

Cooking beyond Frames: A Stereo Event Camera Dataset in the Kitchen

Chengming Feng¹, Hesam Araghi¹, Liming Zheng¹, Julien Dupeyroux²,
Xucong Zhang¹, Jan van Gemert¹, and Nergis Tömen¹

¹ Delft University of Technology, The Netherlands
{c.feng-1, n.tomen}@tudelft.nl

² STMicroelectronics, France

Abstract. Event cameras, also known as neuromorphic cameras, have gained significant attention in recent years due to their high temporal resolution, high dynamic range, and low power consumption. While many studies and datasets in neuromorphic vision have focused on automotive and drone applications, human-centric daily-life scenarios remain largely underrepresented, despite their importance for developing and benchmarking event-based perception systems. Moreover, the few existing event-based human activity datasets are typically recorded with scripted human actions, limiting their ability to capture natural human behaviors. In this paper, we introduce **EventKitchen**, a large-scale stereo event camera benchmark dataset of human cooking activities in the kitchen. EventKitchen is egocentrically collected from 10 participants in 13 diverse kitchens, where the participants wear a helmet with multiple sensors and naturally perform cooking activities, without any scripted actions. EventKitchen comprises 5.5 hours of stereo event recordings with synchronized RGB, depth, and IMU data. We provide human annotations for 10,762 action segments and 13,482 bounding boxes. We train baseline models on EventKitchen to perform multiple event-based tasks, including action recognition, object detection, and stereo depth estimation. By capturing natural, real-world human activities, EventKitchen establishes a challenging benchmark for neuromorphic vision beyond autonomous driving. The dataset and toolkit are available at <https://chengmingf.github.io/EventKitchen.github.io/>

1 Introduction

Unlike standard frame-based RGB cameras, event cameras operate asynchronously, where each pixel independently responds to intensity changes. This sensing paradigm of event cameras provides extremely high temporal resolution, high dynamic range, and low power consumption [19]. As a result, event cameras excel at capturing rapid object movements with reduced motion blur, and preserving fine object features under challenging illumination conditions, significantly outperforming traditional frame-based sensors in dynamic and high-contrast scenarios [19, 22, 33, 60, 64]. Due to these advantages, event cameras are popular in




Fig. 1: (a) **Data collection device:** The device is a wearable bicycle helmet with mounted sensors including two event cameras, two CMOS RGB cameras, and one Intel RealSense depth camera to allow multi-modal sensing; (b) **Raw data:** We show seven example raw data streams collected with the device in three out of 13 diverse kitchens, columns ‘IMU’, ‘D-RGB’, and ‘Depth’ are from Intel RealSense, ‘Left RGB’ and ‘Right RGB’ from two CMOS stereo cameras, ‘Left Event’ and ‘Right Event’ from stereo event cameras; (c) **Human annotations:** We show example human annotations including action labels and bounding boxes, which enables the tasks of event-based action recognition and object detection. The collected depth maps can be rectified to both left and right event cameras to support the event-based stereo depth estimation.

autonomous driving [1, 8, 13, 22, 23, 27, 57, 71]. This, in turn, led to the development of several large-scale, real-world datasets, such as 1Mpx [46] and GEN1 [13], which target event-based object detection, as well as multi-task datasets like DSEC [23] and MVSEC [71]. In addition, the low-cost computation makes event cameras favorable in robotics and drone research [8, 14, 17, 18, 45, 71].

In contrast, the study of event-based human daily-life activities remains significantly underexplored, despite their importance for developing and benchmarking neuromorphic perception systems in human-centric environments [10].

Although event cameras have demonstrated advantages in capturing fine-grained and fast human motion for action recognition [2, 9, 50, 68, 70] and pose estimation [25, 39, 40], existing datasets are limited in scale, realism, and environmental diversity. Most of the event-based human activity datasets either are collected under scripted laboratory settings that fail to capture the variability and spontaneity of natural human behavior [2, 3, 20, 36, 38, 61, 67, 68], or from fixed viewpoints that restrict scene dynamics and yield sparse background events [2, 3, 5, 20, 36, 38]. To facilitate dataset construction, event simulators [21, 28, 53] have been widely used to generate synthetic event streams from RGB inputs. For instance, N-EPIC-Kitchens [50] creates simulated events by passing EPIC-KITCHENS [11] videos through ESIM [53]. However, prior works [7, 48, 62] have reported notable sim-to-real gaps, with models trained on simulated events often exhibiting degraded performance when evaluated on real-world recordings. Moreover, simulated event data are inherently constrained by the frame rate and dynamic range of the source RGB videos, cannot faithfully reproduce real-world sensor noise, and fail to fully capture the distinctive properties of real event streams. Other efforts, such as UCF-Crime-DVS [51], capture real event streams by recording RGB videos displayed on monitors. While this setup employs real event sensors, the captured data remain constrained by the parameters of the monitor, such as refresh rate, panel type, and dynamic range, which makes them less representative of real-world dynamics [36, 58]. These limitations highlight the need for larger, diverse, and complex benchmark datasets that represent human activity in real-world daily-life scenarios, enabling large-scale pre-training and systematic benchmarking for neuromorphic vision beyond the autonomous driving domain.

In this paper, we introduce EventKitchen: a large-scale, real-world, egocentric, stereo event camera dataset. The key recording device is a pair of stereo Prophesee Gen4 event cameras, which provide high-resolution event data. To provide reliable ground truth for the event data, our setup is multi-modal: alongside an event-based stereo stream, it features an RGB-based stereo stream, an RGB-depth pair, and an Inertial Measurement Unit (IMU), as shown in Fig. 1. All sensors are mounted on a wearable helmet to support egocentric data collection. Compared to fixed-view setups, the egocentric configuration introduces rich scene motion that generates abundant background events, which are highly informative for event-based perception tasks [19, 39, 50].

Following established paradigms in RGB data [11, 30, 47], we focus on cooking scenarios to capture natural human activities, as kitchens inherently involve diverse human-object interactions, rapid motion dynamics, and complex visual patterns. The dataset is collected in 13 diverse kitchen environments with 10 participants, resulting in 5.5 hours of stereo event recordings. We ask each participant to wear the data collection device and naturally perform several common cooking activities in the kitchen without any instructions, which ensures the activities are unscripted. With stereo cameras and ground-truth depth maps, EventKitchen enables accurate 3D perception of human motions in indoor environments. The wearable device setup of EventKitchen also supports research in

Table 1: Comparison of our EventKitchen and event-based human activity datasets in terms of real-world capture, stereo event cameras, egocentric setting, event camera resolution, unscripted recording (UR), supporting multiple tasks (MT), object of activity, number of activity classes (#Cls), number of object bounding boxes (#BB), number of depth maps (#DM), and duration per recording (DR). EventKitchen is the only dataset recorded with stereo event cameras in the egocentric, real-world setting, and supports multiple tasks. EventKitchen compares favorably to existing event camera datasets on resolution and number of annotations. The average recording duration of three minutes encourages natural and unscripted behavior during data collection in EventKitchen. Note that we only report the dataset collected with real event cameras.

Dataset	Year	Real-world	Stereo	Ego	Resolution	UR	MT	Activity		#BB	#DM	DR
								Object	#Cls			
ASLAN-DVS [3]	2019	✗	✗	✗	240×180	✗	✗	Action	432	✗	✗	-
UCF-DVS [3, 59]	2019	✗	✗	✗	240×180	✗	✗	Action	101	✗	✗	25s
HMDB-DVS [3, 31]	2019	✗	✗	✗	240×180	✗	✗	Action	51	✗	✗	19s
UCF-Crime-DVS [51]	2025	✗	✗	✗	1280×720	✓	✗	Anomaly	14	✗	✗	avg 242s
DVS-Gesture [2]	2017	✓	✗	✗	128×128	✗	✗	Action	11	✗	✗	6s
ASL-DVS [3]	2019	✓	✗	✗	240×180	✗	✗	Hand	24	✗	✗	0.1s
PAF [38]	2019	✓	✗	✗	246×260	✗	✗	Action	10	✗	✗	5s
DailyAction [36]	2021	✓	✗	✗	346×260	✗	✗	Action	12	✗	✗	5s
Bullying10K [15]	2023	✓	✗	✗	346×260	✗	✗	Action	10	✗	✗	2-20s
HARDVS [68]	2023	✓	✗	✗	346×260	✗	✗	Action	300	✗	✗	5s
THU ^{MV-EACT} - 50 [20]	2024	✓	✗	✗	1280×800	✗	✗	Action	50	✗	✗	2-5s
DailyDVS-200 [67]	2024	✓	✗	✗	320×240	✗	✗	Action	200	✗	✗	1-20s
SeAct [70]	2024	✓	✗	✗	346×260	✗	✗	Action	58	✗	✗	-
EE3D-R [39]	2024	✓	✗	✓	640×480	✗	✗	Pose	10	✗	✗	-
EventKitchen (ours)	2026	✓	✓	✓	1280×720	✓	✓	Action	268	13,482	297,457	avg 180s

Augmented Reality, Virtual Reality, Human-Computer Interaction, and wearable AI [49, 55].

To demonstrate the potential and complexity of our dataset, we train preliminary baseline models on event-based action recognition, object detection, and stereo depth estimation. Our main contributions are:

- We collect the first large-scale stereo event camera benchmark dataset in a real-world daily life and egocentric setting for multiple tasks.
- We provide annotations of human action labels, action segments, object labels, and object bounding boxes.
- We supply the ground truth depth maps and calibration matrices among sensors.
- We present challenges on our EventKitchen dataset by evaluating seven baseline models.

2 Related Work

We compare EventKitchen with representative human activity datasets captured using real event cameras in Table 1. Existing event camera datasets are generally acquired either by recording monitors that display RGB data [3, 29, 32, 44, 51], or by directly capturing real-world scenes with event sensors [2, 3, 20, 36, 38, 67, 68, 70]. However, most of them are designed for a single task. In contrast, our

EventKitchen can be used to benchmark multiple tasks involving natural human activities within the real-world, stereo, and egocentric setting.

Synthetic datasets. Due to the scarcity and cost of event cameras, the scale of event camera datasets remains far smaller than the RGB datasets. To address this, a common approach is to record RGB data displayed on monitors using an event camera, as adopted in event-based object recognition datasets N-ImageNet [29], N-MNIST [44], N-Caltech101 [44], event-based action recognition datasets ASLAN-DVS [3], UCF-DVS [3], HMDB-DVS [3], and event-based anomaly detection dataset UCF-Crime-DVS [51]. However, it remains limited by the refresh rate and brightness range of the display setup, thus failing to fully reflect real-world dynamics.

Real-world datasets. THU^{MV-EACT}-50 [20] collects human actions with six event cameras at a high resolution of 1280×800 . Unfortunately, it is script-driven with 50 predefined actions. DailyDVS-200 [67] and HARDVS [68] provide a wide range of human actions of 200 and 300 classes each, captured with low resolutions of 320×240 and 346×260 , respectively, in diverse environments and illumination conditions. Still, their action classes are scripted. SeAct [70], DailyAction [36], DVS-Gesture [2], ASL-DVS [3], Bullying10K [15], and PAF [38] focus on action recognition, but with limited predefined action classes and low-resolution samples. EE3D-R [39] targets 3D human pose estimation with a wearable helmet, yet it only includes 10 scripted motions captured by a single event camera.

Complexity limitation. Currently, popular event camera datasets [2, 3, 38, 44, 58] are limited in complexity. As reported in [29], methods developed prior to 2021 already achieved over 90% object recognition accuracy on [2, 3, 32, 44, 58], suggesting that these datasets no longer pose sufficient challenges for modern models. This highlights the need for new event camera datasets that better reflect the complexity and variability of real-world environments.

To bridge the gap, our EventKitchen dataset provides high-resolution, multi-task, real-world data of natural human activities in a multi-modality, stereo, and egocentric setup.

3 The EventKitchen Dataset

We build a data collection device with multiple sensors, capture natural human actions with multiple participants, and annotate data with human annotators.

3.1 The Data Collection Device

Our data collection device is a wearable system mounted on a bicycle helmet, integrating multiple sensors as shown in Fig. 1(a). Two HD Prophesee Gen4 event cameras are mounted on a 3D-printed structure to form a stereo pair. We mount two CMOS RGB cameras next to each event camera to form the RGB-EVENT unit. To acquire reliable ground truth depth measurements, an Intel RealSense depth camera is positioned above the stereo configuration, adding another RGB




Fig. 2: Top: Rectified and annotated frames from four different time points $t_0, \dots, t_3$. Frames depict the ground truth depth maps and corresponding events aggregated from the left and right event cameras, along with human-annotated object bounding boxes. RGB references highlight the difficulty of annotating on events directly; **Middle:** Video streams from the depth camera, and left and right event cameras. Event colors (red, blue) show the event polarity; **Bottom:** Human-annotated action segments.

camera (D-RGB), depth maps, and a 6-axis IMU. (See the supplement for further device details.)

In summary, EventKitchen incorporates diverse multimodal data: two event streams, three RGB streams (two CMOS RGB and D-RGB), one depth stream, and one IMU stream.

Synchronization. We connect all sensors to a data-recording laptop using USB-3 ports. We employ the Robot Operating System (ROS) [52] as the framework for managing and operating all sensors in our data collection system. All sensors are synchronized with the timestamps provided in ROS.

Calibration. Due to the inherent asynchrony of event cameras, conventional calibration techniques based on image-based corner detection are not directly applicable. We adopt E2Calib [42] to reconstruct event streams into grayscale frames and apply the standard OpenCV [4] calibration framework to perform intrinsic calibration for all cameras. For extrinsic calibration, we designate the depth camera as the reference and perform stereo calibration with the remaining cameras, respectively. To facilitate event-based stereo depth estimation, the left event camera is used as the reference for extrinsic calibration with both the right event camera and the depth camera. After calibration, both intrinsic and extrinsic re-projection errors are below two pixels for all cameras. Detailed calibration results are given in the supplement.




Fig. 3: Bounding box annotation pipeline: 1. Obtain human-annotated bounding boxes on the RGB frame; 2. Construct 3D points of each bounding box based on its synchronized depth map; 3. Project 3D points to the left and right event camera using intrinsic and extrinsic matrices. **Depth map rectification pipeline:** 1. Construct 3D points for the depth map; 2. Project all 3D points to the target camera domain (left or right event camera) to get the projected depth map; 3. Rectify the projected depth with left and right event cameras.

3.2 Data collection

Variability and task complexity are important assets of a benchmark dataset. EventKitchen includes recordings with 10 participants of different nationalities, genders, and ages in eight private kitchens and five public kitchens. The participant statistics are reported in the supplement. As the data collection system integrates multiple sensors and is relatively heavy, we carefully considered participant comfort and safety during recording, in accordance with the HERC approval. Accordingly, we selected 13 cooking activities with moderate duration, from which participants could freely choose, which include 1. *Cut bread*, 2. *Cut cake*, 3. *Fry bacon*, 4. *Fry egg*, 5. *Fry pepper*, 6. *Make cereal bowl*, 7. *Make coffee*, 8. *Make fruit salad*, 9. *Make lemon water*, 10. *Make a sandwich*, 11. *Make tea*, 12. *Make vegetable salad*, and 13. *Wash dish*. At the start of the recording session, each participant voluntarily chooses a set of activities based on their preference. The data is recorded by the participant wearing the data collection helmet and a backpack containing the data recording laptop. We provide no instructions on how they perform these activities. We thus ensure that the collected data is natural and complex enough to reflect daily life activities in a real-world and unscripted setting.

Before the recording, participants are instructed to remove any identifiable items to ensure privacy protection. We perform sensor calibration before and after data collection.




Fig. 4: Left: Distribution of the 32 annotated action verbs in EventKitchen; **Right:** Distribution of the 268 annotated actions in EventKitchen. The two distributions demonstrate the variance of action verbs and action classes in EventKitchen, and the realistic long tails present further challenges, including few-shot learning and generalization.

3.3 Annotation and Ground Truth

We provide human annotations of action segments and bounding boxes with VGG Image Annotator (VIA) [16] tool on the D-RGB video, and ground-truth depth maps from the depth camera. We project the annotations onto event data using the calibration matrices. Human annotators do not label event data directly since accurately identifying objects and actions is too challenging, as shown in Fig. 2 (D-RGB shown in orange dashed boxes).

Action Segments To minimize annotator bias, we use the verb list of EPIC-KITCHENS [11] as reference and label actions as sets of $\{verb, noun\}$. Three different, randomly assigned human annotators label action segments in three steps: the first annotator defines the action labels in time order, and the second annotator annotates the corresponding action segment. The last annotator verifies the action segments from the second step. Annotators are permitted to modify labels from the previous step to ensure accuracy and consistency. We use the precise synchronization between the D-RGB and event cameras to align the start and end timestamps for each action.

Bounding Box Annotations Two human annotators are appointed to annotate the bounding boxes in each recording. For accuracy and consistency, one annotator is randomly assigned to annotate bounding boxes on the D-RGB frames, while the second annotator reviews and verifies the annotations. Given the large number and diversity of objects in kitchen environments, we focus primarily on kitchenware categories, as these objects are typically rigid and frequently encountered during cooking activities. In contrast, food items are often deformable and suffer significant appearance changes during preparation and cooking, which can introduce ambiguity in annotation [12, 56]. Therefore, we restrict the object classes to the following twelve categories: 1. *Bowl*, 2. *Box*, 3. *Chopping board*, 4. *Cup*, 5. *Fork*, 6. *Knife*, 7. *Pan*, 8. *Plastic packaging*, 9. *Plate*, 10. *Spatula*, 11. *Spoon*, and 12. *Teapot*.




Fig. 5: Left: Distribution of the 12 object classes in terms of number of samples in EventKitchen; **Middle:** Distribution of the central coordinate of all bounding boxes; **Right:** Distribution of the size of all bounding boxes. These figures demonstrate our projection pipeline is not biased, as the left and right cameras have a similar number of bounding boxes per class. Moreover, the bounding boxes are meaningfully distributed across the frames with diverse variances in size.

Bounding box projection. As the D-RGB frames and depth maps are pixel-to-pixel aligned, we construct a 3D point for each pixel in the D-RGB frames (Fig. 3). Given a human-annotated bounding box on the D-RGB frame, we retrieve all 3D depth points inside the bounding box. Then, we project each set of 3D points to the left and right event cameras to get the projected bounding box. As refinement, we apply a smoothing process by averaging the y-axis values of all points along each width edge and the x-axis values of all points along each height edge. Bounding boxes that are out of the field of view (FoV) of the event cameras are removed.

Depth map Depth maps are captured using the Intel RealSense depth camera with an error of $< 2\%$ at two meters. To enable stereo depth estimation, the ground truth needs to be rectified with the stereo system. Due to the FoV difference between the depth and event cameras, we first project the raw depth map on to the FoV of the left event camera by retrieving 3D points for all pixels, then rectify the projected depth with the event stereo. We illustrate the pipeline along with a set of rectified data in Fig. 3.

3.4 Dataset Statistics

We record 14 sessions comprising 110 sequences in 13 unique kitchen environments, totaling 5.5 hours from 10 participants, as reported in Tab. 2. With multi-sensor data, EventKitchen provides 5.5 hours of stereo event recordings, 5.5 hours of stereo RGB videos from CMOS cameras at 30 fps, 5.5 hours of D-RGB videos from the depth camera at 15 fps, 5.5 hours of 16-bit depth videos at 15 fps, and 5.5 hours of 6-axis IMU data at 200 fps. The average sequence duration of approximately three minutes further indicates that our recording setup preserves natural and unscripted human motions.

Action Segments. We report the verb and action distributions in Fig. 4. In total, there are 268 action classes with 32 verb classes annotated in EventKitchen,




Fig. 6: **Left:** Depth histogram from all 10,800 samples is in the range from 200 mm to 1500 mm; **Right:** Mean and standard deviation of depth per session. Depth information is within an expected range for kitchen activities where objects are typically about 700 mm away.

Table 2: Statistics of the 14 recording sessions (S1–S14). The EventKitchen dataset provides 14 recording sessions with a total duration of 5.5 hours. Each session consists of multiple sequences, with each sequence corresponding to a cooking activity. The average duration of each sequence is approximately three minutes.

Session ID	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	S11	S12	S13	S14	All
# Sequence	7	9	8	4	7	8	10	9	5	8	5	12	10	8	110
Duration (s)	1978	2443	1333	537	1003	1444	1865	1746	491	1198	875	2092	1746	999	19750

associated with 10,762 segments. We find that our dataset and the large-scale, unscripted egocentric EPIC-KITCHENS dataset share common high-frequency verbs (Table 3). This implies that our recording and annotation pipeline preserves the naturality of human actions.

Bounding Boxes. Following the annotation pipeline described in Section 3.3, we annotate the D-RGB frames at a rate of 1 frame per 4 seconds, and project the labeled bounding boxes to the left and right event cameras. Eventually, there are 13,482 bounding boxes with 12 different classes in the event domain. We present the distribution of each object class, along with the size and central coordinates of all bounding boxes in Fig. 5.

Depth Map. Our dataset includes 297,547 raw ground truth depth maps. Fig. 6 shows the statistics of 10,800 raw depth maps sampled from each sequence.

We compare our EventKitchen with other event-based human activity datasets in Table 1. EventKitchen is the only event-based human activity dataset that is recorded naturally and unscripted, and supports multiple tasks.

4 Baseline Models

EventKitchen’s multi-modal stereo setup enables several challenging tasks in the event domain. We establish baselines on the following event-based tasks: 1) action recognition, 2) object detection, 3) stereo depth estimation (Fig. 2). For all baseline models, we sum the event representations from the left and right event

Table 3: Frequency ranking of common verbs in EventKitchen and EPIC-KITCHENS [11]. We count ‘*Rinse*’ (7th) and ‘*Soap*’ (16th) separately, while EPIC-KITCHENS includes the two verbs in ‘*Wash*’ (3rd). The similar frequency ranking between the two datasets indicates they are unscripted actions.

Dataset	Action Frequency Ranking						
	Put	Take	Open	Close	Move	Cut	Rinse/Wash
EventKitchen (ours)	1 st	2 nd	3 rd	4 th	5 th	6 th	7 th
EPIC-KITCHENS [11]	1 st	2 nd	4 th	5 th	10 th	6 th	3 rd

Table 4: The dataset split of EventKitchen. The split among both the data and annotations is roughly 0.8 : 0.2. #BB: number of bounding box, #AC: number of action class, #AS: number of action segment, #DM: number of depth map

Split	#Kitchen	#Session	#Sequence	Duration (s)	#BB	#AC	#AS	#DM
Train	9	10	82	15683	11065	252	8598	235271
Test	4	4	28	4067	2417	168	2164	62276

cameras and use the combined event input for both training and evaluation. See the supplement for implementation details.

We perform a kitchen-level split, allocating nine kitchens to the training set and reserving the remaining four kitchens for testing. This protocol ensures that the test environments are entirely unseen during training. We report the dataset split in Table 4, which results in a rough 0.8:0.2 split ratio for training and test sets in terms of duration (79.4% vs 20.6%), annotated action segments (78.4% vs 21.6%), annotated bounding boxes (82.1% vs 17.9%), and number of depth maps (79.1% vs 20.9%).

4.1 Action Recognition

Description. We define our action recognition objective as follows: considering a sequence of events $E_i^{t_0:t_n} = \{e_i^{t_0}, e_i^{t_1}, \dots, e_i^{t_n}\}$, the model should predict the corresponding action class $C_a = \{c_v + c_n; c_v \in C_v, c_n \in C_n\}$, where c_v is the verb class, and c_n is the noun class.

Methods. We implement two frame-based action recognition algorithms, Temporal Shift Module (TSM) [34] and Video Swin Transformer (Swin) [37], both of which achieve the top-2 performance on the large-scale event-based action recognition dataset HARDVS [68] and DailyDVS-200 [67], to evaluate their performance on EventKitchen. TSM is an efficient and lightweight method for video action recognition. Specifically, we use a ResNet-50 [26] as the backbone for TSM. Swin is a powerful and flexible method that computes self-attention globally with spatial-temporal factorization. We use the largest Swin-Base as the backbone for Swin. Both TSM and Swin models are pretrained on Kinetics 400 [6]. We finetune the models on the selected 69 action classes C_a^{69} , which have at least five samples in the test set. Additionally, we finetune the models to recognize the corresponding verbs of C_a^{69} , which comprise 18 independent classes C_v^{18} .

Table 5: Baseline results for action recognition. The low performances indicate the difficulty of EventKitchen.

Method	Actions: C_a^{69}		Verbs: C_v^{18}	
	Top-1	Top-5	Top-1	Top-5
TSM [34]	19.23	42.26	38.17	87.20
Swin [37]	24.69	56.48	46.08	90.67

Table 6: Test accuracy of the 10 most frequent action and verb classes in the training set. OT: open tap; CT: close tap; TK: take knife; PV: put vegetable; PK: put knife; TV: take vegetable; TP: take plastic packaging; OD: open drawer; OC: open cupboard; CD: close drawer.

Method	10 most frequent (in train set) actions										10 most frequent (in train set) verb									
	OT	CT	TK	PV	PK	TV	TP	OD	OC	CD	Put	Take	Open	Close	Move	Rinse	Cut	Fry	Pour	Throw
TSM [34]	20.0	69.6	24.3	17.6	23.4	0.0	32.7	25.0	34.8	15.0	54.8	42.3	42.2	37.7	5.0	20.6	44.6	13.9	4.5	21.4
Swin [37]	50.0	34.8	28.4	23.5	31.3	0.0	32.7	43.8	21.7	5.0	66.4	41.3	48.5	22.7	11.7	8.8	87.8	16.7	54.5	0.0

Criteria. We use the top-1 and top-5 accuracy to evaluate the action class and verb class recognition. The accuracy is reported on the test sets of C_a^{69} and C_v^{18} .

Results. The results are presented in Table 5. Both methods achieve performance clearly above random chance, with Swin consistently outperforming TSM. Verb classification proves to be an easier task than action classification. However, neither method achieves high performance on either task. We further provide the per-class metrics of the 10 most frequent classes in C_a^{69} and C_v^{18} in Table 6. Both results strongly highlight the challenging nature of our EventKitchen.

4.2 Object Detection

Description. We define the object detection challenge as follows: within a time window of Δt , given a sequence of events $E_i^{t_0:t_0+\Delta t} = \{e_i^{t_0}, e_i^{t_1}, \dots, e_i^{t_0+\Delta t}\}$, we aim to localize and classify objects of interest into 12 classes.

Methods. We evaluate the state-of-the-art frame-based object detection algorithm YOLOv10 [65], event-based object detection algorithm Recurrent Vision Transformers (RVT) [24] and EvRT-DETR [63] on EventKitchen. We implement the largest YOLOv10-x model pre-trained on MS-COCO [35] and the RVT-Base model pretrained on 1 Mpx [46]. We follow the EvRT-DETR pipeline: we first train the RT-DETR-B frame detector initialized from a 1 Mpx pretrained model; we then train the EvRT-DETR-B model by initializing it from the trained RT-DETR-B model.

Criteria. For evaluation we report the average precision AP [35] with IoU threshold between 0.5 and 0.95 with steps of 0.05, AP_{50} with IoU threshold at 0.5, and AP_{05} with IoU threshold at 0.05.

Results. Table 7 shows the baseline results. YOLOv10 achieves an AP of 16.2% on EventKitchen, while it achieves an AP of 54.4% on the COCO dataset. RVT achieves an AP of 7.6% on EventKitchen, while it achieves an AP of 47.4% on the 1 Mpx dataset. EvRT-DETR achieves an AP of 8.5% on EventKitchen, while

Table 7: Baseline results for object detection. The column ‘All’ represents the mean average precision (mAP) over the 12 object classes. CB: chopping board; PP: Plastic packaging.

Method	Metric	Bowl	Box	CB	Cup	Fork	Knife	Pan	PP	Plate	Spatula	Spoon	Teapot	All
YOLOv10 [65]	AP	26.0	6.3	32.2	18.1	1.1	8.2	36.5	10.3	26.8	14.9	1.0	13.4	16.2
	AP_{50}	44.6	13.3	56.1	37.4	1.9	19.3	47.0	19.4	52.5	38.2	5.2	23.4	29.9
	AP_{05}	49.3	18.3	66.4	54.5	2.2	42.1	49.4	26.2	55.5	51.0	16.0	25.9	38.1
RVT [24]	AP	18.3	1.8	14.5	8.3	0	0.2	25.0	3.8	15.9	0.4	0.1	3.3	7.6
	AP_{50}	30.6	4.7	27.1	25.2	0	0.6	49.8	13.1	36.5	1.6	0.5	4.0	16.1
	AP_{05}	33.6	7.8	31.3	34.0	0	10.4	59.6	20.9	43.7	6.9	17.6	6.1	22.7
EvRT-DETR [63]	AP	20.5	1.0	20.4	6.4	0	3.2	27.5	6.0	14.2	1.5	0.2	0.5	8.5
	AP_{50}	37.7	2.0	41.0	18.1	0	8.0	37.9	11.8	35.0	6.1	1.0	1.4	16.7
	AP_{05}	40.2	6.9	51.1	27.5	0.2	25.3	40.0	18.0	37.4	16.9	6.5	2.5	22.7

it achieves an AP of 50.1% on the 1 Mpx dataset. Even for AP_{05} , YOLOv10 reaches only 38.1%, and both RVT and EvRT-DETR reach only 22.7%. Although RVT and EvRT-DETR are designed for event-based object detection, they still perform below YOLOv10 on EventKitchen due to their specialization for high-frequency (non-human) annotations of 1Mpx automotive datasets. And we observe that the detection rates of both baselines are especially low for the classes ‘fork’ and ‘spoon’, as the two classes are smaller in size, often occluded by hands and food, and involve fast actions such as ‘stir’. These findings demonstrate the object detection complexity in our dataset

4.3 Stereo Depth Estimation

Description. We define the stereo depth estimation challenge as follows: given event sequences from the left event camera $E_l^{t_0:t_0+\Delta t} = \{e_l^t\}_{t=t_0}^{t_0+\Delta t}$ and the right event camera $E_r^{t_0:t_0+\Delta t} = \{e_r^t\}_{t=t_0}^{t_0+\Delta t}$ that within a time window Δt , we estimate the depth map $D_l^{t_0:t_0+\Delta t}$ on the left event camera domain. Note that it can also be used to estimate the depth map of the right event camera.

Methods. We train a state-of-the-art event-based stereo depth estimation algorithm SE-CFF [43] from scratch on EventKitchen. We evaluate the model in ‘Event-only’ setting without using the RGB images. For event representation, the method uses Stacking by Number (SBN) [41, 43, 66]. Additionally, we test an RGB stereo foundation model for zero-shot stereo matching [69], where we first reconstruct left and right events to grayscale images using E2VID [54], then feed the reconstructed images to the foundation model. For both methods, we only train and test in the depth range from 200 mm to 1500 mm to filter invalid depth values.

Criteria. We report the Root Mean Square Error (RMSE) and Mean Absolute Error (MAE) of the two methods, averaged over depth images. As a reference for estimation performance, we also include the average standard deviation (STD) and mean absolute deviation (MAD) of depth values in the ground truth. All metrics are measured in millimeters.

Table 8: Depth estimation performance (in mm). Lower RMSE and MAE indicate better accuracy. Errors smaller than STD and MAD in SE-CFF suggest the algorithm is effective, which is not the case for the direct utilization of FoundationStereo model. $aM@bHz$ denotes a maximum stack size of a million events at a sampling rate of b Hz.

Method	Training Setting	RMSE ↓	MAE ↓
SE-CFF [43]	5M@3Hz	88.19	59.32
	5M@1Hz	88.21	57.38
	15M@1Hz	84.91	54.84
FoundationStereo [69]	—	155.06	123.85
Ground Truth	—	STD: 106.26 MAD: 82.00	

Results. Table 8 summarizes results for both methods. For SE-CFF, we report performance using the default configuration with a stack size of 5 million events at two sampling rates (1 Hz and 3 Hz), and a 15 million event stack size at 1 Hz. We find that increasing the number of events has a larger impact on error than increasing the sampling rate, which we attribute to the high spatial resolution of EventKitchen. Comparing the test estimation errors with the ground truth deviations (STD and MAD), the depth estimation algorithm is effective. By contrast, the higher error rates of the RGB FoundationStereo model highlight the need for adaptation strategies to better convert event data into frame-based inputs for foundation models. In both methods, the magnitude of the errors on our EventKitchen indicates the overall difficulty of the dataset.

4.4 Discussion

Our baseline models across three distinct tasks highlight the novel challenges and complexity of EventKitchen. Given that EventKitchen is naturally collected in 13 diverse kitchens, the data inhabits complex, variable, and dynamic background and foreground information. Thus, different scenes may map to the same action class: turning on the electrical vs. gas hobs, opening the single-lever vs. rotated handle taps, opening embedded vs. stand-alone fridges, and so on. Different participants have different cooking styles, which bring various human actions to EventKitchen. Take the ‘fry egg’ activity as an example: 2 participants use chopsticks to fry eggs, 5 participants use a spatula to fry eggs, and 1 participant does not use a tool to fry eggs. These challenges demonstrate that EventKitchen is more complex than fixed viewpoint, studio environments, and scripted action setups, encouraging the community to develop more robust, generalizable solutions.

Limitations. As a naturally collected real-world dataset, we observe the long-tail distributions in both object classes and action classes of EventKitchen in Fig. 4. Class imbalance remains a significant challenge for baseline models in object detection and action recognition. EventKitchen encourages future research to address this limitation. Besides, although EventKitchen provides valuable human-annotated labels, the quantity remains limited compared to the 1Mpx datasets with 25 million non-human annotations. Expanding the volume of annotations presents an important direction for future development.

5 Conclusion

In this work, we present a large-scale, stereo event camera dataset, **Event-Kitchen**, to support human activity-related research for the event community. Unlike existing event-based datasets, EventKitchen captures natural human actions across 13 diverse kitchen environments, offering 14 sessions of rich multi-sensor recordings from 10 participants. It provides 5.5 hours of stereo event recordings with synchronized RGB, depth, and IMU data, along with detailed human annotations of 10,762 action segments and 13,482 bounding boxes. Our baseline results demonstrate that EventKitchen offers complex new challenges for event-based 1) action recognition, with natural actions, dynamic backgrounds, and varying environments, 2) object detection, from an egocentric perspective and small, fast-moving objects, and 3) stereo depth estimation in an indoor, near-field setting. By capturing natural, real-world human activities, EventKitchen establishes a challenging benchmark for neuromorphic vision beyond autonomous driving. Overall, our results showcase the potential of EventKitchen in advancing research on daily human activities using event cameras.

Acknowledgements

We thank Ruud de Jong from INSY, TU Delft, for technical support with the hardware, and Qingru Li from ImPhys, TU Delft, for assistance with the optical lenses and filters.

References

1. Aliminati, M.R., Chakravarthi, B., Verma, A.A., Vaghela, A., Wei, H., Zhou, X., Yang, Y.: Sevd: Synthetic event-based vision dataset for ego and fixed traffic perception. arXiv preprint arXiv:2404.10540 (2024)
2. Amir, A., Taba, B., Berg, D., Melano, T., McKinstry, J., Di Nolfo, C., Nayak, T., Andreopoulos, A., Garreau, G., Mendoza, M., et al.: A low power, fully event-based gesture recognition system. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7243–7252 (2017)
3. Bi, Y., Chadha, A., Abbas, A., Bourtsoulatze, E., Andreopoulos, Y.: Graph-based spatio-temporal feature learning for neuromorphic vision sensing. *IEEE Transactions on Image Processing* **29**, 9084–9098 (2020)
4. Bradski, G.: The opencv library. *Dr. Dobb’s Journal: Software Tools for the Professional Programmer* **25**(11), 120–123 (2000)
5. Calabrese, E., Taverni, G., Awai Easthope, C., Skriabine, S., Corradi, F., Longinotti, L., Eng, K., Delbruck, T.: Dhp19: Dynamic vision sensor 3d human pose dataset. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (June 2019)
6. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6299–6308 (2017)

7. Chanda, K., Verma, A., Vaghela, A., Yang, Y., Chakravarthi, B.: Event quality score (eqs): Assessing the realism of simulated event camera streams via distance in latent space. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 5105–5113 (2025)
8. Chaney, K., Cladera, F., Wang, Z., Bisulco, A., Hsieh, M.A., Korpela, C., Kumar, V., Taylor, C.J., Daniilidis, K.: M3ed: Multi-robot, multi-sensor, multi-environment event dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4016–4023 (2023)
9. Chen, J., Yang, Y., Deng, S., Teng, D., Pan, L.: Spikmamba: When snn meets mamba in event-based human action recognition. In: Proceedings of the 6th ACM International Conference on Multimedia in Asia. pp. 1–8 (2024)
10. Cohen, G., Marcireau, A.: Land: A longitudinal analysis of neuromorphic datasets. arXiv preprint arXiv:2602.15973 (2026)
11. Damen, D., Doughty, H., Farinella, G.M., Fidler, S., Furnari, A., Kazakos, E., Moltisanti, D., Munro, J., Perrett, T., Price, W., et al.: Scaling egocentric vision: The epic-kitchens dataset. In: Proceedings of the European conference on computer vision (ECCV). pp. 720–736 (2018)
12. Darkhalil, A., Shan, D., Zhu, B., Ma, J., Kar, A., Higgins, R., Fidler, S., Fouhey, D., Damen, D.: Epic-kitchens visor benchmark: Video segmentations and object relations. *Advances in Neural Information Processing Systems* **35**, 13745–13758 (2022)
13. De Tournemire, P., Nitti, D., Perot, E., Migliore, D., Sironi, A.: A large scale event-based detection dataset for automotive. arXiv preprint arXiv:2001.08499 (2020)
14. Delmerico, J., Cieslewski, T., Rebecq, H., Faessler, M., Scaramuzza, D.: Are we ready for autonomous drone racing? the uzh-fpv drone racing dataset. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 6713–6719. IEEE (2019)
15. Dong, Y., Li, Y., Zhao, D., Shen, G., Zeng, Y.: Bullying10k: A large-scale neuromorphic dataset towards privacy-preserving bullying recognition. *Advances in Neural Information Processing Systems* **36**, 1923–1937 (2023)
16. Dutta, A., Zisserman, A.: The via annotation software for images, audio and video. In: Proceedings of the 27th ACM international conference on multimedia. pp. 2276–2279 (2019)
17. Falanga, D., Kim, S., Scaramuzza, D.: How fast is too fast? the role of perception latency in high-speed sense and avoid. *IEEE Robotics and Automation Letters* **4**(2), 1884–1891 (2019)
18. Falanga, D., Kleber, K., Scaramuzza, D.: Dynamic obstacle avoidance for quadrotors with event cameras. *Science Robotics* **5**(40), eaaz9712 (2020)
19. Gallego, G., Delbrück, T., Orchard, G., Bartolozzi, C., Taba, B., Censi, A., Leutenegger, S., Davison, A.J., Conradt, J., Daniilidis, K., et al.: Event-based vision: A survey. *IEEE transactions on pattern analysis and machine intelligence* **44**(1), 154–180 (2020)
20. Gao, Y., Lu, J., Li, S., Li, Y., Du, S.: Hypergraph-based multi-view action recognition using event cameras. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **46**(10), 6610–6622 (2024). <https://doi.org/10.1109/TPAMI.2024.3382117>
21. Gehrig, D., Gehrig, M., Hidalgo-Carrió, J., Scaramuzza, D.: Video to events: Recycling video datasets for event cameras. In: IEEE Conf. Comput. Vis. Pattern Recog. (CVPR) (June 2020)
22. Gehrig, D., Scaramuzza, D.: Low-latency automotive vision with event cameras. *Nature* **629**(8014), 1034–1040 (2024)

23. Gehrig, M., Aarents, W., Gehrig, D., Scaramuzza, D.: Dsec: A stereo event camera dataset for driving scenarios. *IEEE Robotics and Automation Letters* **6**(3), 4947–4954 (2021)
24. Gehrig, M., Scaramuzza, D.: Recurrent vision transformers for object detection with event cameras. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 13884–13893 (2023)
25. Goyal, G., Di Pietro, F., Carissimi, N., Glover, A., Bartolozzi, C.: Moveenet: Online high-frequency human pose estimation with an event camera. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4024–4033 (2023)
26. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
27. Hu, Y., Binas, J., Neil, D., Liu, S.C., Delbruck, T.: Ddd20 end-to-end event camera driving dataset: Fusing frames and events with deep learning for improved steering prediction. In: *2020 IEEE 23rd international conference on intelligent transportation systems (ITSC)*. pp. 1–6. IEEE (2020)
28. Hu, Y., Liu, S.C., Delbruck, T.: v2e: From video frames to realistic dvs events. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 1312–1321 (2021)
29. Kim, J., Bae, J., Park, G., Zhang, D., Kim, Y.M.: N-imagenet: Towards robust, fine-grained object recognition with event cameras. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 2146–2156 (2021)
30. Kuehne, H., Arslan, A., Serre, T.: The language of actions: Recovering the syntax and semantics of goal-directed human activities. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 780–787 (2014)
31. Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., Serre, T.: Hmdb: a large video database for human motion recognition. In: *2011 International conference on computer vision*. pp. 2556–2563. IEEE (2011)
32. Li, H., Liu, H., Ji, X., Li, G., Shi, L.: Cifar10-dvs: an event-stream dataset for object classification. *Frontiers in neuroscience* **11**, 309 (2017)
33. Li, W., Wan, P., Wang, P., Li, J., Zhou, Y., Liu, P.: Benerf: neural radiance fields from a single blurry image and event stream. In: *European Conference on Computer Vision*. pp. 416–434. Springer (2024)
34. Lin, J., Gan, C., Han, S.: Tsm: Temporal shift module for efficient video understanding. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 7083–7093 (2019)
35. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *Computer vision—ECCV 2014: 13th European conference, zurich, Switzerland, September 6–12, 2014, proceedings, part v 13*. pp. 740–755. Springer (2014)
36. Liu, Q., Xing, D., Tang, H., Ma, D., Pan, G.: Event-based action recognition using motion information and spiking neural networks. In: *IJCAI*. pp. 1743–1749 (2021)
37. Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H.: Video swin transformer. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 3202–3211 (2022)
38. Miao, S., Chen, G., Ning, X., Zi, Y., Ren, K., Bing, Z., Knoll, A.: Neuromorphic vision datasets for pedestrian detection, action recognition, and fall detection. *Frontiers in neurorobotics* **13**, 38 (2019)

39. Millerdurai, C., Akada, H., Wang, J., Luvizon, D., Theobalt, C., Golyanik, V.: Eventego3d: 3d human motion capture from egocentric event streams. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1186–1195 (2024)
40. Millerdurai, C., Luvizon, D., Rudnev, V., Jonas, A., Wang, J., Theobalt, C., Golyanik, V.: 3d pose estimation of two interacting hands from a monocular event camera. In: 2024 International Conference on 3D Vision (3DV). pp. 291–301. IEEE (2024)
41. Mostafavi, M., Nam, Y., Choi, J., Yoon, K.J.: E2sri: Learning to super-resolve intensity images from events. *IEEE transactions on pattern analysis and machine intelligence* **44**(10), 6890–6909 (2021)
42. Muglikar, M., Gehrig, M., Gehrig, D., Scaramuzza, D.: How to calibrate your event camera. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1403–1409 (2021)
43. Nam, Y., Mostafavi, M., Yoon, K.J., Choi, J.: Stereo depth from events cameras: Concentrate and focus on the future. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6114–6123 (2022)
44. Orchard, G., Jayawant, A., Cohen, G.K., Thakor, N.: Converting static image datasets to spiking neuromorphic datasets using saccades. *Frontiers in neuroscience* **9**, 437 (2015)
45. Paredes-Vallés, F., Hagenaaars, J.J., Dupeyroux, J., Stroobants, S., Xu, Y., de Croon, G.C.: Fully neuromorphic vision and control for autonomous drone flight. *Science Robotics* **9**(90), eadi0591 (2024)
46. Perot, E., De Tournemire, P., Nitti, D., Masci, J., Sironi, A.: Learning to detect objects with a 1 megapixel event camera. *Advances in Neural Information Processing Systems* **33**, 16639–16652 (2020)
47. Perrett, T., Darkhalil, A., Sinha, S., Emara, O., Pollard, S., Parida, K.K., Liu, K., Gatti, P., Bansal, S., Flanagan, K., et al.: Hd-epic: A highly-detailed egocentric video dataset. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 23901–23913 (2025)
48. Planamente, M., Plizzari, C., Cannici, M., Ciccone, M., Strada, F., Bottino, A., Matteucci, M., Caputo, B.: Da4event: towards bridging the sim-to-real gap for event cameras using domain adaptation. *IEEE Robotics and Automation Letters* **6**(4), 6616–6623 (2021)
49. Plizzari, C., Goletto, G., Furnari, A., Bansal, S., Ragusa, F., Farinella, G.M., Damen, D., Tommasi, T.: An outlook into the future of egocentric vision. *International Journal of Computer Vision* **132**(11), 4880–4936 (2024)
50. Plizzari, C., Planamente, M., Goletto, G., Cannici, M., Gusso, E., Matteucci, M., Caputo, B.: E2 (go) motion: Motion augmented event stream for egocentric action recognition. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 19935–19947 (2022)
51. Qian, Y., Ye, S., Wang, C., Cai, X., Qian, J., Wu, J.: Ucf-crime-dvs: A novel event-based dataset for video anomaly detection with spiking neural networks. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39, pp. 6577–6585 (2025)
52. Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., Wheeler, R., Ng, A.Y., et al.: Ros: an open-source robot operating system. In: ICRA workshop on open source software. vol. 3, p. 5. Kobe (2009)
53. Rebecq, H., Gehrig, D., Scaramuzza, D.: Esim: an open event camera simulator. In: Conference on robot learning. pp. 969–982. PMLR (2018)

54. Rebecq, H., Ranftl, R., Koltun, V., Scaramuzza, D.: Events-to-video: Bringing modern computer vision to event cameras. *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)* (2019)
55. Rodin, I., Furnari, A., Mavroeidis, D., Farinella, G.M.: Predicting the future from first person (egocentric) vision: A survey. *Computer Vision and Image Understanding* **211**, 103252 (2021)
56. Saini, N., Wang, H., Swaminathan, A., Jayasundara, V., He, B., Gupta, K., Shrivastava, A.: Chop & learn: Recognizing and generating object-state compositions. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 20247–20258 (2023)
57. Shariff, W., Dilmaghani, M.S., Kieley, P., Moustafa, M., Lemley, J., Corcoran, P.: Event cameras in automotive sensing: A review. *IEEE Access* **12**, 51275–51306 (2024)
58. Sironi, A., Brambilla, M., Bourdis, N., Lagorce, X., Benosman, R.: Hats: Histograms of averaged time surfaces for robust event-based object classification. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1731–1740 (2018)
59. Soomro, K., Zamir, A.R., Shah, M.: Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402* (2012)
60. Sun, L., Sakaridis, C., Liang, J., Jiang, Q., Yang, K., Sun, P., Ye, Y., Wang, K., Gool, L.V.: Event-based fusion for motion deblurring with cross-modal attention. In: *European conference on computer vision*. pp. 412–428. Springer (2022)
61. Tan, G., Wang, Y., Han, H., Cao, Y., Wu, F., Zha, Z.J.: Multi-grained spatio-temporal features perceived network for event-based lip-reading. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 20094–20103 (2022)
62. Tan, K., Chakravarthi, B., et al.: How real is carlas dynamic vision sensor? a study on the sim-to-real gap in traffic object detection. *arXiv preprint arXiv:2506.13722* (2025)
63. Torbunov, D., Ren, Y., Ghose, A., Dim, O., Cui, Y.: Evrt-detr: Latent space adaptation of image detectors for event-based vision. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 9812–9821 (2025)
64. Tulyakov, S., Boicchio, A., Gehrig, D., Georgoulis, S., Li, Y., Scaramuzza, D.: Time lens++: Event-based frame interpolation with parametric non-linear flow and multi-scale fusion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 17755–17764 (2022)
65. Wang, A., Chen, H., Liu, L., Chen, K., Lin, Z., Han, J., et al.: Yolov10: Real-time end-to-end object detection. *Advances in Neural Information Processing Systems* **37**, 107984–108011 (2025)
66. Wang, L., Ho, Y.S., Yoon, K.J., et al.: Event-based high dynamic range image and very high frame rate video generation using conditional generative adversarial networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10081–10090 (2019)
67. Wang, Q., Xu, Z., Lin, Y., Ye, J., Li, H., Zhu, G., Ali Shah, S.A., Bennamoun, M., Zhang, L.: Dailydvs-200: A comprehensive benchmark dataset for event-based action recognition. In: *European Conference on Computer Vision*. pp. 55–72. Springer (2024)
68. Wang, X., Wu, Z., Jiang, B., Bao, Z., Zhu, L., Li, G., Wang, Y., Tian, Y.: Hardvs: Revisiting human activity recognition with dynamic vision sensors. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 38, pp. 5615–5623 (2024)

69. Wen, B., Trepte, M., Aribido, J., Kautz, J., Gallo, O., Birchfield, S.: Foundation-stereo: Zero-shot stereo matching. CVPR (2025)
70. Zhou, J., Zheng, X., Lyu, Y., Wang, L.: Exact: Language-guided conceptual reasoning and uncertainty estimation for event-based action recognition and more. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18633–18643 (2024)
71. Zhu, A.Z., Thakur, D., Özaslan, T., Pfrommer, B., Kumar, V., Daniilidis, K.: The multivehicle stereo event camera dataset: An event camera dataset for 3d perception. IEEE Robotics and Automation Letters **3**(3), 2032–2039 (2018)