

Scaling Open-Vocabulary Action Detection

Zhen Hao Sia Yogesh Singh Rawat

{zhenhao.sia, yogesh}@ucf.edu
 University of Central Florida

Abstract

In this work, we focus on scaling open-vocabulary action detection. Existing approaches for action detection are predominantly limited to closed-set scenarios and rely on complex, parameter-heavy architectures. Extending these models to the open-vocabulary setting poses two key challenges: (1) the lack of large-scale datasets with many action classes for robust training, and (2) parameter-heavy adaptations to a pretrained vision-language contrastive model to convert it for detection, risking overfitting the additional non-pretrained parameters to base action classes. Firstly, we introduce an encoder-only multimodal model for video action detection, reducing the reliance on parameter-heavy additions for video action detection. Secondly, we introduce a simple weakly supervised training strategy to exploit an existing closed-set action detection dataset for pre-training. Finally, we depart from the ill-posed base-to-novel benchmark used by prior works in open-vocabulary action detection and devise a new benchmark to evaluate on existing closed-set action detection datasets without ever using them for training, showing novel results to serve as baselines for future work. Our code is available at https://siatheindochinese.github.io/sia_act_page/.

1. Introduction

Spatiotemporal action detection has traditionally focused on the closed-set scenario, where models are trained on fully-supervised, predefined action classes and can only recognize actions encountered during training. While effective within controlled environments, this approach is restrictive for real-world applications, where human actions and interactions are inherently diverse and constantly evolving. From public surveillance and autonomous systems to sports analytics and assistive technologies, the range of possible actions is vast and unpredictable, often extending far beyond the limited set captured in any single dataset. Open-vocabulary action detection addresses this limitation by allowing models to detect and identify novel action classes unseen during




Figure 1. **Overview of existing approaches** to convert pretrained vision/video-language models for action detection (left, middle) vs ours (right).

training, providing a more scalable and adaptable solution for real-world scenarios.

Developing open-vocabulary models for action detection is challenging due to two main factors: (1) while vision-language modeling offers a promising approach to extend closed-set models [6, 10, 12, 33, 46, 51] for detecting novel actions, naively applying this to action detection remains challenging due to the scarcity of annotated datasets covering a large number of actions, as opposed to open-vocabulary object detection where large-scale datasets with sufficient number of object classes already exist [15], and (2) while existing work [3, 17] tries to alleviate this issue by adapting a pretrained, frozen vision/video-language contrastive model inside a detection architecture, having a significant increase in the number of non-pretrained modules/parameters leads to the risk that the additional parameters overfit to the base action classes, as well as incurring additional computation.

In this work, we address these two challenges by (1) proposing a weakly-supervised approach to significantly inflate the number of action classes seen during pretraining to deal with the scarcity of action classes, and (2) introducing a

single-stage, encoder-only, open-vocabulary action detection model that does not rely on parameter-heavy additions such as a decoder or an external human detector to avoid overfitting these new parameters to base action classes (Figure 1).

The abundance of action classes from our training scheme allows us to avoid freezing pretrained vision/video-language models, enabling our model to adapt more effectively to novel actions without sacrificing generalization. Additionally, by avoiding parameter-heavy modules attached to a frozen vision/video encoder, our approach remains lightweight, reducing computational overhead and making it more practical for real-world deployment, as well as allowing it to be trained end-to-end.

We perform extensive experiments on six different existing action detection datasets: AVA [13], AVA-Kinetics (AVA-K) [25], UCF-101-24 [39], JHMDB51-21 [19], MultiSports [26], and UCF-MAMA [30], demonstrating the open-vocabulary capability of our model.

In summary, we have the following contributions:

- We introduce SiA, a *simple architecture* for open-vocabulary action detection which is multimodal and lightweight.
- We introduce a weakly-supervised training scheme to exploit the largest existing closed-set action detection dataset by inflating the number of action classes seen during training from 80 to more than 700, a significant departure from prior work that has only seen less than 20 actions during pretraining.
- We depart from the ill-posed base-to-novel benchmark used by prior works in zero-shot/open-vocabulary action detection and set a new benchmark for the task of open-vocabulary action detection by showing open-vocabulary results on UCF-101-24, JHMDB, MultiSports and UCF-MAMA without training our model on these datasets, setting a novel baseline.

2. Related Work

Spatio-Temporal Action Detection fall under two categories: frame-level [44] and clip-level [44]. Clip-level action detectors output spatiotemporal tubelets with their corresponding action in a given video clip [44], whereas frame-level action detectors output human boxes and their actions only for a given keyframe in a video clip [44], relying on postprocessing methods to build spatiotemporal tubelets. Clip-level action detection is usually more computationally expensive compared to frame-level action detection. We focus on frame-level task in this work.

Adapting Recognition/Classification Transformer Backbones for Detection involve either: (1) directly regressing [PATCH] tokens [28, 29, 33], (2) adding an extra sequence of trainable [DET] detection tokens to the input and regressing those tokens at the output [9], (3) using them as a backbone

to produce feature maps, following a two-stage/FasterRCNN detection scheme [11, 41, 42], and (4) using them as the encoder in an encoder-decoder architecture similar to DETR or AdaMixer [12, 46, 51]. As of current, (2) has not yet been explored for video transformers for action detection, nor has it been explored for multimodality; we explore this scheme for videos as well as its utility in the open-vocabulary setting.

Open-Vocabulary Learning has dominated tasks in the image domain, such as image classification [35], object detection [28, 29], and image segmentation [22]. In the video domain, most open-vocabulary works revolve around video classification/retrieval [1, 18, 27, 36, 43, 45, 47] and specifically for human actions, temporal action detection [21, 31, 32, 37] which only localizes the start and end times of an action, not the spatial location of people and their individual actions. In the more spatially fine-grained task of action detection in videos, existing works are predominantly closed-set.

Open-vocabulary action detection remains an ill-posed task primarily due to the lack of large-scale action detection datasets. Commonly used action detection datasets, such as UCF-101-24 [39], JHMDB [19], AVA [13], and AVA-Kinetics [25] have between 21-80 action classes. Prior methods to deal with the lack of large scale data revolve around adapting existing vision-language contrastive models which are already pretrained on large image-text/video-text datasets for action detection and using an ill-posed base-to-novel scheme to split the datasets into base actions for training and novel actions for evaluation.

iCLIP [17] is one of the first work to extend action detection into the vision-language domain by adapting frozen CLIP [35] image and text encoders within an Asynchronous Interaction Aggregation (AIA) network [40]. The model employs a complicated pipeline: frames from an input video clip are processed by a pretrained closed-set object detector, then object proposals are cropped and encoded by the frozen CLIP image encoder before passing through the AIA network. This multi-stage design introduces inefficiencies and a key bottleneck: the closed-set object detector. Since this detector is limited to objects it was trained on, it may fail to capture novel or unlabeled objects, leading to potential information loss. OpenMixer [3] extends the closed-set encoder-decoder action detection model STMixer [46] to the open-vocabulary setting by adding a frozen video backbone and text encoder from a video-language contrastive model.

These existing approaches are pretrained on small-scale action detection datasets, limiting their exposure to a maximum of only 18 action classes during training, which risks overfitting the model to these actions.

In contrast, (1) our training method allows more than 700 action classes to be seen during training, leading to (2) not having to rely on freezing the pretrained vision and language encoders to preserve learned semantics, and (3) our model is

encoder-only; we do not rely on adding an external human detector or a parameter-heavy decoder.

Scaling Open-Vocabulary Detection with Weak Supervision in the task of object detection in images involve extending detection capabilities to more classes without exhaustive manual labeling. DETIC [53] introduces a method to use image classification datasets that have no annotated bounding boxes by implementing several heuristics to generate a pseudobox for each image. OWLv2 [29], 3Ways [2] and RegionCLIP [52] rely on self-training by collecting pseudoboxes from their own detections on large-scale image-text datasets and further training their model on these boxes.

In contrast, our method does not increase the number of videos, nor do we use pseudoboxes. To our knowledge, there is no existing work focused on weakly-supervised scaling for open-vocabulary action detection.

3. Methodology

Problem Formulation Given a video $V = (v_1, v_2, \dots, v_L)$ with L frames, the task of frame-level action detection is to train a model to output a set of human bounding boxes on the keyframe v_K and classify the actions associated with them. We adopt the open-vocabulary definition from the well-studied object detection problem in images [28], where the model is trained on a set of base action classes and is expected to generalize to both base and novel action classes during testing. In the following sections, we introduce our model and our weakly supervised training scheme centered on AVA-Kinetics to expand the number of action classes. An overview of our approach is shown in Figure 2.

3.1. SiA Architecture

Our model consists of a video encoder and text encoder which has been initially contrastively pretrained on a large scale video-text dataset for open-vocabulary video classification/retrieval.

Video Encoder: In the interest of avoiding overfitting newly added parameters to base action classes, we seek to avoid designing and attaching additional parameter-heavy modules to the pretrained video encoder. We attempt two schemes on the video encoder for action detection as following:

1. Temporally average pool [PATCH] tokens at the output and regress them for action detection, which has already been explored by BMViT [33].
2. Remove the [CLS] token and add 100 [DET] tokens to the input sequence and regress them at the output for action detection. To our knowledge, this scheme has not yet been attempted for video action detection.

In our ablations, we find that the [DET] token scheme is a better alternative, and choose it for our final model.

Following transformer-based detection architectures [4, 9, 12, 33, 38, 46, 51], we use two MLPs to regress the output

Table 1. **Number of action classes** in existing closed-set action detection datasets versus after both of our weakly-supervised training recipes: Naive Weak Supervision (NWS) and Assignment-based Weak Supervision (AWS). (UCF-MAMA-H: only human actions)

Dataset	# Action Classes	Multi-label Actions
JHMDB	21	×
UCF-101-24	24	×
UCF-MAMA	35	×
UCF-MAMA-H	27	×
MultiSports	66	×
AVA	80	✓
AVA-Kinetics	80	✓
+NWS (ours)	700+	✓
+AWS (ours)	700+	✓

tokens to obtain bounding boxes and actor scores, and a projection layer projects these tokens into vision embeddings; for each output token, the modified video encoder outputs a triplet of bounding box coordinates, actor probability and a vision embedding, $(\mathbf{b}, \mathbf{p}_{act}, \mathbf{e}_v)$. Tokens with $\mathbf{p}_{act}$ more than 0.5 are considered to have an actor in them, whereas tokens with $\mathbf{p}_{act}$ less than 0.5 are considered background tokens and filtered out.

Text Encoder: We utilize LoRA [16] for the MLPs in each transformer block of the text encoder, keeping the original weights frozen and finetuning the LoRA modules to better align the output text embeddings $embed_t$ for region-specific actions. We show the importance of LoRA-finetuning the text encoder in our ablations.

Detecting Actions: For any given input clip, we designate the middle frame as the keyframe and specifically train our model to detect humans and their actions within that frame. For each output token with a positive human detection at the output of the video encoder, we determine the actions of the detected individual by calculating the cosine similarity S between $\mathbf{e}_v$, and the encoded text embedding, $\mathbf{e}_t$, of the target action, where $S = \frac{\mathbf{e}_v \cdot \mathbf{e}_t}{\|\mathbf{e}_v\| \|\mathbf{e}_t\|}$. S closer to 1 indicates a higher likelihood that the detected person is performing the specified target action, and vice versa for values close to -1.

The text encoder is able to encode any action as a textual input, expanding detection capabilities to actions unseen during training, provided that a sufficient number of human actions are used during training. In contrast to previous models (iCLIP and OpenMixer), our model is single-stage and end-to-end trainable, avoiding the need for additional parameter-heavy modules attached to the vision encoder or external detectors to generate human proposals.

3.2. Scaling Action Classes Using Weak Supervision

To address the limited number of action classes in existing action detection datasets as shown in Table 1, we aim to significantly increase the number of action classes by exploiting and unlocking the full potential of AVA-Kinetics [25] with




Figure 2. *Overview of SiA architecture (top) and our weak-supervision scheme (bottom)*: Our architecture consists of a contrastive-pretrained video and text encoder. We add [DET] tokens to the input sequence and regress them at the output to convert the model for action detection. We train it to specifically detect humans and their actions only in the keyframe, which we set to the middle frame in any given input video clip. Our weak-supervision scheme naively appends the global action of Kinetics-700 videos to their AVA annotations; refined annotations are obtained by using our model trained on these naive annotations to assign the global action to the right boxes. In the shown example, the highlighted boxes denotes that the global action has been assigned to the human box; the Kinetics-700 action ‘cracking back’ is properly assigned to the only person cracking the back of another person.

two weakly-supervised approaches outlined below:

Naive Weak Supervision (NWS): AVA-Kinetics is a dataset that combines the original AVA dataset with a subset of Kinetics-700 [5] videos, annotated with the 80 action classes of AVA. Since the human boxes in these Kinetics-700 videos are already annotated following the AVA format with AVA classes, we introduce a weakly-supervised approach to expand the number of action classes from 80 to over 700. This is achieved by appending the global action label of each Kinetics-700 video to the multi-action ground-truth labels for the human boxes within that video, treating these appended labels as pseudo-labels. We refer to this method as NWS.

Assignment-based Weak Supervision (AWS): The NWS scheme outlined above presents two main issues: (1) not all actors in a Kinetics-700 video may be performing the global action assigned to that video (as shown in Figure 2), and (2) the global action may not occur in the frames surrounding the AVA-annotated keyframe. To address these limitations, we introduce an enhanced approach based on self-training, named AWS. AWS training proceeds in two stages:

1. **Initial Training with NWS:** In the first step, we train the

model using the NWS strategy.

2. **Assigning Pseudolabels to Ground Truth Boxes:** The NWS model is used to assign the global Kinetics-700 action class to the most relevant ground-truth human boxes in each Kinetics-700 video within AVA-Kinetics. Subsequently, the model is trained on these refined pseudolabels.

Unlike traditional self-training approaches in object detection, we do not use the output boxes of the model as pseudoboxes for self-training. Instead, we use Hungarian matching on the output to allocate the global Kinetics-700 action class to the ideal ground-truth human box in each instance. This process ensures more accurate pseudolabels as shown in Figure 2 and improves the overall performance of action detection, as well as eliminating the additional uncertainty incurred by using pseudoboxes.

3.3. Training Objective

Our model is trained using a bipartite matching loss following transformer-based detection models [4, 9, 12, 33, 38, 46, 51]. For all predicted triplets $(\mathbf{b}, \mathbf{p}_{act}, \mathbf{e}_v)$, only actor scores $\mathbf{p}_{act}$ and bounding box coordinates $\mathbf{b}$ are used for the

initial Hungarian matching step to match predictions with ground truth labels, as the only object of interest is the human figure. Finally, our bipartite loss function is as follows: $\mathcal{L}_{loss} = \lambda_{actor}CE_{actor} + \lambda_{box}\mathcal{L}_{box} + \lambda_{action}CE_{action}$ where CE_{actor} , $\mathcal{L}_{box}$ and CE_{action} represent the actor classification loss, bounding box loss, and action classification loss, respectively. Following OWL-ViT [28, 29], each λ is set to 2. More details on training can be found in the supplementary.

4. Experimental Setup

Implementation Details: We initialize the video and text encoder from ViCLIP-B16 pretrained on InternVid-10m-FLT [45]. More details on training configurations, video sampling strategy and hyperparameters can be found in the supplementary.

GPT4-assisted Text Augmentation: Following recent works [3, 20] in text augmentation for open-vocabulary detection, we use GPT4 to generate 16 descriptors for each action class to alleviate generalization issues. During training, we randomly sample one descriptor for each action class that appears in a training batch. During evaluation, for a given action, we average the cosine similarity for all 16 descriptors to obtain one final cosine similarity for that action.

Evaluation Metric: We quantify our results using frame-level mean average precision (f-mAP) with Intersection-over-Union (IoU) threshold at 0.5 following previous works in action detection [6, 17, 33].

4.1. Datasets

We use six closed-set action detection datasets for our experiments: AVA [13], AVA-Kinetics (AVA-K) [25], UCF-101-24 [39], JHMDB51-21 [19], Multisports [26], and UCF-MAMA [30]. AVA contains 299 videos each lasting 15 minutes with keyframe annotations at every second. There are 80 atomic action classes in AVA, and the annotations are multi-label in nature. AVA-K contains 238,476 Kinetics-700 videos in addition to the original 299 AVA videos, annotated with 80 AVA classes in a similar manner. Similar to AVA, the annotations are multi-label in nature; currently, AVA and AVA-Kinetics are the only action detection datasets with multi-label human boxes. For both datasets, we only use AVA2.2 annotations for AVA videos.

UCF-101-24 contains 2284 temporally untrimmed videos for training and 923 for testing distributed amongst 24 action classes. Annotations are in the form of spatiotemporal tubes.

JHMDB51-21 consists of 21 action classes split across 600 temporally trimmed videos for training and 300 for testing, with annotations in the form of spatiotemporal tubes.

MultiSports consists of 4 sports and each sport consists of a set of fine-grained sport-specific actions, totaling to 66 action classes. It consists of 1574 untrimmed videos for

training and 555 for validation. Similar to UCF-101-24 and JHMDB51-21, annotations in the form of spatiotemporal tubes.

UCF-MAMA consists of high-resolution, temporally cropped videos from VIRAT [34] and MEVA [7] in a surveillance-style footage that depict humans and vehicles at a long range, totalling to 35 action classes. The annotations also include non-human actions, which we remove during training. Specifically, we remove vehicle actions (e.g. ‘Vehicle Turning Left’, ‘Vehicle Turning Right’), reducing the number of action classes to 27.

4.2. Training and Evaluation

We evaluate our model using two settings: 1) *Base-to-Novel*: A single dataset is partitioned into base and novel categories, following the approach established in previous open-vocabulary object detection works, such as [8, 49]. 2) *Cross-dataset*: Training is performed on one dataset, while evaluation is conducted on separate downstream datasets. This setup aligns with the cross-dataset approach used in open-vocabulary object detection [8]. We discuss the two setups for different datasets in detail below.

Base-to-Novel: Following iCLIP [17] and OpenMixer [3], we use UCF-101-24 [39] or JHMDB [19] and randomly split their videos into base classes for training and novel classes for zero-shot inference. The base-to-novel ratio is either 75%-25% or 50%-50%. We use our weights pretrained on AVA-Kinetics + AWS before training our model in this setup. **The issue with base-to-novel:** JHMDB and UCF-101-24 are small action detection datasets with only 21 and 24 actions respectively. Splitting a set of novel actions from these datasets will result in fewer action classes for training and even fewer for evaluation, rendering it ill-posed. Instead of relying on this benchmark to evaluate open-vocabulary capabilities, we devise two cross-dataset schemes to evaluate on all action classes of any given downstream dataset without ever using them for training, similar to how all actions are evaluated in a closed-set setting.

Cross-Dataset: 1) **AVA-Kinetics:** Our primary contribution lies in this setting. Our model is trained on the 80 action classes from AVA, and we further employ NWS/AWS methods to increase the number of base action classes for training from 80 to over 700. For downstream evaluation, we evaluate on all action classes from UCF-101-24, JHMDB, MultiSports, and UCF-MAMA. 2) **UCF:JHMDB:** We use the UCF-101-24 [39] classes as base classes for training and treat JHMDB [19] classes as novel classes for zero-shot inference. We use this setup in our ablations.

5. Results and Analysis

Baseline: In the absence of prior open-vocabulary results for our main benchmark in Table 2, we devise a 3-stage baseline using off-the-shelf components. Our baseline is as follows:




Figure 3. *Open-Vocabulary Qualitative Results* on UCF-101-24 (1st row) and JHMDB (2nd row) from the output of our model, trained on AVA-Kinetics with AWS; this model is not trained on UCF-101-24 or JHMDB. Green boxes/labels denote predictions and red boxes/labels denote the ground truth. The multi-label nature of our model is also able to determine actions that are not labeled in the ground truth of JHMDB. Additionally, class confusion occurs for specific UCF-101-24 actions, notably ‘Basketball’ and ‘Basketball Dunk’.

Table 2. *Open-vocabulary results on training with AVA-Kinetics* without Kinetics-700 labels, with NWS and with AWS. * denotes that we only use human boxes/actions; non-human actors such as vehicles and their actions are removed.

Baseline	UCF f@0.5	JHMDB f@0.5	MultiSports f@0.5	UCF-MAMA* f@0.5
LanguageBind [54]	25.2	36.2	0.0	0.0
X-CLIP [27]	19.6	33.2	0.4	0.1
ViCLIP [45]	25.5	39.9	0.1	0.2
Ours	UCF f@0.5	JHMDB f@0.5	MultiSports f@0.5	UCF-MAMA* f@0.5
SiA-B16	33.7	39.9	1.3	0.5
+ NWS	36.7	51.5	0.2	0.6
+ AWS	42.6	57.1	0.8	0.6

Table 3. *JHMDB base-to-novel results* for both 75-25 and 50-50 splits (%).

Model	75-25 Split		50-50 Split	
	Base@0.5	Novel@0.5	Base@0.5	Novel@0.5
iCLIP [17]	-	66.8	-	45.2
OpenMixer [3]	-	77.1	-	-
SiA-B16	81.4	83.2	87.5	61.0

Table 4. *UCF-101-24 base-to-novel results* for both 75-25 and 50-50 splits (%).

Model	75-25 Split		50-50 Split	
	Base@0.5	Novel@0.5	Base@0.5	Novel@0.5
iCLIP [17]	-	72.5	-	60.3
SiA-B16	97.0	97.1	94.7	75.1

1) For a given input clip, human detections are obtained for all frames using a pretrained human detector. 2) Human tubelets are built across the clip using ByteTrack [50]. 3) The action associated with each human tubelet is obtained by cropping the tubelet from the input clip and classified by an off-the-shelf video-language model [27, 45, 54].

Table 5. *Closed-set results* after full-finetuning on the downstream datasets. * denotes that we only use human actions.

Model	UCF f@0.5	JHMDB f@0.5	MultiSports f@0.5	UCF-MAMA f@0.5
YOWO [23]	75.7	80.4	-	-
TubeR [51]	81.3	-	-	-
STMixer [46]	83.7	86.7	-	-
YOWOv2 [48]	87.0	-	-	-
HIT [10]	84.8	83.8	33.3	-
EVAD [6]	85.1	90.2	-	-
BMViT [33]	90.7	88.4	-	-
STAR [12]	90.3	92.1	59.3	-
VCN-MA [30]	-	-	-	0.4
SiA-B16	88.5	88.5	28.8	4.0*

5.1. Open-Vocabulary Evaluation

As shown in our new open-vocabulary benchmark in Table 2, our model consistently exceeds the performance of the training-free baselines in all four downstream datasets, both with and without NWS and AWS.

For the base-to-novel benchmarks on JHMDB and UCF-101-24 in Table 3 and 4 respectively, our method outperforms iCLIP on both 75:25 and 50:50 splits, as well as OpenMixer on the 75:25 split for JHMDB.

5.2. Closed-Set Evaluation

We show that our model also performs sufficiently in a closed-set setting by comparing against the latest closed-set action detection models. We initialize our model from AWS-pretrained weights and finetune them on each downstream closed-set action detection dataset. No text augmentation is applied and all available actions are passed into the text encoder to emulate closed-set training. Our results are shown in Table 5. Closed-set performance is higher than open-vocabulary, consistent with findings in image object detection [14, 24].




Figure 4. *t-SNE plot of text embeddings* for our model trained on AVA-Kinetics with (a) frozen text encoder, (b) LoRA-finetuned text encoder, and (c) LoRA-finetuned text encoder with AWS to include the 700 actions from Kinetics-700. More downstream classes lie within the cluster of base (AVA) classes after LoRA-finetuning the text encoder, and even more lie within the AVA and Kinetics-700 classes after introducing AWS.




Figure 5. *Visualizations of AWS pseudolabel assignment* showing multi-human instances where the global action has been allocated to the correct person.

5.3. Analysis of NWS and AWS

Quantitative comparison: NWS results in a notable performance boost on UCF-101-24, JHMDB, and UCF-MAMA, as opposed to only training on 80 AVA actions.

While AWS cannot correct every single noisy pseudolabel from NWS, the performance increase compared to NWS suggests that most of the erroneous pseudolabels from NWS have been rectified, as shown in Table 2.

Qualitative analysis: We must highlight that AVA-Kinetics annotations for Kinetics-700 videos have 1.2 human boxes on average; most of the erroneous pseudolabels are caused by annotations with more than 1 person, which is only a small proportion of the dataset (28%). Nevertheless, for certain videos, AWS is able to correct NWS pseudolabels as shown in Figure 5.

Visualization of actions in embedding space: From the t-SNE plots in Figure 4, we can observe that introducing NWS and AWS to include Kinetics-700 actions alongside AVA actions further expands the convex hull of base action embeddings, and majority of the downstream actions from the aforementioned datasets lie within this hull.

5.4. Ablations

Encoder-only design: [PATCH] vs [DET] token regression: As shown in Figure 6, regressing [DET] tokens yields

better downstream results than using [PATCH] tokens.

Furthermore, within the first 100 training iterations, the model with the [DET] token design demonstrates faster convergence compared to the model using only [PATCH] tokens, as shown in Figure 7.

In summary, the [DET] token scheme is a more effective approach for converting the video encoder of ViCLIP for action detection, and we finalize our design on this scheme.

Impact of the number of [DET] tokens: From Figure 7, we observe that increasing the number of [DET] tokens is detrimental to downstream performance. Nevertheless, to accommodate real-world use-cases where many people can be present in a surveillance-style footage such as UCF-MAMA, we choose to use 100 [DET] tokens as the default.

Importance of finetuning the text encoder: From Table 6 we find that introducing LoRA to the frozen text encoder yields a significant increase in performance, as opposed to keeping it frozen. This highlights the need to adapt the embeddings from the video-language pretrained text-encoder to be region specific.

From the t-SNE plots in Figure 4, we observe that LoRA-finetuning the text encoder expands the convex hull of the base (AVA) action class embeddings, which starts to include more downstream actions from UCF-101-24, JHMDB, MultiSports and UCF-MAMA within the cluster of base action




Figure 6. **Ablation:** [PATCH] vs [DET] token regression on training with UCF-101-24 and evaluation on JHMDB.




Figure 7. **Ablation:** f-mAP@0.5 for JHMDB during the first 100 iterations of training on UCF-101-24.

Table 6. **Ablation:** Impact of LoRA-finetuning the text encoder.

Dataset	Text	Base	UCF-101-24	JHMDB
		f@0.5	f@0.5	f@0.5
AVA	Frozen	10.2	5.1	6.1
AVA	LoRA	20.1	19.0	34.4
AVA-K	Frozen	12.3	6.0	14.3
AVA-K	LoRA	23.2	33.7	39.9

embeddings.

GPT4-assisted text augmentation: In Table 7, we observe that the use of GPT4-generated descriptors for action classes yields a significant increase in detection performance compared to simply using class names.

6. Discussion

Weak-supervision: scaling the number of videos vs number of actions: Unlike existing weakly supervised scaling methods in object detection for images which prioritize increasing the number of images, our approach focuses on expanding the number of action classes, as humans are the

Table 7. **Ablation:** Impact of GPT4-assisted text augmentation (with a frozen text encoder).

Dataset	GPT4	Base	UCF-101-24	JHMDB
		f@0.5	f@0.5	f@0.5
AVA	×	7.4	0.6	2.7
AVA	✓	10.2	5.1	6.1
AVA-K	×	10.8	4.8	6.1
AVA-K	✓	12.3	6.0	14.3

sole object of interest, which we achieve by adding more action labels to the already-annotated human boxes in an existing action detection dataset, providing our model with a more comprehensive understanding of potential actions without the need to increase the number of videos.

Domain-specific actions: Fine-grained actions from the MultiSports dataset require sport-specific domain knowledge (e.g., aerobic kick jump vs. aerobic straddle jump). Closed-set performance of our model on this dataset is significantly higher than open-vocabulary, as shown in Tables 2 and 5. We conclude that domain-specific actions are better handled with a closed-set training approach.

The issue with single-action datasets: Our model effectively detects multiple simultaneous actions, even in datasets annotated with only a single action label per instance (e.g., JHMDB), as shown in Figure 3. This highlights a fundamental flaw in such datasets: they impose an artificial constraint by labeling each video with only one action, despite real-world scenarios where multiple actions co-occur. This forces models to ignore secondary actions during evaluation. Additionally, single-action datasets that contain ambiguous label hierarchies (e.g., UCF-101-24 includes both "basketball" and "basketball dunk,") unfairly penalizes models that recognize broader activities but fail to predict the most specific label. Such inconsistencies distort performance metrics and hinder the development of truly generalizable models.

7. Conclusion

In this work, we addressed the challenges of open-vocabulary action detection by introducing SiA, a single-stage, encoder-only model trained end-to-end for action detection and a weakly supervised training strategy that enables SiA to see more than 700 action classes during training, a significant departure from prior work that involve complicated adaptations to pretrained vision-language models for detection that sees less than 18 actions during training. Finally, we introduce a new cross-dataset benchmark to evaluate open-vocabulary action detection to replace the ill-posed base-to-novel benchmark on small action detection datasets used by prior works in zero-shot/open-vocabulary action detection, showing novel results to serve as baselines for future work.

8. Acknowledgment

This material is based upon work supported by the National Science Foundation (NSF) Accelerating Research Translation (ART) program under Grant No. 2331319. In addition, this research has benefited from the Microsoft Accelerating Foundation Models Research (AFMR) grant program.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF or Microsoft.

References

- [1] Shahzad Ahmad, Sukalpa Chanda, and Yogesh S Rawat. Ez-clip: Efficient zeroshot video action recognition. *arXiv preprint arXiv:2312.08010*, 2023. 2
- [2] Relja Arandjelović, Alex Andonian, Arthur Mensch, Olivier J Hénaff, Jean-Baptiste Alayrac, and Andrew Zisserman. Three ways to improve feature alignment for open vocabulary detection. *arXiv preprint arXiv:2303.13518*, 2023. 3
- [3] Wentao Bao, Kai Li, Yuxiao Chen, Deep Patel, Martin Renqiang Min, and Yu Kong. Exploiting vlm localizability and semantics for open vocabulary action detection. *arXiv preprint arXiv:2411.10922*, 2024. 1, 2, 5, 6
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 3, 4
- [5] Joao Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *arXiv preprint arXiv:1907.06987*, 2019. 4
- [6] Lei Chen, Zhan Tong, Yibing Song, Gangshan Wu, and Limin Wang. Efficient video action detection with token dropout and context refinement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10388–10399, 2023. 1, 5, 6
- [7] Kellie Corona, Katie Osterdahl, Roderic Collins, and Anthony Hoogs. Meva: A large-scale multiview, multimodal video dataset for activity detection. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1060–1068, 2021. 5
- [8] Ruohuan Fang, Guansong Pang, and Xiao Bai. Simple image-level classification improves open-vocabulary object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1716–1725, 2024. 5
- [9] Yuxin Fang, Bencheng Liao, Xinggang Wang, Jiemin Fang, Jiyang Qi, Rui Wu, Jianwei Niu, and Wenyu Liu. You only look at one sequence: Rethinking transformer in vision through object detection. *arXiv preprint arXiv:2106.00666*, 2021. 2, 3, 4
- [10] Gueter Josmy Faure, Min-Hung Chen, and Shang-Hong Lai. Holistic interaction transformer network for action detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3340–3350, 2023. 1, 6
- [11] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6202–6211, 2019. 2
- [12] Alexey A Gritsenko, Xuehan Xiong, Josip Djolonga, Mostafa Dehghani, Chen Sun, Mario Lucic, Cordelia Schmid, and Anurag Arnab. End-to-end spatio-temporal action localisation with video transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18373–18383, 2024. 1, 2, 3, 4, 6
- [13] Chunhui Gu, Chen Sun, David A Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, et al. Ava: A video dataset of spatio-temporally localized atomic visual actions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6047–6056, 2018. 2, 5
- [14] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921*, 2021. 6
- [15] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5356–5364, 2019. 1
- [16] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. 3
- [17] Wei-Jhe Huang, Jheng-Hsien Yeh, Min-Hung Chen, Gueter Josmy Faure, and Shang-Hong Lai. Interaction-aware prompting for zero-shot spatio-temporal action detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 284–293, 2023. 1, 2, 5, 6
- [18] Xiaohu Huang, Hao Zhou, Kun Yao, and Kai Han. Froster: Frozen clip is a strong teacher for open-vocabulary action recognition. *arXiv preprint arXiv:2402.03241*, 2024. 2
- [19] Hueihan Jhuang, Juergen Gall, Silvia Zuffi, Cordelia Schmid, and Michael J. Black. Towards understanding action recognition. In *2013 IEEE International Conference on Computer Vision*, pages 3192–3199, 2013. 2, 5
- [20] Sheng Jin, Xueying Jiang, Jiaxing Huang, Lewei Lu, and Shijian Lu. Llm meets vlm: Boost open vocabulary object detection with fine-grained descriptors. *arXiv preprint arXiv:2402.04630*, 2024. 5
- [21] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. In *European Conference on Computer Vision*, pages 105–124. Springer, 2022. 2
- [22] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 2
- [23] Okan Köpüklü, Xiangyu Wei, and Gerhard Rigoll. You only watch once: A unified cnn architecture for real-time spatiotemporal action localization. *arXiv preprint arXiv:1911.06644*, 2019. 6
- [24] Weicheng Kuo, Yin Cui, Xiuye Gu, AJ Piergiovanni, and Anelia Angelova. F-vm: Open-vocabulary object detection

- upon frozen vision and language models. *arXiv preprint arXiv:2209.15639*, 2022. 6
- [25] Ang Li, Meghana Thotakuri, David A Ross, João Carreira, Alexander Vostroikov, and Andrew Zisserman. The avakinetix localized human actions video dataset. *arXiv preprint arXiv:2005.00214*, 2020. 2, 3, 5
- [26] Yixuan Li, Lei Chen, Runyu He, Zhenzhi Wang, Gangshan Wu, and Limin Wang. Multisports: A multi-person video dataset of spatio-temporally localized sports actions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 13536–13545, 2021. 2, 5
- [27] Yiwei Ma, Guohai Xu, Xiaoshuai Sun, Ming Yan, Ji Zhang, and Rongrong Ji. X-clip: End-to-end multi-grained contrastive learning for video-text retrieval. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 638–647, 2022. 2, 6
- [28] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. Simple open-vocabulary object detection. In *European Conference on Computer Vision*, pages 728–755. Springer, 2022. 2, 3, 5
- [29] Matthias Minderer, Alexey Gritsenko, and Neil Houlsby. Scaling open-vocabulary object detection. *Advances in Neural Information Processing Systems*, 36, 2024. 2, 3, 5
- [30] Rajat Modi, Aayush Jung Rana, Akash Kumar, Praveen Tirupattur, Shruti Vyas, Yogesh Singh Rawat, and Mubarak Shah. Video action detection: Analysing limitations and challenges. *arXiv preprint arXiv:2204.07892*, 2022. 2, 5, 6
- [31] Sauradip Nag, Xiatian Zhu, Yi-Zhe Song, and Tao Xiang. Zero-shot temporal action detection via vision-language prompting. In *European Conference on Computer Vision*, pages 681–697. Springer, 2022. 2
- [32] Trung Thanh Nguyen, Yasutomo Kawanishi, Takahiro Komamizu, and Ichiro Ide. One-stage open-vocabulary temporal action detection leveraging temporal multi-scale and action label features. *arXiv preprint arXiv:2404.19542*, 2024. 2
- [33] Ioanna Ntinou, Enrique Sanchez, and Georgios Tzimiropoulos. Multiscale vision transformers meet bipartite matching for efficient single-stage action localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18827–18836, 2024. 1, 2, 3, 4, 5, 6
- [34] Sangmin Oh, Anthony Hoogs, Amitha Perera, Naresh Cuntoor, Chia-Chih Chen, Jong Taek Lee, Saurajit Mukherjee, JK Aggarwal, Hyungtae Lee, Larry Davis, et al. A large-scale benchmark dataset for event recognition in surveillance video. In *CVPR 2011*, pages 3153–3160. IEEE, 2011. 5
- [35] 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 *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2
- [36] Hanoona Rasheed, Muhammad Uzair Khattak, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Fine-tuned clip models are efficient video learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6545–6554, 2023. 2
- [37] Vivek Rathod, Bryan Seybold, Sudheendra Vijayanarasimhan, Austin Myers, Xiuye Gu, Vighnesh Birodkar, and David A Ross. Open-vocabulary temporal action detection with off-the-shelf image-text features. *arXiv preprint arXiv:2212.10596*, 2022. 2
- [38] Hwanjun Song, Deqing Sun, Sanghyuk Chun, Varun Jampani, Dongyoon Han, Byeongho Heo, Wonjae Kim, and Ming-Hsuan Yang. VidT: An efficient and effective fully transformer-based object detector. *arXiv preprint arXiv:2110.03921*, 2021. 3, 4
- [39] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 2, 5
- [40] Jiajun Tang, Jin Xia, Xinzhi Mu, Bo Pang, and Cewu Lu. Asynchronous interaction aggregation for action detection. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XV 16*, pages 71–87. Springer, 2020. 2
- [41] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in neural information processing systems*, 35:10078–10093, 2022. 2
- [42] Limin Wang, Bingkun Huang, Zhiyu Zhao, Zhan Tong, Yinan He, Yi Wang, Yali Wang, and Yu Qiao. Videomae v2: Scaling video masked autoencoders with dual masking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14549–14560, 2023. 2
- [43] Mengmeng Wang, Jiazheng Xing, Jianbiao Mei, Yong Liu, and Yunliang Jiang. Actionclip: Adapting language-image pretrained models for video action recognition. *IEEE Transactions on Neural Networks and Learning Systems*, 2023. 2
- [44] Peng Wang, Fanwei Zeng, and Yuntao Qian. A survey on deep learning-based spatio-temporal action detection. *arXiv preprint arXiv:2308.01618*, 2023. 2
- [45] Yi Wang, Yinan 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. *arXiv preprint arXiv:2307.06942*, 2023. 2, 5, 6
- [46] Tao Wu, Mengqi Cao, Ziteng Gao, Gangshan Wu, and Limin Wang. Stmixon: A one-stage sparse action detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14720–14729, 2023. 1, 2, 3, 4, 6
- [47] Hongwei Xue, Yuchong Sun, Bei Liu, Jianlong Fu, Ruihua Song, Houqiang Li, and Jiebo Luo. Clip-vip: Adapting pretrained image-text model to video-language representation alignment. *arXiv preprint arXiv:2209.06430*, 2022. 2
- [48] Jianhua Yang and Kun Dai. Yowov2: A stronger yet efficient multi-level detection framework for real-time spatio-temporal action detection. *arXiv preprint arXiv:2302.06848*, 2023. 6
- [49] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14388–14397, 2020. 5
- [50] Yifu Zhang, Peize Sun, Yi Jiang, Dongdong Yu, Fucheng Weng, Zehuan Yuan, Ping Luo, Wenyu Liu, and Xinggang

- Wang. Bytetrack: Multi-object tracking by associating every detection box. In *European conference on computer vision*, pages 1–21. Springer, 2022. [6](#)
- [51] Jiaojiao Zhao, Yanyi Zhang, Xinyu Li, Hao Chen, Shuai Bing, Mingze Xu, Chunhui Liu, Kaustav Kundu, Yuanjun Xiong, Davide Modolo, et al. Tuber: Tubelet transformer for video action detection. *arXiv preprint arXiv:2104.00969*, 2021. [1](#), [2](#), [3](#), [4](#), [6](#)
- [52] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luowei Zhou, Xiyang Dai, Lu Yuan, Yin Li, et al. Regionclip: Region-based language-image pretraining. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16793–16803, 2022. [3](#)
- [53] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Philipp Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. In *European Conference on Computer Vision*, pages 350–368. Springer, 2022. [3](#)
- [54] Bin Zhu, Bin Lin, Munan Ning, Yang Yan, Jiayi Cui, Hongfa Wang, Yatian Pang, Wenhao Jiang, Junwu Zhang, Zongwei Li, et al. Languagebind: Extending video-language pretraining to n-modality by language-based semantic alignment. *arXiv preprint arXiv:2310.01852*, 2023. [6](#)