

Video Action Detection: Analysing Limitations and Challenges

Rajat Modi[†], Aayush Jung Rana[†], Akash Kumar[†], Praveen Tirupattur[†],
Shruti Vyas[‡], Yogesh Singh Rawat[‡] & Mubarak Shah[‡]
Center for Research in Computer Vision
University of Central Florida, Orlando, Florida, USA

[†]{rajatmodi, aayushjr, akash.k, praveentirupattur}@knights.ucf.edu

[‡]{shruti, yogesh, shah}@crcv.ucf.edu

Abstract

Beyond possessing large enough size to feed data hungry machines (eg, transformers), what attributes measure the quality of a dataset? Assuming that the definitions of such attributes do exist, how do we quantify among their relative existences? Our work attempts to explore these questions for video action detection. The task aims to spatio-temporally localize an actor and assign a relevant action class. We first analyze the existing datasets on video action detection and discuss their limitations. Next, we propose a new dataset, Multi Actor Multi Action (MAMA) which overcomes these limitations and is more suitable for real world applications. In addition, we perform a biasness study which analyzes a key property differentiating videos from static images: the temporal aspect. This reveals if the actions in these datasets really need the motion information of an actor, or whether they predict the occurrence of an action even by looking at a single frame. Finally, we investigate the widely held assumptions on the importance of temporal ordering: is temporal ordering important for detecting these actions? Such extreme experiments show existence of biases which have managed to creep into existing methods inspite of careful modeling. The dataset and code is publicly available for research at ¹

1. Introduction

Video action understanding has been widely explored in recent years in terms of action classification [3, 11, 13, 30, 34–37, 40] and detection [8, 9, 15, 18, 21, 21, 22, 26, 38, 42] for a wide range of actions. Analyzing the videos for action classification task involves feature extraction and prediction on a video-level, with recent methods achieving state-of-the-art performance [3, 11, 34]. Compared to ac-

¹<https://www.crcv.ucf.edu/research/projects/mama-multi-actor-multi-action-dataset-for-action-detection/>




Figure 1. An illustration of relative properties of datasets in activity detection. Size of the bubble corresponds to the no of samples in the dataset. X axis: No of action classes. MAMA has more classes [35] than standard detection datasets. Y axis: No of Instances. MAMA has a lot more actors, including crowded scenarios. Finally, MAMA possesses 10x more samples.

tion classification, action detection is a harder task as it requires the spatio-temporal prediction along with the action class prediction for *each* action in the video. The increased complexity in action detection has garnered more interest recently with significant progress in various datasets [15, 18, 23, 26, 39, 42]. Spatio-temporal detection for action requires datasets with frame-level annotations (bounding-box, pixel-wise) which is costly to produce, resulting in limited datasets that can be used for this task. Current datasets annotated for video action detection include dense frames annotation (UCF101-24 [31], JHMDB-21 [17], UCF-Sports [28], VIRAT [25], MEVA [5]) and sparse frames annotation (A2D [41], DALY [39], AVA [12]), which is limited compared to the vast action classification datasets due to higher spatio-temporal annotation costs. A graphical comparison of these datasets has been illustrated in Figure 1.

Datasets with dense spatio-temporal annotations have

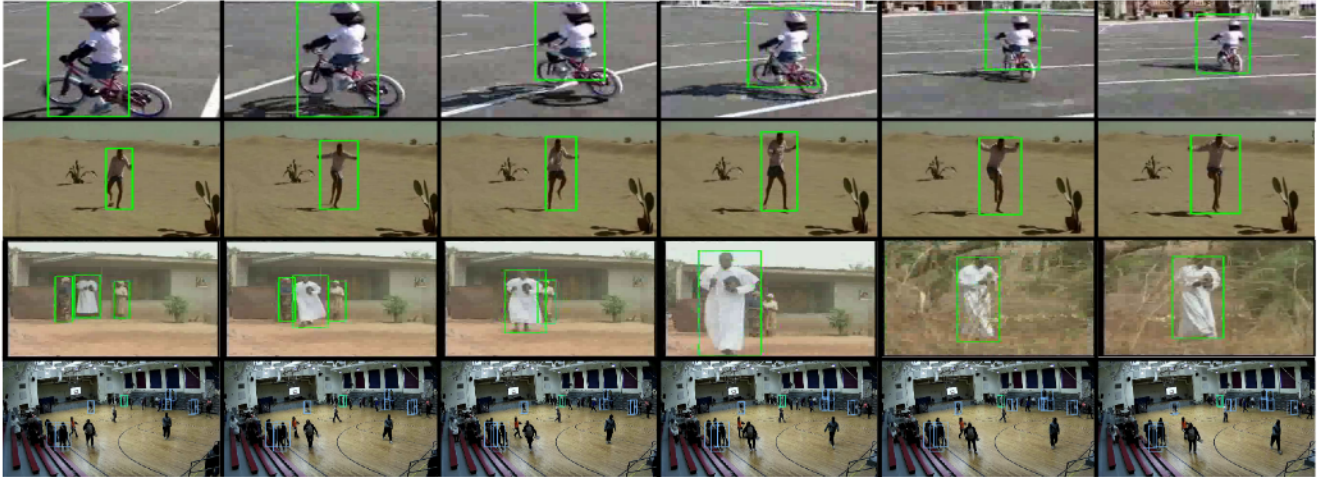


Figure 2. Each row shows samples from UCF101, JHMDB, AVA, and MAMA respectively. UCF101 (First Row): contains a single activity instances per video. Most samples show an actor in far closer perspective than observed in real scenarios. JHMDB (Second Row): has even lesser videos than UCF101. AVA (Third Row) Since most of the frames are annotated at 1 sec intervals, successive annotated frames have significant differences in the actor motion. MAMA (Fourth Row): Shows even far off actors being annotated, in a crowded setting. MAMA can possess multiple activities per video.

bounding-box or pixel-wise annotation for each action/actor for the entire duration of the action (UCF101-24, JHMDB-21, VIRAT, MEVA), with most datasets being untrimmed. While this gives better results on supervised methods, it is costly to annotate all the frames. UCF101-24 and JHMDB-21 are the two widely used dense annotation datasets for action detection and have saturated with recent methods performing exceedingly well [8, 42]. VIRAT and MEVA have larger untrimmed dense annotations from surveillance videos which is significantly more costly to produce compared to UCF101-24, but have low adoption in research community since they make action detection harder due to smaller actors, simultaneous multiple actions and untrimmed long videos [6, 27, 44]. Sparse annotation datasets (A2D, DALY, AVA) reduces annotation cost and allows annotating more samples, however they are harder to train for action detection task with current methods due to not having sufficient temporal annotation for doing action detection [10, 11, 26, 39, 42]. We analyze the properties of these various datasets and bring to light their similarities and differences in order to understand how action detection can be improved as a whole. Each dataset provides different context for the actions in them (eg. movie scenes, surveillance, controlled lab, YouTube videos, sports clips), with benefits and limitations for action detection generalization.

To better understand the importance of temporal annotation along with spatial annotation for better action detection, we analyze the effect temporal information has in different datasets and evaluate properties such as temporal aspect and order. It is observed that temporal information has positive correlation for action classification [3, 13]. We

look into what type of relations temporal information has for action detection task and provide insight into how each dataset contributes to this task. To further improve action detection in real-world scenarios, we introduce the **MAMA** dataset. We aim to provide more real world scenarios with multiple simultaneous actions, mitigating shortcomings of prior detection datasets. In summary, we contribute via (1) a new dataset addressing issues in existing datasets; (2) study of temporal aspect and order for action detection in all datasets; and (3) in depth analysis in similarities/differences of all video action datasets.

2. Video action detection datasets

Initial research on video action understanding focused on action classification where the task is to identify the action in a short, manually trimmed video containing a single action. Some of the popular action classification datasets are HMDB [20], UCF101 [31], Sports-1M [19], Moments in Time [24], TinyVIRAT [7, 33] and Kinetics [3]. Another action understanding problem that received interest is the temporal action localization, where the task is to detect the temporal extents of actions in a long untrimmed video. ActivityNet [2], THUMOS [16], MultiTHUMOS [43] and Charades [29] are some of the popular temporal action localization datasets. Compared to action classification and temporal action localization, action detection is a harder task and it involves finding both the spatial and temporal extents of actions in untrimmed videos. Some of the datasets that provide spatio-temporal annotations required to address this problem are UCF Sports [28], JHMDB [17], UCF101-

Datasets	VIRAT	MEVA	AVA	UCF Sports	UCF101-24	JHMDB	MAMA
#Classes	40	39	80	10	24	21	35
Source	Surveillance	Surveillance	Movies	Sports	Sports	Movies	Surveillance
Resolution	1920 × 1080	1920 × 1080	320 × 400	690 × 450	320 × 240	320 × 240	1920 × 1080
Total Videos	369	1145	430	150	3194	928	32726
Total Frames	1.05M	10.31M	~ 11M	10K	558K	32K	2.7M
Avg. Video Length	1.6 min	5 min	15min	5.8 sec	5.8 sec	1.4 sec	2.6 sec
Avg. Action Duration	10.1 sec	10.3 sec	-	5.8sec	4.5 sec	1.4 sec	2.6 sec
Total #Instances	8K	37K	1.62M	154	4030	928	32726
Multi-Actor	✓	✓	✓	✓	✓	✓	✓
Multi-Label	✓	✓	✓	✗	✗	✗	✓
Annotation Type	Boxes	Boxes	Boxes	Boxes	Boxes	Pixels	Boxes
Spatio-Temporal Ann.	✓	✓	✗	✓	✓	✓	✓
Class Distribution	Long-tail	Long-tail	Long-tail	Uniform	Uniform	Uniform	Long-tail

Table 1. A relative comparison of datasets for brevity. Three columns from the left are grey because although they are standard datasets in action-detection, they don’t meet the criteria for this biasness study. VIRAT/MEVA contain untrimmed videos. AVA only annotates frames at 1 sec intervals, therefore it is not truly spatio-temporal in nature. Notice how MAMA exceeds in several statistics to all the other datasets.

24 [31], AVA [12], VIRAT [25] and MEVA [5].

UCF Sports consists of 150 videos from 10 action classes and JHMDB contains a total of 928 videos with 21 action classes. UCF101-24 consists of 3207 videos with annotations for 24 action classes. These datasets provide annotations for each frame, but contain smaller number of actions, fewer number of video and with videos of shorter duration. The widely AVA dataset contains 430 videos, each video of length 15 minutes, and provides annotations for 80 atomic actions. In this dataset, the annotations are not provided for each frame but for a single frame at one second intervals. VIRAT provides annotations for 40 actions and contains 118 videos in the train/validation set and 246 videos in held-out test set. All the samples in this dataset are of long untrimmed videos of varying length. MEVA consists of 1056 videos, each 5 minutes long, with annotations for 37 actions. Videos in this dataset cover both indoor and outdoor scenes. In Figure 2 we show a sample frame from each of the action detection datasets to highlight the differences.

2.1. Limitations

Commonly considered sources for building action detection datasets are videos of sport activities, movie scenes or surveillance videos. Compared to videos of sport activities, action detection in videos from movies is much more challenging due to variations in background, view point, scale, and occlusion. However, most of the videos from the movies have a narrow field-of-view and the focus is on the actor/actors performing the action. Surveillance videos on the other hand, have a wider field-of-view and the actions in these videos can occur at different spatial locations. Apart from the source, these datasets also vary in the type of actions, number of actors and the granularity of annotations. The datasets UCF-Sports, UCF101-24 contain composite

actions (e.g., pole vaulting) with single actor and bounding box annotations for each frame; AVA is focused on atomic actions (e.g., stand) with multiple actors and provide bounding box annotations for all the actors in a single frame at every one second intervals; JHMDB contains both atomic and composite actions with single actor and pixel-level annotations for every frame. VIRAT and MEVA have both atomic and composite actions with multiple actors and provide annotations for every actor in each frame. Another aspect in which these datasets vary is the distribution of samples. While the small scale datasets (UCF Sports, UCF101-24 and JHMDB) contain similar number of samples for each class, the large-scale datasets (AVA, VIRAT, MEVA) have a long-tail distribution. Please refer to Table 1, for a detailed comparison of the datasets.

3. MAMA dataset

Most of the footage relevant for action detection comes from CCTV cameras typically mounted at higher altitudes. Naturally, the videos obtained are untrimmed which can only be processed in chunks by existing action detection models. However, the activities in such chunks are generally sparse, i.e. in most of the temporally sampled clips of smaller durations, no activity occurs. [eg, very few activities occurs in the night]. Therefore, we propose MAMA dataset which generates shorter temporal crops from untrimmed footage which are more realistic for modelling, and yet sufficiently challenging.

To build the intuition for our dataset, we first explain the temporal trimming protocol. Then, we explore the annotation properties of our dataset, along with the relative comparison of difficulty of MAMA dataset over existing action detection datasets. Finally, we discuss the evaluation proto-




Figure 3. A Bar plot showing the relative instance counts of each of the activities in the MAMA dataset. The bar heights (# instances) are shown in log scale.

col on our dataset which relies on standard detection based metrics.

3.1. Temporal Trimming Protocol

We model the problem of generating trimmed clips as a variant of the classic *scheduling problem*. Given original untrimmed videos from VIRAT/MEVA dataset, we first generate the [start,end] intervals for each of the activities. A min heap is used to store the activities in ascending order of finish times. For each activity interval, activities lasting lesser than 30 frames are dropped. If an activity lasts longer than 150 frames, we randomly generate a temporal crop between [30,150] frames. Finally, we utilize FFMPEG to actually generate clips from untrimmed video based on given temporal slices.

This simple, yet effective protocol guarantees that we make a temporal crop, whenever an activity is just starting. The lower constraint of 30 frame generates clips of ideal length so that enough temporal changes could be captured by a model. Varying lengths of clips in our dataset capture both the atomic actions like *Opening a Car Trunk* and long term actions like *Texting on a Cell Phone*. We believe that this should motivate the development of techniques which attend to different length of temporal slices simultaneously based on whether an action is atomic/long term.

To maintain the fairness of our dataset, i.e. the participants don't synthetically *rejoin* temporal crops of a same video to learn more long term context, we anonymize and shuffle the generated clips using a rotating hex cipher. Finally, we assign an 80/20 split of our generated clips using a weighed sampler for extreme multi labelled data. We ensure that the clips belonging to a same video footage go to only one of the splits.

3.2. Dataset Description

Fig. 3 shows the relative ratios of 35 classes in the MAMA dataset on a logarithmic scale. The most common activity in the MAMA dataset is *People Talking*, with the least common being *Closing Facility Door*. One interesting aspect is that most of the activities lasting for longer time durations are the ones which are more frequently observed. The activities which are instantaneous (eg *closing facility door*) are concentrated on the right end of the histogram. Surprisingly, the distributions of complementary activities like *Opening Facility Door* and *Closing Facility Door* are not identical, which might be a source of bias.

In Tab. 1, we show the statistics on the MAMA dataset. MAMA consists of a total of 32726 video with 25837 videos in the train split, and 6889 videos in the test split. The length of clips in the dataset ranges from 1 sec to 5 sec.

3.3. Difficulty And Diversity in MAMA Dataset

Indoor/Outdoor scenes: MAMA dataset contains the trimmed clips from both VIRAT/MEVA datasets. Due to this, MAMA contains both indoor/outdoor scenes. Certain indoor scenes like *Baseketball Court* are crowded, containing over 44 actors in a single frame. On the other hand, the outdoor scenes are generally relatively sparse. During the course of mapping the activities from VIRAT to MEVA during the dataset construction, we dropped certain labels like that of a *Person Walking*. Our belief is that an activity is characterized only when a person *interacts* with another object/person. Activities like walking just involves a single actor. So, our dataset will help models to learn to focus on *activities of interest* rather than just focusing on 'all' temporal movements in a video. We present some of the samples from our dataset in Figure 6.

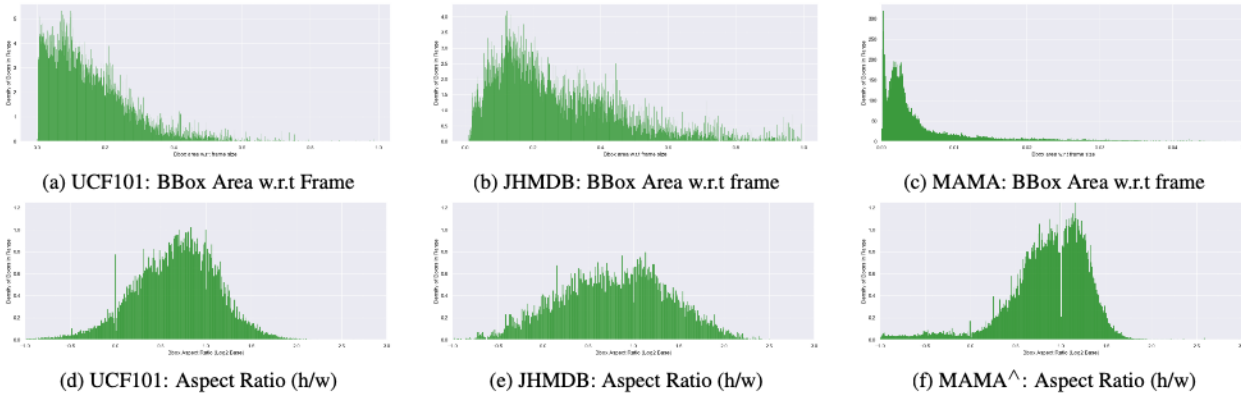


Figure 4. First Row: (X Axis) shows the possible areas of BBoxes w.r.t video frame in several action detection datasets. (Y Axis) shows the normalized density of bbox occurrences. Notice how MAMA’s density is concentrated on the left, signifying that most of the bboxes occupy lower resolution. Second Row: (X Axis): We plot the aspect ratios of the bboxes on a logarithmic scale. All the datasets follow a gaussian distribution, but MAMA has lower variance. Hence MAMA dataset models low resolution conditions with higher probability. [^]: X axis has been scaled to 0.05 due to lower bbox aspect ratios.

Bounding Box Areas and Aspect Ratios: Figure 4a, shows the relative distribution of the bbox areas with respect to the frame area on a logarithmic scale. Most of the samples in the MAMA dataset are constrained to $< 2\%$ of the frame area. Concurrently, Figure 4b, shows the density curve of the bounding box aspect ratios. While the bboxes generally model a nice gaussian distribution, we observe that the curve extends almost equally in both negative and positive ranges. This shows that MAMA dataset contains an almost equal no of wide and long objects.

3.4. Evaluation Protocol

We modify output channels of decoder in VideoCapsuleNet [8] to predict semantic segmentation volume for each of the 35 activities in the MAMA dataset. Then, we run 3D connected components on each of the activity’s volume to isolate the 3D tubes corresponding to an actor temporally. Components with temporal length less than 4 frames and lesser than 20 pixel prediction per frame are dropped. Finally, we fit bboxes on per-pixel actor localization to obtain frame wise detections.

Following the protocol as described in [18], we estimate the fMAP at 0.5 ioU threshold. Then, linking is done temporally to obtain the predicted tubes from the VideoCapsuleNet baseline. [8]. Finally, we estimate the spatio-temporal overlap between predicted and ground truth tubes using 3D ioU and report the vMAP. For classification accuracy, we report balanced accuracy over 35 classes of the MAMA dataset.

4. Analysing datasets

Actions from videos takes into account the position of involved actors in space and time and their relation to the surrounding during the action. From a detection perspective,

one key aspect is to understand how the actor changes with respect to their surrounding in temporal dimension. We look into how the temporal information can be influential for understanding different actions from different datasets.

4.1. Importance of temporal aspect

First we analyze to what degree does action understanding rely on the temporal aspect. Scene bias in videos is an issue which questions the need of temporal aspect in video understanding [4]. Certain actions can be guessed based on the scene alone, which leads to a heavy scene bias during learning. This can limit the use of such dataset for generalizing action detection to real world videos. We evaluate how important the temporal information is for actions across various datasets available for action detection.

4.2. Importance of temporal order

Along with scene bias, we also look into the importance of the temporal order of the information in order to understand the activity. While actions are based on how the actor moves in space and time, the temporal order of the motion can be essential to understand how important the scenes from each dataset are to detect the actions. Getting this relation between temporal order and action understanding can give insight into developing methods for better action detection.

5. Experiments and results

We will briefly describe the statistics for datasets used in our work to compare with the new dataset. Then, we discuss recent approaches for activity detection, their limitations, and our proposed modifications as well.




Figure 5. Overview of the effect of temporal aspect for different datasets. We show the accuracy, f-mAP and v-MAP (@0.5 IoU) for original, reversed and frozen samples across datasets.




Figure 6. Samples from the proposed MAMA dataset are illustrated. The video frames span both indoor/outdoor scenes and multiple actors. MAMA is realistic due to extremely fine temporal changes in actor trajectory. Best viewed on screen.

5.1. Datasets

UCF101-24: The number of training and testing videos in UCF101 is 2.2k and 900 respectively. The videos are untrimmed and the resolution of videos is 320x240.

JHMDB-21: Total number of videos are approximately 900 which is divided into 2:1 for training and testing. Video resolution is same as UCF101-24, however, the videos are trimmed.

AVA: dataset contains long untrimmed videos distributed amongst 80 actions. The resolution of videos are 320x400.

MAMA: The train test split is 25.8k and 6.9k respectively. The video resolution is 1920x1080.

For UCF101-24 and JHMDB-21, videos are resized to 224x224. For AVA, the frame size is 400x400. For MAMA dataset, the input resolution is 256x256.

5.2. Approaches

VCN: [8] is an end-to-end 3d encoder decoder based approach. The features are extracted at multiple checkpoints and its fed into the capsule routing algorithm. Video is classified from the output of encoder. To generate the localization map, we upsample the encoder features also known as activations. The ground truth vector is multiplied by the activations to propagate the class information. However,

this approach works only if the output is binary, meaning presence of a single action across the whole video. This approach is not suitable on AVA and MAMA dataset.

STEP: [42] It progressively refines proposal using regression. Starting with rough estimate of proposals, the network, updates the proposals with each iteration. It also targets to extend the bounding boxes in temporal dimension. The architecture has two branch: Global branch that works on spatio-temporal modeling of the whole input sequence and a local branch that applies bounding box regression on each frame.

VCN-MA: We extend the original work [8] from one channel to multiple channels. Since, we have 35 classes in our dataset, we have in total 36 (35 activity classes + 1 background channel) channels. We also extended the spread loss such that it's applicable for multiple dimensions. [14] promotes the activation of the target class to be far away from the other classes by a margin factor. It does not 'force' the activations of the wrong classes to be far apart. We implement a multi dimensional variant, where all logits corresponding to multiple ground truth action classes are far away with significant margin, and retain this non-separation behaviour in the rest.

5.3. Evaluation metrics

We show performance on three metrics classification accuracy, f-mAP and v-mAP. Given IoU value p , f-mAP metric provides information about how many predicted frames have atleast spatial overlap of p from the ground truth. Similarly, v-mAP provides information about spatio-temporal overlap for different values of IoU.

To further analyze across different video datasets, it's important to ponder over how much scene bias is in the dataset, and, how important is the temporal ordering of frames. Thus, we perform two sets of experiments: 1) Freeze: We record the performance by taking the center frame and repeating it for the total number of frames present in the original video to analyze scene bias, 2) Reversal: We reverse the order of frames along temporal dimension. Then, we measure the absolute $(\sigma - \sigma')$ and relative $((\sigma - \sigma')/\sigma)$ drop from the original scores for accuracy, f-mAP and v-mAP. σ for f-mAP calculation is shown in equation 1.

$$\sigma_{0.5} = \frac{\sum_{i=1}^n (fmap > 0.5)}{n} \quad (1)$$

where $\sigma_{0.5}$ means it assigns 1 if IoU value is greater than 0.5 and n denotes the number of frames. $fmap$ is defined in equation 2. A_P and A_{GT} means predicted area and ground truth area.

$$fmap = \frac{|A_P \cap A_{GT}|}{|A_P \cup A_{GT}|} \quad (2)$$

5.4. Results

After training VCN-MA for 100 epochs, we report accuracy, fmap and vmap scores in Table 3. Additionally, we evaluate classification performance on several datasets in Table 2.

Further, we investigate scene bias and importance of temporal ordering. Firstly, for scene bias, we see the most absolute and relative drop in performance for f-mAP@0.5 for AVA dataset. For v-mAP@0.5, UCF101-24 has the most absolute and relative drop. Since, the network backbone is same, we reverse the order of frames. We flip the clip along the temporal dimension and then compare the performance with the original input. From Tables 4 and 5, we can see that the relative drop in f-mAP@0.5 is most for UCF101-24. In case of JHMDB-21, there's no drop in performance for f-mAP. We see the similar trend for AVA dataset as well. The performance drop is 1.2% but relative drop in performance is high.

Next, we analyze the class accuracy under different types of evaluations. We look into top 10 classes with highest score in normal evaluation. In UCF101-24 (Fig. 7), we see the drop in scores for frames freeze case, when there's a long interaction between actor and object (e.g. polevault, biking, trampoline-jump, and skateboarding). In JHMDB-21 (Fig. 8), we see performance drop specifically for activities that involves fast




Figure 7. UCF101-24 - Frame order class accuracy.

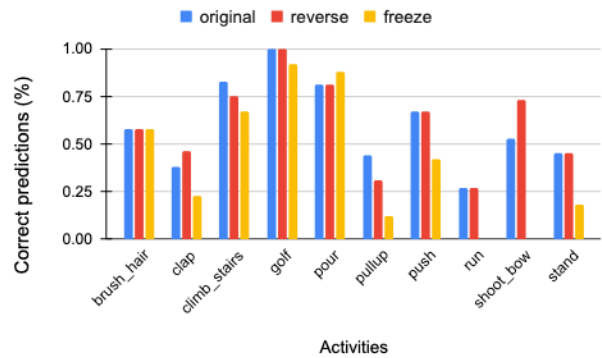


Figure 8. JHMDB-21 - Frame order class accuracy.

Dataset	UCF101-24	JHMDB-21	AVA [†]	MAMA [^]
Acc.	83.2	34.3	21.2	4.1

Table 2. Performance comparison for classifiers trained on different datasets using I3D backbone. [†]: AVA - (STEP - I3D). [^]: MAMA - (VCN-MA - I3D).

motion such as running and shooting bow in case of freezing the frames. In AVA dataset, walking, driving and swimming has almost zero accuracy in case of freezing the frame. If there's little or no motion, then the classification accuracy is comparable with normal evaluation protocol for example sit, lie, stand, and watch. Figure 5 illustrates the relative comparison of freezing and reversing the input frames on several datasets.

5.5. Challenges and Discussions

In this section, we list out several key challenges MAMA dataset list out that corroborates to real-world scenarios:

Video resolution: Most of the video datasets are low resolution, thus, action detection approaches are designed with input size of 224 or 256 resolution. Resizing a 1920x1080 video to 224x224 leads to information loss which makes




Figure 9. MAMA - Frame order class accuracy.

Dataset	Order	Acc	f-mAP	v-mAP
UCF101-24	Original	93.4	74.0	75.1
	Freeze	81.2	61.7	60.9
	Reverse	92.2	65.5	63.9
JHMDB-21	Original	38.8	68.2	68.4
	Freeze	32.6	56.9	56.2
	Reverse	38.1	68.1	66.7
AVA [†]	Original	21.2	20.1	-
	Freeze	7.7	7.4	-
	Reverse	20.8	18.9	-
MAMA [^]	Original	9.2	0.4	0.1
	Freeze	8.1	0.1	0.0
	Reverse	9.2	0.1	0.0

Table 3. Performance comparison for different types of frame ordering at f-mAP@0.5 and v-mAP@0.5. [†]: AVA results are shown for illustration purposes only. [^]: MAMA is a challenging dataset, and our baseline models yields low accuracy.

Dataset	Order	Acc	f-mAP	v-mAP
Absolute	UCF101-24	12.2	12.3	14.2
	JHMDB-21	6.2	11.3	12.2
	AVA [†]	13.5	12.7	-
	MAMA	1.1	0.3	0.07
Relative	UCF101-24	13.1	16.6	18.9
	JHMDB-21	16.0	16.5	17.8
	AVA [†]	63.7	63.2	-
	MAMA	12.0	75	-

Table 4. Absolute and relative drop in performance for scene bias condition at f-mAP@0.5 and v-mAP@0.5. [†]: AVA results are shown for illustration purposes only.

proposed approaches to work efficiently on this dataset.

Class imbalance: Diving deeper into real-world scenario, there’s always a huge class imbalance in activi-

Dataset	Order	Acc	f-mAP	v-mAP
Absolute	UCF101-24	1.2	8.7	11.2
	JHMDB-21	0.7	0.1	1.7
	AVA [†]	0.4	1.2	-
	MAMA	0.0	0.3	0.08
Relative	UCF101-24	1.3	11.8	14.9
	JHMDB-21	1.8	0.0	2.4
	AVA [†]	0.0	6.0	-
	MAMA	0.0	75	-

Table 5. Absolute and relative drop in performance for temporal reverse condition at f-mAP@0.5 and v-mAP@0.5. [†]: AVA results are shown for illustration purposes only.

ties. Some activities are more frequent than others. For example, talking and texting are activities have a higher frequency than activities such as close_trunk, people_embrace, and, close_facility_door.

Instance separation: Existing datasets have been collected under controlled number of instances. However, number of instances vary by a good margin between an indoor scene of a cafeteria vs an outdoor scene of a parking lot. A lot of actors are concentrated in a smaller area vs less actors spread over a large area.

6. Related work

[1] discusses the hierarchical organizations of activities, specifically breaking down an activity into varying types of behaviours exhibited by an actor. Motivated by this, datasets like AVA [12], classify actions as ‘atomic’ and try to capture subtle changes in actor motion by annotating at 1 sec intervals. On the other hand, datasets like UCF101 [31], JHMDB [17] perform annotations at finer spatio temporal scales. However, such datasets do not capture the subtle atomic actions. Similarly, a lot of relevant work [32] [4] has been done on detecting and removing bias by training networks on masked actor regions. However, understanding the sources of bias in the context of action *detection* is a relatively unexplored problem.

7. Conclusion

We have analyzed existing video datasets and shown the importance of temporal aspects for action understanding. Specifically, we have demonstrated how a detection model behaves when the temporal information is frozen or reversed. A reliable metric to compare datasets could be to measure the relative drops of models across datasets on several such properties. Furthermore, to contribute to the ongoing research in action detection, we have presented a new spatio-temporal dataset titled MAMA.

References

- [1] Roger G Barker and Herbert F Wright. Midwest and its children: The psychological ecology of an american town. 1955. **8**
- [2] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–970, 2015. **2**
- [3] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. **1, 2**
- [4] Jinwoo Choi, Chen Gao, Joseph CE Messou, and Jia-Bin Huang. Why can't i dance in the mall? learning to mitigate scene bias in action recognition. *Advances in Neural Information Processing Systems*, 32, 2019. **5, 8**
- [5] Kellie Corona, Katie Osterdahl, Roderic Collins, and Anthony Hoogs. Meva: A large-scale multimodal video dataset for activity detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1060–1068, 2021. **1, 3**
- [6] Ishan Dave, Zachaeus Scheffer, Akash Kumar, Sarah Shiraz, Yogesh Singh Rawat, and Mubarak Shah. Gabriellav2: Towards better generalization in surveillance videos for action detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 122–132, 2022. **2**
- [7] Ugur Demir, Yogesh S Rawat, and Mubarak Shah. Tinyvirat: Low-resolution video action recognition. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 7387–7394. IEEE, 2021. **2**
- [8] Kevin Duarte, Yogesh Rawat, and Mubarak Shah. Videocapsulenet: A simplified network for action detection. In *Advances in Neural Information Processing Systems*, pages 7610–7619, 2018. **1, 2, 5, 6**
- [9] Alaaeldin El-Nouby and Graham W Taylor. Real-time end-to-end action detection with two-stream networks. *arXiv preprint arXiv:1802.08362*, 2018. **1**
- [10] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6824–6835, 2021. **2**
- [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. **1, 2**
- [12] 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. **1, 3, 8**
- [13] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6546–6555, 2018. **1, 2**
- [14] Geoffrey E Hinton, Sara Sabour, and Nicholas Frosst. Matrix capsules with em routing. In *International conference on learning representations*, 2018. **6**
- [15] Rui Hou, Chen Chen, and Mubarak Shah. Tube convolutional neural network (t-cnn) for action detection in videos. In *IEEE International Conference on Computer Vision*, 2017. **1**
- [16] Haroon Idrees, Amir R Zamir, Yu-Gang Jiang, Alex Gorban, Ivan Laptev, Rahul Sukthankar, and Mubarak Shah. The thumos challenge on action recognition for videos “in the wild”. *Computer Vision and Image Understanding*, 155:1–23, 2017. **2**
- [17] Hueihan Jhuang, Juergen Gall, Silvia Zuffi, Cordelia Schmid, and Michael J Black. Towards understanding action recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 3192–3199, 2013. **1, 2, 8**
- [18] Vicky Kalogeiton, Philippe Weinzaepfel, Vittorio Ferrari, and Cordelia Schmid. Action tubelet detector for spatio-temporal action localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4405–4413, 2017. **1, 5**
- [19] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014. **2**
- [20] Hildegard Kuehne, Hueihan Jhuang, Estibaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *2011 International conference on computer vision*, pages 2556–2563. IEEE, 2011. **2**
- [21] Akash Kumar and Yogesh Singh Rawat. End-to-end semi-supervised learning for video action detection. *IEEE conference on computer vision and pattern recognition*, 2022. **1**
- [22] Dong Li, Zhaofan Qiu, Qi Dai, Ting Yao, and Tao Mei. Recurrent tubelet proposal and recognition networks for action detection. In *Proceedings of the European conference on computer vision (ECCV)*, pages 303–318, 2018. **1**
- [23] Bruce McIntosh, Kevin Duarte, Yogesh S Rawat, and Mubarak Shah. Visual-textual capsule routing for text-based video segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9942–9951, 2020. **1**
- [24] Mathew Monfort, Alex Andonian, Bolei Zhou, Kandan Ramakrishnan, Sarah Adel Bargal, Tom Yan, Lisa Brown, Quanfu Fan, Dan Gutfreund, Carl Vondrick, et al. Moments in time dataset: one million videos for event understanding. *IEEE transactions on pattern analysis and machine intelligence*, 42(2):502–508, 2019. **2**
- [25] 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. **1, 3**

- [26] Aayush J Rana and Yogesh S Rawat. We don't need thousand proposals: Single shot actor-action detection in videos. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2960–2969, 2021. **1, 2**
- [27] Mamshad Nayeem Rizve, Ugur Demir, Praveen Tirupattur, Aayush Jung Rana, Kevin Duarte, Ishan R Dave, Yogesh S Rawat, and Mubarak Shah. Gabriella: An online system for real-time activity detection in untrimmed security videos. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 4237–4244. IEEE, 2021. **2**
- [28] Mikel D Rodriguez, Javed Ahmed, and Mubarak Shah. Action mach a spatio-temporal maximum average correlation height filter for action recognition. In *2008 IEEE conference on computer vision and pattern recognition*, pages 1–8. IEEE, 2008. **1, 2**
- [29] Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *European Conference on Computer Vision*, pages 510–526. Springer, 2016. **2**
- [30] Khurram Soomro and Amir R Zamir. Action recognition in realistic sports videos. In *Computer Vision in Sports*, pages 181–208. Springer, 2014. **1**
- [31] 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. **1, 2, 3, 8**
- [32] Enzo Tartaglione, Carlo Alberto Barbano, and Marco Grangetto. End: Entangling and disentangling deep representations for bias correction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13508–13517, 2021. **8**
- [33] Praveen Tirupattur, Aayush J Rana, Tushar Sangam, Shruti Vyas, Yogesh S Rawat, and Mubarak Shah. Tinyaction challenge: Recognizing real-world low-resolution activities in videos. *arXiv preprint arXiv:2107.11494*, 2021. **2**
- [34] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Computer Vision (ICCV), 2015 IEEE International Conference on*, pages 4489–4497. IEEE, 2015. **1**
- [35] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6450–6459, 2018. **1**
- [36] Shruti Vyas, Yogesh S Rawat, and Mubarak Shah. Multi-view action recognition using cross-view video prediction. In *European Conference on Computer Vision*, pages 427–444. Springer, 2020. **1**
- [37] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018. **1**
- [38] Philippe Weinzaepfel, Zaid Harchaoui, and Cordelia Schmid. Learning to track for spatio-temporal action localization. In *Proceedings of the IEEE international conference on computer vision*, pages 3164–3172, 2015. **1**
- [39] Philippe Weinzaepfel, Xavier Martin, and Cordelia Schmid. Human action localization with sparse spatial supervision. *arXiv preprint arXiv:1605.05197*, 2016. **1, 2**
- [40] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *Proceedings of the European conference on computer vision (ECCV)*, pages 305–321, 2018. **1**
- [41] Chenliang Xu, Shao-Hang Hsieh, Caiming Xiong, and Jason J Corso. Can humans fly? action understanding with multiple classes of actors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2264–2273, 2015. **1**
- [42] Xitong Yang, Xiaodong Yang, Ming-Yu Liu, Fanyi Xiao, Larry S Davis, and Jan Kautz. Step: Spatio-temporal progressive learning for video action detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 264–272, 2019. **1, 2, 6**
- [43] Serena Yeung, Olga Russakovsky, Ning Jin, Mykhaylo Andriluka, Greg Mori, and Li Fei-Fei. Every moment counts: Dense detailed labeling of actions in complex videos. *International Journal of Computer Vision*, 126(2):375–389, 2018. **2**
- [44] Lijun Yu, Yijun Qian, Wenhe Liu, and Alexander G Hauptmann. Argus++: Robust real-time activity detection for unconstrained video streams with overlapping cube proposals. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 112–121, 2022. **2**