, encompassing annotations across various granularities. The datasets are categorized based on the task, including action recognition (AR) [12, 38], video-text retrieval (VTR) [2, 11, 14, 88, 101, 107], temporal action localization (TAL) [11, 67, 126], temporal video grounding (TVG) [2, 37, 51, 52, 72, 82], video object segmentation (VOS) [60, 95, 113], referring video object segmentation (RVOS) [26, 84]. The total number of video-text pairs used for pre-training is  $\sim 1\text{M}$ .

### 4.2. Implementation Details

We initialize the spatial attention layers using pre-trained weights from SigLIP [119] with LoRA rank 32, while the temporal attention layers are randomly initialized. To facilitate smooth integration of temporal information, we incorporate a learnable  $\tanh$  gate (initialized to zero) in the temporal attention. This design ensures the model initially preserves SigLIP’s strong spatial encoding capabilities,

and gradually learns temporal dependencies. At training time, we perform over-sampling to balance the data scale of different tasks. In all experiments, we use the base model (siglip-base-patch16-224). We train our model for 20 epochs, with global batch size 128, and a learning rate of  $1e^{-6}$ . By default, we uniformly sample 16 frames from the video as input. After training StreamFormer jointly on multiple tasks, we transfer the model to downstream video tasks by fixing the backbone and training task-specific heads.

### 4.3. Comparison on Downstream Tasks

We freeze StreamFormer and evaluate it on multiple downstream tasks, including online action detection, online video instance segmentation, and video question answering.

#### 4.3.1. Online Action Detection

Given an input video stream, online action detection (OAD) requires the model to classify each frame of the video. Notably, the model can only access the current and historical frames, without seeing the future frames. This characteristic inherently makes OAD a suitable benchmark for evaluating causal video models. We choose MAT [96] as our baseline OAD model, by replacing the TSN feature [97] with our StreamFormer feature. To ensure a fair comparison, we also incorporate optical flow features extracted using the BN-Inception model [47], which are concatenated with visual features and fed into the OAD model. The evaluation metric is per-frame *mean Average Precision (mAP)* on THUMOS-14 dataset [48].

As shown in Table 1, our model achieves state-of-the-art results, demonstrating a significant improvement over the baseline SigLIP without using optical flow (58.9 vs. 68.4), and even achieves comparable performance with early approaches that use flow. Given the time-consuming nature of extracting optical flow [16, 110], our approach clearly offers more efficient online inference without relying on flow, while still ensuring a high accuracy in action detection.

#### 4.3.2. Online Video Instance Segmentation

The goal of online video instance segmentation is the simultaneous detection, segmentation, and tracking of instances in videos, in a frame-by-frame manner. With our causal temporal attention, we can evaluate our video backbone model like an image backbone. We choose CTVIS [116] as our baseline, and keep all hyperparameters identical while changing the backbone into frozen StreamFormer.

Due to the requirement for dense features in this task, we add ViT-Adapter [20, 123] to our backbone. Note that we *do not* initialize model parameters with pre-training on COCO [60], and instead directly train YoutubeVIS-2019 [113] with COCO joint training. We report the Average Precision (AP) and Average Recall on YoutubeVIS-2019 benchmark.

Method	Backbone	THUMOS-14	
		w/ Flow	w/o Flow
TRN [109]	ResNet50	62.1	-
OADTR [42]	ResNet50	65.2	-
Colar [114]	ResNet50	66.9	-
LSTR [110]	ResNet50	69.5	63.5
GateHUB [16]	ResNet50	70.7	66.5
TesTra [129]	ResNet50	71.2	-
TRN [109]	TimeSformer	68.5	-
LSTR [110]	TimeSformer	69.6	-
GateHUB [16]	TimeSformer	72.5	-
MAT [96]	ResNet50	71.6	59.4
MAT [96]	SigLIP ViT	71.1	58.9
MAT [96]	<b>StreamFormer</b>	<b>73.9</b>	<b>68.4</b>

Table 1. Comparison on Online Action Detection.

Method	Res.	Backbone	COCO-init.	YTVIS19	
				AP	AR <sub>1</sub>
CrossVIS [115]	360p	ResNet50	✓	33.7	35.4
MinVIS [44]	360p	ResNet50	✓	47.4	45.7
IDOL [105]	360p	ResNet50	✓	49.5	47.7
DVIS++ [123]	480p	ResNet50	✓	55.5	51.1
CTVIS [116]	360p	ResNet50	✓	55.1	51.9
CTVIS [116]	224p	ResNet50	✗	43.6	42.7
CTVIS [116]	224p	SigLIP ViT	✗	40.9	41.8
CTVIS [116]	224p	<b>StreamFormer</b>	✗	<b>45.1</b>	<b>43.6</b>

Table 2. Comparison on Online Video Instance Segmentation.

Method	Res.	VideoMME	MLVU
LLaMA-VID [58]	224 <sup>2</sup>	25.9	18.1
Video-LLaVA [59]	224 <sup>2</sup>	39.9	18.8
Video-Chat2 [55]	224 <sup>2</sup>	42.3	44.5
LLaVA-NeXT-Video [124]	336 <sup>2</sup>	41.0	36.9
VideoLLaMA2 [21]	336 <sup>2</sup>	<b>47.9</b>	<b>48.4</b>
LLaVA-Next (StreamFormer)	224 <sup>2</sup>	45.0	<b>48.4</b>

Table 3. Comparison on VideoQA. All language models are 7B.

As illustrated in Table 2, StreamFormer outperforms the baseline approaches (*i.e.* ResNet50 and ViT) under the same experimental setups, while approaching prior works that adopt advanced training strategies (*e.g.* high-resolution training or using mask head pre-trained on COCO).

#### 4.3.3. Video Question Answering

Building upon the frameworks established by LLaVA [63] and LLaVA-Next-Video [124], we propose to integrate StreamFormer with large language models for video question answering. For a fair comparison, we replace the CLIP in LLaVA-Next (Vicuna-7b) with StreamFormer (frozen in

all training stages), and keep the remaining architecture and training pipeline unchanged. We use the same image-text pairs in LLaVA-NeXT, and sample 80K video-text pairs from LLaVA-Video-178K [125]. We provide our full training details in the supplementary materials.

Table 3 presents the comparison results on VideoMME [36] and MLVU [132]. StreamFormer exhibits superior performance among methods that use 224×224 resolution, indicating that enhancing temporal motion and spatial detail perception in video backbone enables better VideoQA.

#### 4.4. Ablation Study

In this part, we study the impact of different pre-training tasks and investigate the effects of multiple design choices. Apart from the task ablations (Table 4), we conduct comparative analysis experiments using 10% data by default.

**Ablation study for pre-training tasks.** Table 4 shows the model’s downstream performance when pre-trained on different task combinations. Our baseline is the original image encoder of SigLIP, *i.e.*, without multitask training and temporal modeling. We draw the following conclusions: (1) the baseline achieves relatively good performance on downstream tasks, highlighting the strong image representation capability of SigLIP; (2) the incorporation of any form of video-language pre-training, either global, temporal, or spatial, yields a notable improvement over the baseline on all downstream tasks; (3) both global- and temporal-level tasks yield better results across various downstream tasks, while models trained solely on spatial-level tasks perform sub-optimally in online action detection and video question answering. This may be attributed to the relatively limited amount of data available for video segmentation, leading to an insufficient understanding of verbs/actions; (4) by integrating tasks of varying granularities, the model achieves competitive downstream performance, proving the effectiveness of our multitask training approach.

**Effect of temporal causal modeling.** We investigate the effect of temporal causal attention (Causal TA) by comparing it with two baselines: (1) an image-based model that only employs spatial attention combined with LoRA, without any temporal attention (No TA); (2) a model utilizes bi-directional temporal attention (Bi-directional TA). We report the performance of the models without optical flow to verify the temporal effect of visual features. As observed from Table 5, temporal attention matters in online action detection, while bi-directional attention performs slightly better than causal attention. This is reasonable as each frame can access all frames in bi-directional attention.

However, one notable advantage of causal attention over bi-directional attention is the efficiency at inference time in the streaming setting, particularly when supported by KV-Cache. We compare the inference cost of the backbone (*i.e.*,

Method	Global	Temporal	Spatial	Online Action Detection		Online VIS		VideoQA	
				THUMOS-14		YouTubeVIS-2019		VideoMME	MLVU
				w/ Flow	w/o Flow	AP	AR <sub>1</sub>		
SigLIP				71.1	58.9	40.9	41.8	40.8	46.1
StreamFormer	✓			73.1	65.9	41.9	41.6	43.7	44.6
StreamFormer		✓		72.1	65.5	40.1	41.1	44.0	<b>49.6</b>
StreamFormer			✓	70.3	51.5	42.3	41.6	30.0	29.2
StreamFormer	✓	✓		<b>74.1</b>	<b>69.1</b>	41.9	41.4	43.1	49.2
StreamFormer	✓	✓	✓	73.9	68.4	<b>45.1</b>	<b>43.6</b>	<b>45.0</b>	48.4

Table 4. Ablation study for pre-training tasks.

Method	No TA	Bi-directional TA	Causal TA
w/o Flow	61.7	62.8	62.5

Table 5. Ablation study for temporal attention.




Figure 3. Computational complexity.

latency and GPU memory) in Figure 3, with input frames varying from 1 to 4096. The latency of causal attention (w/ KV-Cache) remains nearly constant as input frame increases, whereas bi-directional attention experiences a significant rise in computational cost. This highlights the scalability of StreamFormer in streaming video processing.

**Data efficiency.** In this work, we leverage existing video benchmark datasets, including rich annotations of various granularities for training StreamFormer. In contrast, most existing video foundation models [61, 102, 104, 111, 127] are trained on extensive (usually >50M) video-text pairs [6, 70, 103], which is prohibitively expensive. For comparison, we also train our backbone on the WebVid [6] dataset, using the video-language contrastive loss.

We randomly select 1M video-text pairs from WebVid-10M, which is comparable to the scale of our pre-training data. As shown in Table 6, the models trained on WebVid-1M exhibit relatively low performance, possibly due to insufficient pre-training data for video-text contrastive learning. In comparison, our approach outperforms WebVid-1M model using only 0.1M pre-training data, significantly reducing the computational cost for training. Our method effectively leverages the existing high-quality annotated video datasets to enhance the backbone’s capabilities, which suggests there is certainly a trade-off between computation and annotation cost.

Dataset	OAD		VideoQA	
	w/ Flow	w/o Flow	VideoMME	MLVU
WebVid-1M	69.3	57.2	37.4	46.6
Ours-0.1M	73.0	62.5	43.1	<b>49.6</b>
Ours-1M	<b>73.9</b>	<b>68.4</b>	<b>45.0</b>	48.4

Table 6. Comparison on data efficiency.

## 5. Conclusions

In this paper, we present StreamFormer, a streaming video backbone that incorporates causal temporal attention and spatial low-rank adaptation into a pre-trained SigLIP image encoder. To facilitate the training of StreamFormer, we devise a multitask visual-language alignment framework that unifies diverse video understanding tasks, enabling it to learn global-level semantics, temporal-level dynamics and spatial-level perception simultaneously. We conduct extensive experiments on online action detection, online video instance segmentation and video question answering. StreamFormer exhibits promising online video understanding ability while preserving high efficiency, demonstrating its huge potential in future real-time applications.

## Acknowledgement

Weidi would like to acknowledge the funding from Scientific Research Innovation Capability Support Project for Young Faculty (ZY-GXQNJSKYCXNLZCXM-I22).

## References

- [1] Hassan Akbari, Liangzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong. Vatt: Transformers for multimodal self-supervised learning from raw video, audio and text. In *NeurIPS*, 2021. 2
- [2] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *ICCV*, 2017. 2, 6, 1
- [3] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *ICCV*, 2021. 2, 4



- [4] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023. 2
- [5] Yutong Bai, Xinyang Geng, Karttikeya Mangalam, Amir Bar, Alan L Yuille, Trevor Darrell, Jitendra Malik, and Alexei A Efros. Sequential modeling enables scalable learning for large vision models. In *CVPR*, 2024. 2
- [6] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, 2021. 2, 8
- [7] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, 2021. 1, 2, 3, 4
- [8] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022. 1
- [9] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *NeurIPS*, 2020. 3
- [10] Shyamal Buch, Victor Escorcia, Chuanqi Shen, Bernard Ghanem, and Juan Carlos Niebles. Sst: Single-stream temporal action proposals. In *CVPR*, 2017. 1
- [11] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, 2015. 2, 6, 1
- [12] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017. 2, 6, 1
- [13] Rich Caruana. Multitask learning. *Machine learning*, 1997. 3
- [14] David Chen and William B Dolan. Collecting highly parallel data for paraphrase evaluation. In *ACL*, 2011. 6, 1
- [15] Guo Chen, Yin-Dong Zheng, Jiahao Wang, Jilan Xu, Yifei Huang, Junting Pan, Yi Wang, Yali Wang, Yu Qiao, Tong Lu, et al. Videollm: Modeling video sequence with large language models. *arXiv preprint arXiv:2305.13292*, 2023. 3
- [16] Junwen Chen, Gaurav Mittal, Ye Yu, Yu Kong, and Mei Chen. Gatehub: Gated history unit with background suppression for online action detection. In *CVPR*, 2022. 3, 6, 7
- [17] Joya Chen, Zhaoyang Lv, Shiwei Wu, Kevin Qinghong Lin, Chenan Song, Difei Gao, Jia-Wei Liu, Ziteng Gao, Dongxing Mao, and Mike Zheng Shou. Videollm-online: Online video large language model for streaming video. In *CVPR*, 2024. 1
- [18] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pre-training from pixels. In *ICML*, 2020. 2
- [19] Ting Chen, Saurabh Saxena, Lala Li, David J Fleet, and Geoffrey Hinton. Pix2seq: A language modeling framework for object detection. In *ICLR*, 2022. 3
- [20] Zhe Chen, Yuchen Duan, Wenhai Wang, Junjun He, Tong Lu, Jifeng Dai, and Yu Qiao. Vision transformer adapter for dense predictions. In *ICLR*, 2023. 6
- [21] Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi Zhang, Ziyang Luo, Deli Zhao, et al. Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-llms. *arXiv preprint arXiv:2406.07476*, 2024. 7
- [22] Kashyap Chitta, Aditya Prakash, Bernhard Jaeger, Zehao Yu, Katrin Renz, and Andreas Geiger. Transfuser: Imitation with transformer-based sensor fusion for autonomous driving. *TPAMI*, 2022. 1
- [23] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 2023. 3
- [24] Roeland De Geest, Efstratios Gavves, Amir Ghodrati, Zhenyang Li, Cees Snoek, and Tinne Tuytelaars. Online action detection. In *ECCV*, 2016. 2
- [25] Shangzhe Di, Zhelun Yu, Guanghao Zhang, Haoyuan Li, Hao Cheng, Bolin Li, Wanggui He, Fangxun Shu, Hao Jiang, et al. Streaming video question-answering with in-context video kv-cache retrieval. In *ICLR*, 2025. 1, 3
- [26] Henghui Ding, Chang Liu, Shuting He, Xudong Jiang, and Chen Change Loy. Mevis: A large-scale benchmark for video segmentation with motion expressions. In *ICCV*, 2023. 6, 1
- [27] 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 *ICLR*, 2021. 2
- [28] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Ayaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, et al. Palm-e: An embodied multimodal language model. 2023. 1
- [29] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *ICCV*, 2015. 3
- [30] Alaaeldin El-Nouby, Michal Klein, Shuangfei Zhai, Miguel Angel Bautista, Alexander Toshev, Vaishaal Shankar, Joshua M Susskind, and Armand Joulin. Scalable pre-training of large autoregressive image models. In *ICML*, 2024. 2
- [31] Junming Lin et al. Streamingbench: Assessing the gap for mllms to achieve streaming video understanding, 2024. 2
- [32] Viorica Pătrăucean et al. Trecvit: A recurrent video transformer, 2024. 2
- [33] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *ICCV*, 2021. 2
- [34] Christoph Feichtenhofer, Axel Pinz, and Richard P Wildes. Spatiotemporal multiplier networks for video action recognition. In *CVPR*, 2017. 2
- [35] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *CVPR*, 2019. 1, 2

- [36] Chaoyou Fu, Yuhao Dai, Yondong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, et al. Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis. In *CVPR*, 2025. 7
- [37] Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. Tall: Temporal activity localization via language query. In *ICCV*, 2017. 2, 6, 1
- [38] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The “something something” video database for learning and evaluating visual common sense. In *ICCV*, 2017. 2, 6, 1
- [39] Kristen Grauman, Andrew Westbury, Lorenzo Torresani, Kris Kitani, Jitendra Malik, Triantafyllos Afouras, Kumar Ashutosh, Vijay Baiyya, Siddhant Bansal, Bikram Boote, et al. Ego-exo4d: Understanding skilled human activity from first-and third-person perspectives. In *CVPR*, 2024. 1
- [40] Tengda Han, Dilara Gokay, Joseph Heyward, Chuhan Zhang, Daniel Zoran, Viorica Patraucean, Joao Carreira, Dima Damen, and Andrew Zisserman. Learning from streaming video with orthogonal gradients. In *CVPR*, 2025. 3
- [41] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. 2
- [42] Lukas Hedegaard, Arian Bakhtiarnia, and Alexandros Iosifidis. Continual transformers: redundancy-free attention for online inference. In *ICLR*, 2023. 7, 1
- [43] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *ICLR*, 2022. 4
- [44] De-An Huang, Zhiding Yu, and Anima Anandkumar. Minvis: A minimal video instance segmentation framework without video-based training. In *NeurIPS*, 2022. 7
- [45] Yifei Huang, Jilan Xu, Baoqi Pei, Yuping He, Guo Chen, Lijin Yang, Xinyuan Chen, Yaohui Wang, Zheng Nie, Jinyao Liu, et al. Vinci: A real-time embodied smart assistant based on egocentric vision-language model. *arXiv preprint arXiv:2412.21080*, 2024. 1, 3
- [46] Zhenpeng Huang, Xinhao Li, Jiaqi Li, Jing Wang, Xiangyu Zeng, Cheng Liang, Tao Wu, Xi Chen, Liang Li, and Limin Wang. Online video understanding: A comprehensive benchmark and memory-augmented method. *arXiv preprint arXiv:2501.00584*, 2024. 1, 3
- [47] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 6, 1
- [48] Y.-G. Jiang, J. Liu, A. Roshan Zamir, G. Toderici, I. Laptev, M. Shah, and R. Sukthankar. THUMOS challenge: Action recognition with a large number of classes. <http://csrcv.ucf.edu/THUMOS14/>, 2014. 6
- [49] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *CVPR*, 2018. 3
- [50] B Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Manion, Ahmad A Al Sallab, Senthil Yogamani, and Patrick Pérez. Deep reinforcement learning for autonomous driving: A survey. *TPAMI*, 2021. 1
- [51] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *ICCV*, 2017. 6, 1
- [52] Jie Lei, Tamara L Berg, and Mohit Bansal. Detecting moments and highlights in videos via natural language queries. In *NeurIPS*, 2021. 6, 1
- [53] Kunchang Li, Yali Wang, Yizhuo Li, Yi Wang, Yinan He, Limin Wang, and Yu Qiao. Unmasked teacher: Towards training-efficient video foundation models. In *ICCV*, 2023. 1, 2
- [54] Kunchang Li, Yali Wang, Junhao Zhang, Peng Gao, Guanglu Song, Yu Liu, Hongsheng Li, and Yu Qiao. Uniformer: Unifying convolution and self-attention for visual recognition. *TPAMI*, 2023. 2
- [55] Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, et al. Mvbench: A comprehensive multi-modal video understanding benchmark. In *CVPR*, 2024. 7
- [56] Xinhao Li, Yuhao Zhu, and Limin Wang. Zero2v: Zero-cost adaptation of pre-trained transformers from image to video. In *ECCV*, 2024. 4
- [57] Yanghao Li, Chao-Yuan Wu, Haoqi Fan, Karttikeya Mangalam, Bo Xiong, Jitendra Malik, and Christoph Feichtenhofer. Mv2: Improved multiscale vision transformers for classification and detection. In *CVPR*, 2022. 2
- [58] Yanwei Li, Chengyao Wang, and Jiaya Jia. Llama-vid: An image is worth 2 tokens in large language models. In *ECCV*, 2024. 7
- [59] Bin Lin, Yang Ye, Bin Zhu, Jiayi Cui, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning unified visual representation by alignment before projection. In *EMNLP*, 2023. 7
- [60] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 6, 1
- [61] Ziyi Lin, Shijie Geng, Renrui Zhang, Peng Gao, Gerard De Melo, Xiaogang Wang, Jifeng Dai, Yu Qiao, and Hongsheng Li. Frozen clip models are efficient video learners. In *ECCV*, 2022. 2, 8
- [62] Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*, 2024. 2
- [63] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023. 7
- [64] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *CVPR*, 2024. 1, 2
- [65] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge. <https://llava-vl.github.io/blog/2024-01-30-llava-next/>, 2024. 2

- [66] Jihao Liu, Zhiding Yu, Shiyi Lan, Shihao Wang, Rongyao Fang, Jan Kautz, Hongsheng Li, and Jose M Alvarez. Streamchat: Chatting with streaming video. *arXiv preprint arXiv:2412.08646*, 2024. 1
- [67] Yi Liu, Limin Wang, Yali Wang, Xiao Ma, and Yu Qiao. Fineaction: A fine-grained video dataset for temporal action localization. *TIP*, 2022. 6, 1
- [68] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *CVPR*, 2022. 2
- [69] Jiasen Lu, Christopher Clark, Rowan Zellers, Roozbeh Mottaghi, and Aniruddha Kembhavi. Unified-io: A unified model for vision, language, and multi-modal tasks. In *ICLR*, 2023. 3
- [70] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *ICCV*, 2019. 2, 8
- [71] Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. Cross-stitch networks for multi-task learning. In *CVPR*, 2016. 3
- [72] Andreea-Maria Oncescu, Joao F Henriques, Yang Liu, Andrew Zisserman, and Samuel Albanie. Queryd: A video dataset with high-quality text and audio narrations. In *ICASSP*, 2021. 6, 1
- [73] Chenbin Pan, Burhaneddin Yaman, Tommaso Nesti, Abhirup Mallik, Alessandro G Allievi, Senem Velipasalar, and Liu Ren. Vlp: Vision language planning for autonomous driving. In *CVPR*, 2024. 1
- [74] Junting Pan, Ziyi Lin, Xiatian Zhu, Jing Shao, and Hongsheng Li. St-adapter: Parameter-efficient image-to-video transfer learning. In *NeurIPS*, 2022. 4
- [75] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In *ICML*, 2018. 2
- [76] Jiyang Qi, Yan Gao, Yao Hu, Xinggang Wang, Xiaoyu Liu, Xiang Bai, Serge Belongie, Alan Yuille, Philip HS Torr, and Song Bai. Occluded video instance segmentation: A benchmark. In *IJCV*, 2022. 2
- [77] Rui Qian, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Shuangrui Ding, Dahua Lin, and Jiaqi Wang. Streaming long video understanding with large language models. In *NeurIPS*, 2024. 1
- [78] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 2, 1
- [79] Ilija Radosavovic, Tete Xiao, Stephen James, Pieter Abbeel, Jitendra Malik, and Trevor Darrell. Real-world robot learning with masked visual pre-training. In *CoRL*, 2023. 1
- [80] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*, 2020. 3
- [81] Jathushan Rajasegaran, Ilija Radosavovic, Rahul Ravishanker, Yossi Gandelsman, Christoph Feichtenhofer, and Jitendra Malik. An empirical study of autoregressive pre-training from videos. *arXiv preprint arXiv:2501.05453*, 2025. 2
- [82] Anna Rohrbach, Marcus Rohrbach, Wei Qiu, Annemarie Friedrich, Manfred Pinkal, and Bernt Schiele. Coherent multi-sentence video description with variable level of detail. In *GCPR*, 2014. 6, 1
- [83] Sebastian Ruder. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*, 2017. 3
- [84] Seonguk Seo, Joon-Young Lee, and Bohyung Han. Urvos: Unified referring video object segmentation network with a large-scale benchmark. In *ECCV*, 2020. 6, 1
- [85] Peize Sun, Yi Jiang, Shoufa Chen, Shilong Zhang, Bingyue Peng, Ping Luo, and Zehuan Yuan. Autoregressive model beats diffusion: Llama for scalable image generation. *arXiv preprint arXiv:2406.06525*, 2024. 2
- [86] Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*, 2024. 3
- [87] Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling: Scalable image generation via next-scale prediction. In *NeurIPS*, 2024. 2
- [88] Atousa Torabi, Niket Tandon, and Leonid Sigal. Learning language-visual embedding for movie understanding with natural-language. *arXiv preprint arXiv:1609.08124*, 2016. 6, 1
- [89] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 2
- [90] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. 2
- [91] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015. 2
- [92] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, 2018. 2
- [93] Du Tran, Heng Wang, Lorenzo Torresani, and Matt Feiszli. Video classification with channel-separated convolutional networks. In *ICCV*, 2019. 2
- [94] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixcnn decoders. In *NeurIPS*, 2016. 2
- [95] Haochen Wang, Cilin Yan, Shuai Wang, Xiaolong Jiang, Xu Tang, Yao Hu, Weidi Xie, and Efstratios Gavves. Towards open-vocabulary video instance segmentation. In *ICCV*, 2023. 6, 1

- [96] Jiahao Wang, Guo Chen, Yifei Huang, Limin Wang, and Tong Lu. Memory-and-anticipation transformer for online action understanding. In *ICCV*, 2023. 3, 6, 7, 1
- [97] 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 *ECCV*, 2016. 2, 6
- [98] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *ICML*, 2022. 3
- [99] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024. 2
- [100] Wenhai Wang, Zhe Chen, Xiaokang Chen, Jiannan Wu, Xizhou Zhu, Gang Zeng, Ping Luo, Tong Lu, Jie Zhou, Yu Qiao, et al. Visionllm: Large language model is also an open-ended decoder for vision-centric tasks. In *NeurIPS*, 2023. 3
- [101] Xin Wang, Jiawei Wu, Junkun Chen, Lei Li, Yuan-Fang Wang, and William Yang Wang. Vatex: A large-scale, high-quality multilingual dataset for video-and-language research. In *ICCV*, 2019. 6, 1
- [102] Yi Wang, Kunchang Li, Yizhuo Li, Yanan He, Bingkun Huang, Zhiyu Zhao, Hongjie Zhang, Jilan Xu, Yi Liu, Zun Wang, et al. Internvideo: General video foundation models via generative and discriminative learning. *arXiv preprint arXiv:2212.03191*, 2022. 1, 2, 8
- [103] Yi Wang, Yanan He, Yizhuo Li, Kunchang Li, Jiashuo Yu, Xin Ma, Xinhao Li, Guo Chen, Xinyuan Chen, Yaohui Wang, et al. Internvid: A large-scale video-text dataset for multimodal understanding and generation. In *ICLR*, 2024. 2, 8
- [104] Yi Wang, Kunchang Li, Xinhao Li, Jiashuo Yu, Yanan He, Guo Chen, Baoqi Pei, Rongkun Zheng, Zun Wang, Yansong Shi, et al. Internvideo2: Scaling foundation models for multimodal video understanding. In *ECCV*, 2024. 2, 8
- [105] Junfeng Wu, Qihao Liu, Yi Jiang, Song Bai, Alan Yuille, and Xiang Bai. In defense of online models for video instance segmentation. In *ECCV*, 2022. 3, 7
- [106] Jiannan Wu, Muyan Zhong, Sen Xing, Zeqiang Lai, Zhaoyang Liu, Zhe Chen, Wenhai Wang, Xizhou Zhu, Lewei Lu, Tong Lu, et al. Visionllm v2: An end-to-end generalist multimodal large language model for hundreds of vision-language tasks. In *NeurIPS*, 2024. 3
- [107] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*, 2016. 2, 6, 1
- [108] Jilan Xu, Yifei Huang, Junlin Hou, Guo Chen, Yuejie Zhang, Rui Feng, and Weidi Xie. Retrieval-augmented ego-centric video captioning. In *CVPR*, 2024. 2
- [109] Mingze Xu, Mingfei Gao, Yi-Ting Chen, Larry S Davis, and David J Crandall. Temporal recurrent networks for online action detection. In *ICCV*, 2019. 3, 7
- [110] Mingze Xu, Yuanjun Xiong, Hao Chen, Xinyu Li, Wei Xia, Zhuowen Tu, and Stefano Soatto. Long short-term transformer for online action detection. In *NeurIPS*, 2021. 1, 3, 6, 7
- [111] Shen Yan, Tao Zhu, Zirui Wang, Yuan Cao, Mi Zhang, Soham Ghosh, Yonghui Wu, and Jiahui Yu. Videococa: Video-text modeling with zero-shot transfer from contrastive captioners. *arXiv preprint arXiv:2212.04979*, 2022. 2, 8
- [112] Ziang Yan, Zhilin Li, Yanan He, Chenting Wang, Kunchang Li, Xinhao Li, Xiangyu Zeng, Zilei Wang, Yali Wang, Yu Qiao, et al. Task preference optimization: Improving multimodal large language models with vision task alignment. *arXiv preprint arXiv:2412.19326*, 2024. 3
- [113] Linjie Yang, Yuchen Fan, and Ning Xu. Video instance segmentation. In *ICCV*, 2019. 2, 6, 1
- [114] Le Yang, Junwei Han, and Dingwen Zhang. Colar: Effective and efficient online action detection by consulting exemplars. In *CVPR*, 2022. 7
- [115] Shusheng Yang, Yuxin Fang, Xinggang Wang, Yu Li, Chen Fang, Ying Shan, Bin Feng, and Wenyu Liu. Crossover learning for fast online video instance segmentation. In *ICCV*, 2021. 7
- [116] Kaining Ying, Qing Zhong, Weian Mao, Zhenhua Wang, Hao Chen, Lin Yuanbo Wu, Yifan Liu, Chengxiang Fan, Yunzhi Zhuge, and Chunhua Shen. Ctviz: Consistent training for online video instance segmentation. In *ICCV*, 2023. 3, 6, 7, 1
- [117] Amir R Zamir, Alexander Sax, William Shen, Leonidas J Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *CVPR*, 2018. 3
- [118] Rowan Zellers, Jiasen Lu, Ximing Lu, Youngjae Yu, Yanyan Zhao, Mohammadreza Salehi, Aditya Kusupati, Jack Hessel, Ali Farhadi, and Yejin Choi. Merlot reserve: Neural script knowledge through vision and language and sound. In *CVPR*, 2022. 2
- [119] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *ICCV*, 2023. 2, 4, 5, 6, 1
- [120] Jun Zhan, Junqi Dai, Jiasheng Ye, Yunhua Zhou, Dong Zhang, Zhigeng Liu, Xin Zhang, Ruibin Yuan, Ge Zhang, Linyang Li, et al. Anygpt: Unified multimodal llm with discrete sequence modeling. In *ACL*, 2024. 3
- [121] Haoji Zhang, Yiqin Wang, Yansong Tang, Yong Liu, Jiashi Feng, Jifeng Dai, and Xiaojie Jin. Flash-vstream: Memory-based real-time understanding for long video streams. *arXiv preprint arXiv:2406.08085*, 2024. 1, 3
- [122] Kaichen Zhang, Bo Li, Peiyuan Zhang, Fanyi Pu, Joshua Adrian Cahyono, Kairui Hu, Shuai Liu, Yuanhan Zhang, Jingkang Yang, Chunyuan Li, and Ziwei Liu. Lmms-eval: Reality check on the evaluation of large multimodal models, 2024. 2
- [123] Tao Zhang, Xingye Tian, Yikang Zhou, Shunping Ji, Xuebo Wang, Xin Tao, Yuan Zhang, Pengfei Wan, Zhongyuan Wang, and Yu Wu. Dvis++: Improved decoupled framework for universal video segmentation. *arXiv preprint arXiv:2312.13305*, 2023. 3, 6, 7



- [124] Yuanhan Zhang, Bo Li, haotian Liu, Yong jae Lee, Liangke Gui, Di Fu, Jiashi Feng, Ziwei Liu, and Chunyuan Li. Llava-next: A strong zero-shot video understanding model. <https://llava-vl.github.io/blog/2024-04-30-llava-next-video/>, 2024. 2, 7, 1
- [125] Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. Video instruction tuning with synthetic data. *arXiv preprint arXiv:2410.02713*, 2024. 7, 2
- [126] Hang Zhao, Antonio Torralba, Lorenzo Torresani, and Zhicheng Yan. Hacs: Human action clips and segments dataset for recognition and temporal localization. In *ICCV*, 2019. 6, 1
- [127] Long Zhao, Nitesh B Gundavarapu, Liangzhe Yuan, Hao Zhou, Shen Yan, Jennifer J Sun, Luke Friedman, Rui Qian, Tobias Weyand, Yue Zhao, et al. Videoprism: A foundational visual encoder for video understanding. In *ICML*, 2024. 2, 8
- [128] Qingsong Zhao, Yi Wang, Jilan Xu, Yinan He, Zifan Song, Limin Wang, Yu Qiao, Cairong Zhao, et al. Does video-text pretraining help open-vocabulary online action detection? *NeurIPS*, 2024. 3
- [129] Yue Zhao and Philipp Krähenbühl. Real-time online video detection with temporal smoothing transformers. In *ECCV*, 2022. 1, 7
- [130] Yue Zhao, Ishan Misra, Philipp Krähenbühl, and Rohit Girdhar. Learning video representations from large language models. In *CVPR*, 2023. 2
- [131] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. In *NeurIPS*, 2023. 1
- [132] Junjie Zhou, Yan Shu, Bo Zhao, Boya Wu, Shitao Xiao, Xi Yang, Yongping Xiong, Bo Zhang, Tiejun Huang, and Zheng Liu. Mlvu: A comprehensive benchmark for multi-task long video understanding. In *CVPR*, 2025. 7
- [133] Xingyi Zhou, Anurag Arnab, Shyamal Buch, Shen Yan, Austin Myers, Xuehan Xiong, Arsha Nagrani, and Cordelia Schmid. Streaming dense video captioning. In *CVPR*, 2024. 1, 3
- [134] Xizhou Zhu, Jinguo Zhu, Hao Li, Xiaoshi Wu, Hongsheng Li, Xiaohua Wang, and Jifeng Dai. Uni-perceiver: Pre-training unified architecture for generic perception for zero-shot and few-shot tasks. In *CVPR*, 2022. 3