

This WACV 2021 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

MEVA: A Large-Scale Multiview, Multimodal Video Dataset for Activity Detection

Kellie Corona, Katie Osterdahl, Roderic Collins, Anthony Hoogs Kitware, Inc. 1712 Route 9, Suite 300, Clifton Park, NY 12065 USA {firstname.lastname}@kitware.com



Figure 1. Examples of MEVA data, showing an approximately synchronized view in four data modalities. Left-to-right:a 4x3 montage of RGB and thermal IR cameras; GPS locations of about 90 actors overlaid on five minutes of GPS tracks; a cropped view from a UAV.

Abstract

We present the Multiview Extended Video with Activities (MEVA) dataset[6], a new and very-large-scale dataset for human activity recognition. Existing security datasets either focus on activity counts by aggregating public video disseminated due to its content, which typically excludes same-scene background video, or they achieve persistence by observing public areas and thus cannot control for activity content. Our dataset is over 9300 hours of untrimmed, continuous video, scripted to include diverse, simultaneous activities, along with spontaneous background activity. We have annotated 144 hours for 37 activity types, marking bounding boxes of actors and props. Our collection observed approximately 100 actors performing scripted scenarios and spontaneous background activity over a threeweek period at access-controled venue, collecting in multiple modalities with overlapping and non-overlapping indoor and outdoor viewpoints. The resulting data includes video from 38 RGB and thermal IR cameras, 42 hours of UAV footage, as well as GPS locations for the actors. 122 hours of annotation are sequestered in support of the NIST Activity in Extended Video (ActEV) challenge; the other 22 hours of annotation and the corresponding video are available on our website, along with an additional 306 hours of ground camera data, 4.6 hours of UAV data, and 9.6 hours of GPS logs. Additional derived data includes camera models geo-registering the outdoor cameras and a dense 3D point cloud model of the outdoor scene. The data was collected with IRB oversight and approval and released under a CC-BY-4.0 license.

1. Introduction

It has been estimated that in 2019, 180 million security cameras were shipped worldwide [9], while the attention span of a human camera operator has been estimated at only 20 minutes [5, 3]. The gap between the massive volume of data available and the scarce capacity of human analysts has been closed, but not eliminated, by the rapid advancement of computer vision techniques, particularly deep-learning based methods. Fundamental progress is often spurred by datasets such as ImageNet [15] and MS COCO [8] for object recognition and MOT16 [10] and Caltech Pedestrian [2] for pedestrian detection and tracking.

However, as discussed in Section 2, datasets for action recognition typically do not address many of the needs of the public safety community. Datasets such as AVA [4], Moments in Time [11], and YouTube-8m [1] present videos which are short, high-resolution, and temporally and spatially centered on the activities of interest. Rigorous research and evaluation of activity detection in public safety and security video data requires a dataset with realistic spatial and temporal scope, yet containing sufficient instances of relevant activities. In support of evaluations such as the NIST Activities in Extended Video (ActEV) [12], we designed the MEVA dataset to explicitly capture video of large groups of people conducting scripted activities in realistic settings in multiple camera views, both indoor and outdoor. We defined 37 activity types, shown in Figure 2, ranging from simple, atomic single-actor instances to complex, multi-actor activities. Our fundamental resources included approximately one hundred actors on-site at an access-controlled facility with indoor and outdoor venues for a total of around three weeks, recorded by 38 groundlevel cameras and two UAVs. Additionally, actors were provided with GPS loggers. We conducted extensive precollect planning, including a pilot collection exercise, to develop the appropriate level of actor direction required to maximize instances of our 37 activity types while maintaining realism and avoiding actor fatigue.

The final dataset contains over 9300 hours of groundcamera video, 42 hours of UAV video, and over three million GPS trackpoints. In support of the ActEV evaluation, we have annotated 144 hours of this video for the 37 activity types. While most of this data is sequestered for ActEV, we have released 328 hours of ground-camera data, 4.6 hours of UAV data, and 22 hours of annotations via the MEVA website. The dataset design and collection underwent rigorous IRB oversight, with all actors signing consent forms. The data is fully releasable under a Creative Commons Attribution 4.0 (CC-BY-4.0) license, and we believe represents by far the largest dataset of its kind available to the research community.

Figure 1 shows a sample of data available from the MEVA website, approximately synchronized in time, col-

lected during a footrace scenario. The left side shows a montage of 12 of the 29 released RGB and thermal IR cameras, illustrating the diversity in locations, settings (indoors and outdoors), as well as overlapping viewpoints. The middle illustration plots approximately 90 GPS locations on a background image. The right image is a crop from UAV footage.

The paper is organized as follows. Sections 2 places our work in context with similar efforts. Section 3 discusses how we designed the dataset to maximize realism and activity counts. Section 4 describes our annotation process. Section 5 briefly describes the ActEV leaderboard results on MEVA.

2. Prior Work

We distinguish "focused" activity recognition datasets such as [4, 11, 1], which typically contain single, short activities, from security-style video, which is typically longduration and ranges from long stretches of low or no activity to busy periods with high counts of overlapping activities. Figure 3 shows how MEVA advances the state of the art along several security dataset factors, notably: duration, number of persistent fields of view (both overlapping and singular), modalities (EO, thermal IR, UAV, hand-held cameras, and GPS loggers), and annotated hours. The UCF-Crime dataset [17] presents real incidents from real security cameras, but only from a single viewpoint at the point of activity; little or no background data is available. The VIRAT Video Dataset [13] contains both scripted and spontaneous activities, but without overlapping viewpoints. The Duke MTMC dataset [14] presents security-style video from multiple viewpoints, but with only spontaneous, unscripted activity.

3. Dataset Design

The MEVA dataset was designed to capture human activities, both spontaneous and scripted, in an environment as close as possible to that encountered by deployed CCTV systems in the real world. The dataset was designed to be realistic, natural and challenging for security and public safety video research in terms of its resolution, background clutter, visual diversity, and activity content. As discussed in Section 3.1, the data is fully releasable under a Creative Commons Attribution 4.0 (CC-BY-4.0) license, facilitated by a detailed human subjects research plan. Activity diversity was achieved through pre-collect scripting for scenarios and activity types (Section-3.2), and collected via an ambitious camera plan with 38 ground-level cameras and two UAVs (Section-3.3). Scene diversity was accomplished by scripting scenarios to occur at different times of day and year to create variations in environmental conditions, and directing demographically diverse actors to perform activi-

person_abandons_package *	hand_interacts_with_person	person_reads_document	vehicle_picks_up_person
person_carries_heavy_object	person_interacts_with_laptop	person_rides_bicycle	vehicle_reverses
person_closes_facility_door	person_loads_vehicle	person_puts_down_object	vehicle_starts
person_closes_trunk	person_transfers_object	person_sits_down	vehicle_stops
person_closes_vehicle_door	person_opens_facility_door	person_stands_up	vehicle_turns_left
person_embraces_person	person_opens_trunk	person_talks_on_phone	vehicle_turns_right
person_enters_scene_through_structure	person_opens_vehicle_door	person_texts_on_phone	vehicle_makes_u_turn
person_enters_vehicle	person_talks_to_person	person_steals_object *	
person_exits_scene_through_structure	person_picks_up_object	person_unloads_vehicle	
person_exits_vehicle	person_purchases	vehicle_drops_off_person	

Figure 2. List of the 37 activities defined and annotated in MEVA. Activities marked with an asterisk are threat-based activities.

	VIRAT [13]	UCF-101 Untrimmed [16]	Duke MTMC [14]	UCF-Crimes [17]	MEVA
Number of activity types	23	101	-	13	37
Range of samples per type	10-1500	90-170	-	50-150	5-750 (1200)
Incidental objects and activities	yes	no	yes	yes	yes
Natural background behavior	yes	no	yes	yes	yes
Tight bounding boxes	yes	no	no	no	yes
Max resolution	1920x1080	320x240	1920x1080	320x240	1920x1080
Sensor modalities	1	1	1	2	5
Security	yes	no	yes	yes	yes
Number of FOV	17	unique-per-clip	8	unique-per-clip	28
Overlapping FOV	no	no	yes	no	yes
Indoor & outdoor	no	yes	no	yes	yes
Availability	direct	reference	retracted	direct	direct
Clip length	2-3 minutes	1-71 seconds	-	4 minutes	5 minutes
Dataset duration (hours)	29	26.6	10	128	9300 / 144

Figure 3. Comparison of characteristics of activity and security datasets to MEVA. The reported range of activity counts per type is only for MEVA released data while the average number of samples in parenthesis includes both released and sequestered annotations. The 5 sensor modalities included in the MEVA dataset are static EO, UAV EO, Thermal IR, Handheld or Body-Worn Cameras (BWC), and GPS. There are 28 unique FOV with additional FOVs offered by drone, handheld and BWC footage. The overall duration of video collected as part of MEVA is 9300 hours while 144 hours were annotated.



Figure 4. Site map with released cameras and approximate fieldsof-view. Indoor cameras are blue circles; triangles are outdoor cameras. Red are co-located EO/IR; blue is fixed-focal-length; green is PTZ in stationary mode; purple is PTZ in patrol mode. Pink fields of view are outdoors; yellow are indoors.

ties in multiple locations with varied natural behaviors. The ground-camera plan includes a variety of camera types, indoor and outdoor viewpoints, and overlapping and singular fields of view. Realism was enhanced by scripting to include naturally occurring activity instances, frequent incidental moving objects and background activities, as discussed in Section-3.2. Finally, the quantity of data collected was enabled by planning for a multi-week data collection effort to allow for a large set of activity types with a large number of instances per activity class as discussed in Section-3.3.

3.1. Releasability

A critical requirement for the MEVA dataset is broad releasability to the research community. Possible restrictions include prevention of commercial use, which discourages commercial companies from using the dataset for research or collaborating on teams which do use the dataset. Another restriction is when data containing human activity is collected without oversight of an Institutional Review Board (IRB) or collected outside of its initially IRB-approved conditions. While the data may address other needs of the research community, lack of proper IRB oversight or participant consent can render the data unusable, wasting significant resources in the collection and curation process.

The MEVA dataset was carefully crafted to avoid such restrictions to maximize its value to the research community. The data collection effort was overseen by an IRB to ensure that the data collected would be usable. We wrote a human subjects research protocol and informed consent documents that detailed the collection process, what types of data would be collected, how the data will be released and may be used, and the risks and benefits to anyone consenting to participate in the study. The informed consent document was provided to prospective participants in advance and reviewed in an informed consent briefing session prior to beginning the data collection. Included in the protocol and consent briefing was a sample license that may be used for the data, a CC-BY 4.0, to make it fully clear the intent and potential future use of the data. The consent presentation was followed by small group discussions and question/answer periods to thoroughly address and answer any participant concerns prior to them consenting to participate the research. Once the individuals who voluntarily chose to participate signed the consent forms, we were able to begin the data collection process.

To further ensure broad releasability of the data, the collection occurred at an access-controlled facility to minimize non-participating individuals entering the fields of view of the cameras and appearing in the data.

3.2. Scripting

To guarantee collection of minimum activity instance counts, ensure diversity in the data, and produce the most realistic video data, MEVA data was collected for both scripted and unscripted activities. The scripting challenges for the data collection effort can broadly be broken down into two categories: (1) scripting for satisfying program requirements such as quantity and diversity, and (2) scripting for realism in activity behavior.

It was critical to design the data collection to represent a variety of realistic scenarios in ground-based security video collection. We determined seven overarching scenarios of interest to the public safety community. For example a footrace with multiple phases including registration, participant arrival, race and cleanup. The scenarios were collected over a period of three weeks on different days and at different times of days to capture video variations due to weather and lighting. At the extreme, two sub-collects in March and May produced dramatic contrast in weather from snow to extreme heat, and thus diversity in natural human behavior (e.g., clustering together in the cold) and wardrobe (i.e., jackets in the winter and t-shirts in the summer).

Scripted into these scenarios were 37 primitive, complex and threat-based human activity types to ensure minimum



Figure 5. Two approximately synchronized views from the gym.

instance counts useful to the program; these are listed in Figure 2. These activities were scripted to be performed by actors with different age, sex and ethnicity as allowed by the actor pool. Activities involving vehicles or props were specifically scripted to rotate the vehicle pool and theatrical property. Scripted activities occurred across various scenarios and at multiple locations in the facility. Activities were scripted in overlapping and singular fields of view, indoor and outdoor cameras, and various camera types (e.g., stationary EO and IR, roaming EO). Activities were scripted to be performed by the same individual in different, singular cameras with the goal of associating the same individual across different cameras. This is relevant when a central activity is detected involving a subject, and then a derived complex activity is to detect other individuals that meet with the original subject of interest. Scenarios also included confusers for the scripted activity types to add value for performer evaluation in the final dataset. Figure 5 shows two approximately synchronized views from the gym.

Though scenarios had an overarching theme, multiple iterations of the same scenario (called takes) were designed to produce visibly distinct but comparable results by varying the parameters described above. Additionally, actor density at various location in the facility and traffic flow were modified to produce differences in subsequent takes of the same scenario. Scripting the scenarios and activities in this level of detail guaranteed the diversity and quantity of data satisfied program requirements.

To address the second broad goal, we hired professional actors to act scenarios in a natural manner to create an equivalent challenge for computer vision event detectors, as if the events occurred in real-life scenarios without acting. The actors were divided into squads which each had a squad lead in charge of managing a group of 5-8 individuals. The squads were designed, grouped and shuffled to produce variations in demographics, behaviors and inter-actor behaviors. Each squad was given a direction card which provided information on timing and locations for activities to be performed, and group behavior such as social groupings within the squad. By giving the squad lead the ability to designate roles within the group, natural social dynamics, such as those observed in married couples or families, were preserved and increased the realism of the collected data. The squad lead was responsible for assigning roles to the actors, managing vehicle and prop assignment, providing feedback to individual actors for corrections on activity behavior, and reporting descriptions of complex or threatbased activities for use in expediting later annotation.

We found that providing scripting at the squad rather than individual level produced the most realistic actor behavior. When actor behavior was scripted on the individual level at high temporal and spatial resolution, the actors, overwhelmed by the details, focused on achieving all activities at the specified time and location rather than performing activities that were natural in appearance. Using a small pilot collect of data, we were able to determine which activities occurred naturally in the data collection without direction (e.g., person_opens_door). Taking this into account when scripting, lower instances of these activities were scripted and a mix of scripted and unscripted occurrences were collected. On the other hand, rare and threat-based activities needed to be specifically scripted to ensure minimum counts of these activities were achieved and later located for annotation. However, allowing the actors to have artistic licence with how to accomplish activities, especially complex activities, produced more a natural and diverse dataset. Actor familiarity with the scripted scenario also increased realism. Most scenarios varied in time from 15 minutes to 1 hour. Multiple takes could be run consecutively with minor modifications to scripted behaviors and activity instances, and minimal reset time (5 to 10 minutes) to efficiently produce varied instances of a scenario and communicate modifications to the squad leads.

In addition to activity-focused scripting, scene-level behaviors also needed to be scripted to ensure realism. For example, train and bus arrival schedules were designed to reflect public transit schedules scaled for the population depicted by our actors. Vehicle traffic was scripted to have ebbs and flows associated with similar patterns to real-world versions scripted scenarios. For example, an increase in ve-



Figure 6. Enrollment photos were taken daily of all actors, with (a) and without (b) outerwear, from the front and back. Photos included actor numbers, which correspond to their unique GPS logger number for reidentification purposes.

hicles dropping off people at the bus station just prior to a bus's scheduled arrival. Driving routes, vehicles and drivers were shuffled to provide variations between different scenarios and takes. Again, these tasks were assigned at the squad level and managed by squad leads.

In addition to scripted scenarios, periods of completely unscripted data were collected with the aim of collecting naturally occurring common activities with no actor direction. We also took advantage of unscripted data to interject threat-based and rare activities to capture undirected response by casual observers. These unscripted times were naturally occurring and added value to reset times between scenario takes.

3.3. Collection

The MEVA video dataset was collected over a total of three weeks at the Muscatatuck Urban Training Center (MUTC) with a team of over 100 actors and 10 researchers. MUTC is an access-controlled facility run by the Indiana National Guard that offers a globally unique, urban operating environment. The real and operational physical infrastructure, including curbed roads and used buildings, set it apart from other access-controlled facilities for collecting realistic video data.

The camera infrastructure included commercial-off-theshelf EO cameras; thermal IR cameras as part of several IR-EO pairs; two DJI Inspire 1 v2 drones, and a range of video and still images from handheld cameras used by the actors. The fields of view, both overlapping and non-overlapping, capture person and vehicle activities in indoor and outdoor settings.

Onsite staging and actor instruction also improved the realism in scenes for the MEVA scenario collection. Areas were staged with props, such as signs, tables, chairs, and trashcans, to provide scenario-specific context necessary to increase the natural appearance and usefulness of the area to both actors and recording cameras. Additionally, areas were



Figure 7. The actor in the purple jacket, actor 544 in Figure 6, is visible in multiple cameras during the scenario. Her height is (a) 301 pixels, (b) 118 pixels, (c) 176 pixels, and (d) 89 pixels in each of the respective fields of view.

staged to look similar but distinct between different scenarios to add diversity to the dataset collected. For example, the foyer of one building was converted into an operational cafe where the actors were able to grab drinks and snacks during scripted or unscripted camera time.

Actors were selected to provide a diverse pool of individuals in background, ethnicity and gender. As part of our iterative process to receive and incorporate feedback, all actors attended briefings in which camera fields of view and common issues (e.g., erratic driving or aimless walking) were discussed in a group setting, raising awareness of key factors required to collaboratively produce realistic behaviors in video. Squad leads also held smaller briefings, providing feedback to the scripting team and receiving instructions to pass to their squad members. During morning briefings, all actors and MEVA team members were registered with a unique identifying number matched to their GPS logger unit and wardrobe enrollment photograph(s) as shown in Figure 6. The photos, paired with the GPS logs, enable reidentification of individual actors across fields of view in a single scenario, as seen in Figure 7. Each day before filming, squad leads were given direction cards to assign roles to their actors and inform prop distribution, then data collection was performed. Shorter scenario takes of less than 1 hour were performed in rapid succession with only minor modifications relayed during a brief reset. One day short-scenario data collection would contain between 4 and 8 takes; for longer scenarios of approximately 8 hours, the scenario would be repeated across multiple sequential days. Direction and course correction was provided by both directly interjecting a MEVA team member into the scenario for a brief interval or incorporating actor briefings naturally into the scenario schedule.

The final dataset contains 9,303 hours of ground-camera video (both EO and thermal IR), 42 hours of UAV footage, and 46 hours from hand-held and body-worn cameras. Additionally we collected over 2.7 million GPS trackpoints from 108 unique loggers.

4. Annotation

Annotating a large video dataset requires balancing quality and cost. Annotating for testing and evaluation, as is the case for MEVA, must prioritize the quality of ground truth annotations, as any evaluation will only be as reliable as the ground truth derived from the annotations. Ground truth annotations must minimize missed and incorrect activities to reduce the potential for penalizing correct detections, strictly adhere to activity definitions, and address corner cases to reduce ambiguity in scoring. We annotated MEVA by first localizing spatio-temporal activities and then creating bounding boxes for objects involved in the activity. We optimized a multi-step process for quality and cost: (1) an annotation step to identify the temporal bounds for activities and objects involved, (2) an audit step by experts to ensure completeness and accuracy of the activity annotations, (3) a crowdsourced method for bounding box track creation for objects, and (4) a custom interface for rapid remediation of corner cases by experts, with automated checks for common human errors between each of these steps.

4.1. Activity Annotation

We annotated MEVA through a sequential process of activity identification by a trained, dedicated team of thirdparty annotators, followed by a quality control audit from an internal team of MEVA experts intimately involved in defining the activities. The initial activity identification step detects all MEVA activities in a 5 minute video clip, specifying the start and end times for each activity, and identifying the initial and final location for each actor or object involved. The subsequent audit step catches any missed or



Figure 8. Half an hour of activities overlaid on a background image. 112 activities, 17 types, ranging from 1 to 14 instances.

incorrectly labeled activity instances, and confirms strict adherence to MEVA definitions.

Ambiguity in activity definitions is a fundamental annotation challenge. To minimize this, we define activities very explicitly and require that events are always annotated on pixel evidence and not human interpretation. We developed detailed annotation guidelines for each activity, including definition of the start and end times for an activity, and the actors and objects involved. Definitions also include extensive discussion of corner cases.

We explored several methods to increase annotation efficiency and quality, including completely crowdsourced methods on Amazon Mechanical Turk (AMT) and a solely in-house team of experts. The optimized annotation process used a team of third-party annotators dedicated to MEVA annotation to lower the overall cost and supply surge capacity via dynamic team scaling while still guaranteeing quality results. After comparing the results of multiple annotators performing activity annotation in parallel and combined results in a post-processing step with annotating in series with a quality assurance step, serial annotation was selected to enable efficiently annotating a larger dataset.

Annotation consistency is essential to providing highquality ground truth annotations. When using a diverse team of annotators, it was essential to have a clear set of guidelines and iterative definition modifications to include corner cases to ensure annotator agreement across the team. Additionally, we found that project-specific training, including effective use of the annotation tools, in-depth discussion of definitions including examples, and iterative feedback on annotation quality produced improvements in annotation quality and speed. Finally, a quality audit from an internal team of MEVA annotators was used to guarantee no missed or incorrect instances. We observed a 3-fold increase in audit efficiency over a two month period of working in a tight feedback loop with these procedures and a project-dedicated team of annotators.

4.2. Track Annotation

Once activities and participating objects were identified by MEVA experts, objects were tracked with bounding boxes for the duration of their involvement in the activity. The annotation of bounding boxes was primarily conducted through crowdsourced annotation on AMT, via an iterative process of bounding box creation and quality review.

In the first step of bounding box refinement, videos were broken up into segments that corresponded to a single activity annotation. An AMT task was created for each activity, displaying a set of start and end boxes with linearly interpolated boxes on intervening frames for all objects involved in the activity. Workers were instructed to complete the tracks by annotating bounding boxes for the interpolated frames which were as tight as possible around the object; for example, all the visible limbs of a person should be in the bounding box while minimizing "buffer" pixels. These two characteristics are required in the MEVA dataset to ensure that high-quality activity examples with minimal irrelevant pixels are provided for testing and evaluation. These traits were enforced by using other AMT annotators to assess the quality of the resulting tracks.

In the quality review step, the bounding box annotations were shown to several AMT workers who were asked to evaluate them as acceptable according our guidelines. If agreement was achieved between the AMT workers, then the results were considered acceptable and complete; otherwise, a new AMT task would be created for bounding box refinement. The results of the next round of refinement would be provided for quality review, and the process would repeat until acceptable annotations were produced or a threshold number of iterations were performed. If the threshold number of iterations were performed, the clips would be vetted and edited manually as needed by an MEVA auditor. The percentage of activities requiring expert intervention was less than 5%. Specialized web interfaces were developed for each of these tasks to allow workers (AMT or in-house) to easily provide the necessary results. In order to eliminate most systematic low quality jobs, annotations were sampled and workers were allowed to continue based on quality.

Figure 9 illustrates bounding box annotations for tracks from a variety of activities and fields of view that demonstrate the quality produced by our track annotation procedure. Box size varied dramatically due to the scale variation in the video, with a mean area of 13559 ± 23799 . Annotations span 5 track types (person, vehicle, other, bag, and bicycle) with a distribution of 90.71%, 4.5%, 4.51%, 0.15%, and 0.05%, respectively. As part of quality control, we compared MEVA annotations against high-confidence performer false alarms from the ActEV evaluation; our false negative rate (i.e., confirmed missed instances to total activity count in reviewed clips) is less than 0.6%.



Figure 9. Examples of activities and tracks from diverse fields of view. The font size in each images is consistent, indicating varying scale of the annotations.



Figure 10. Visualization of the fine-grained 3D point cloud model of the collection site.

Additional Data In addition to the annotations, we have provided supplemental data such as camera models for the camera models which register into a common georeferenced coordinate system. We have also provided a 3D model of the outdoor component of the collection site, provided as a PLY file and visualized in Figure 10.

5. Baseline Activity Detection Results

The NIST ActEV challenge [12] is using the MEVA dataset for its ongoing Sequestered Data Leaderboard (SDL). The ActEV challenge defines three tasks: Activity Detection (AD) with no spatial localization in the video, Activity and Object Detection (AOD) where the activity and participating objects are spatially localized within a frame but not necessarily correlated across frames, and Activity / Object Detection and Tracking (AODT), which extends AOD to establish real-world activity and object identity across frames. The current leaderboard is for AD, scored using Probability of Miss (pMiss), the proportion of activ-



Figure 11. The NIST ActEV leaderboard as of 9 Nov 2020, computed on MEVA data. Better performance is lower and to the left.

ities which were not detected for at least one second, vs. Time-based False Alarm (TFA), the proportion of time the system detected an activity when there was none. Figure 11 shows results current at time of writing for nine teams plus a baseline implementation [7] based on RC3D [18]. As the scores indicate, the MEVA dataset is very difficult compared to related datasets, presenting abundant opportunities for innovation and advancement in activity detection, tracking, re-identification and other problems.

6. Conclusion

We have presented the MEVA dataset, a large-scale, realistic video dataset containing annotation of a diverse set of visual activity types. The MEVA video dataset surpasses existing activity detection datasets in hours of video, number of cameras providing overlapping and singular fields of view, variety of sensor modalities, and broad releasability. The dataset also provides a substantial 144 hours of evaluation-quality activity annotations of scripted and naturally occurring activities. We believe that with these traits the MEVA dataset will stimulate diverse research within the computer vision community. The MEVA dataset is available at: http://mevadata.org.

Acknowledgement: This work is supported by the Intelligence Advanced Research Projects Activity (IARPA) via contract 2017-16110300001. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA or the U.S. Government.

References

- Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. Youtube-8m: A large-scale video classification benchmark, 2016.
- [2] Piotr Dollár, Christian Wojek, Bernt Schiele, and Pietro Perona. Pedestrian detection: A benchmark. In 2009 IEEE Conference on Computer Vision and Pattern Recognition, pages 304–311. IEEE, 2009.
- [3] Mary W Green. The appropriate and effective use of security technologies in US schools: A guide for schools and law enforcement agencies. US Department of Justice, Office of Justice Programs, National Institute of ..., 1999.
- [4] Chunhui Gu, Chen Sun, David A. Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, Cordelia Schmid, and Jitendra Malik. AVA: A video dataset of spatio-temporally localized atomic visual actions. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [5] Niels Haering, Péter L Venetianer, and Alan Lipton. The evolution of video surveillance: an overview. *Machine Vision* and Applications, 19(5-6):279–290, 2008.
- [6] Kitware. Multiview Extended Video with Activities. http: //mevadata.org.
- [7] Kitware. R-C3D Baseline for ActEV. https://gitlab. kitware.com/kwiver/R-C3D.
- [8] 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 *European conference on computer vision*, pages 740–755. Springer, 2014.
- [9] IHS Markit. Security technologies top trends for 2019. https://technology.informa.com/607069/ video-surveillance-installed-base-report-2019.
- [10] Anton Milan, Laura Leal-Taixé, Ian Reid, Stefan Roth, and Konrad Schindler. Mot16: A benchmark for multi-object tracking. arXiv preprint arXiv:1603.00831, 2016.
- [11] 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.
- [12] NIST. ActEV: Activities in Extended Video. https:// actev.nist.gov.
- [13] Sangmin Oh, Anthony Hoogs, Amitha Perera, Naresh Cuntoor, C.-C. Chen, Jong Taek Lee, Saurajit Mukherjee, J. K. Aggarwal, Hyungtae Lee, Larry Davis, Eran Swears, Xioyang Wang, Qiang Ji, Kishore Reddy, Mubarak Shah, Carl Vondrick, Hamed Pirsiavash, Deva Ramanan, Jenny Yuen, Antonio Torralba, Bi Song, Anesco Fong, Amit Roy-Chowdhury, and Mita Desai. A Large-scale Benchmark Dataset for Event Recognition in Surveillance Video. In Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2011.

- [14] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *European Conference* on Computer Vision, pages 17–35. Springer, 2016.
- [15] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [16] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *CoRR*, abs/1212.0402, 2012.
- [17] Waqas Sultani, Chen Chen, and Mubarak Shah. Real-world anomaly detection in surveillance videos. In *Proceedings* of the IEEE Conference on Computer Vision and Pattern Recognition, pages 6479–6488, 2018.
- [18] Huijuan Xu, Abir Das, and Kate Saenko. R-c3d: Region convolutional 3d network for temporal activity detection. In *Proceedings of the IEEE international conference on computer vision*, pages 5783–5792, 2017.