

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.

# The MECCANO Dataset: Understanding Human-Object Interactions from Egocentric Videos in an Industrial-like Domain

Francesco Ragusa IPLAB, University of Catania XGD-XENIA s.r.l., Acicastello, Catania, Italy francesco.ragusa@unict.it

> Salvatore Livatino University of Hertfordshire s.livatino@herts.ac.uk

# Abstract

Wearable cameras allow to collect images and videos of humans interacting with the world. While human-object interactions have been thoroughly investigated in third person vision, the problem has been understudied in egocentric settings and in industrial scenarios. To fill this gap, we introduce MECCANO, the first dataset of egocentric videos to study human-object interactions in industrial-like settings. MECCANO has been acquired by 20 participants who were asked to build a motorbike model, for which they had to interact with tiny objects and tools. The dataset has been explicitly labeled for the task of recognizing human-object interactions from an egocentric perspective. Specifically, each interaction has been labeled both temporally (with action segments) and spatially (with active object bounding boxes). With the proposed dataset, we investigate four different tasks including 1) action recognition, 2) active object detection, 3) active object recognition and 4) egocentric human-object interaction detection, which is a revisited version of the standard human-object interaction detection task. Baseline results show that the MECCANO dataset is a challenging benchmark to study egocentric humanobject interactions in industrial-like scenarios. We publicy release the dataset at https://iplab.dmi.unict. it/MECCANO/.

# 1. Introduction

Being able to analyze human behavior from egocentric observations has many potential applications related to the recent development of wearable devices [1, 2, 3] which range from improving the personal safety of workers in a factory [10] to providing assistance to the visitors of a muAntonino Furnari IPLAB, University of Catania furnari@dmi.unict.it

Giovanni Maria Farinella IPLAB, University of Catania gfarinella@dmi.unict.it



Figure 1. Toy model built by subjects interacting with 2 tools, 49 components and the instructions booklet. Better seen on screen.

seum [22, 49, 11]. In particular, with the rapid growth of interest in wearable devices in industrial scenarios, recognizing human-object interactions can be useful to prevent safety hazards, implement energy saving policies and issue notifications about actions that may be missed in a production pipeline [55].

In recent years, progress has been made in many research areas related to human behavior understanding, such as action recognition [23, 54, 25, 66, 34, 39], object detection [27, 26, 51, 50] and human-object interaction detection [28, 31, 52, 43]. These advances have been possible thanks to the availability of large-scale datasets [35, 40, 13, 37, 31, 8] which have been curated and often associated with dedicated challenges. In the egocentric vision domain, in particular, previous investigations have considered the contexts of kitchens [13, 37, 17], as well as daily living activity at home and in offices [45, 57, 16, 44]. While these contexts provide interesting test-beds to study user behavior in general, egocentric human-object



Figure 2. Examples of Human-Object Interactions in third person vision (first row) and first person vision (second row)<sup>1</sup>

interactions have not been previously studied in industrial environments such as factories, building sites, mechanical workshops, etc. This is mainly due to the fact that data acquisition in industrial domains is difficult because of privacy issues and the need to protect industrial secrets.

In this paper, we present MECCANO, the first dataset of egocentric videos to study human-object interactions in industrial-like settings. To overcome the limitations related to data collection in industry, we resort to an industrial-like scenario in which subjects are asked to build a toy model of a motorbike using different components and tools (see Figure 1). Similarly to an industrial scenario, the subjects interact with tools such as a screwdriver and a wrench, as well as with tiny objects such as screws and bolts while executing a task involving sequential actions (e.g., take wrench, tighten bolt, put wrench). Despite the fact that this scenario is a simplification of what can be found in real industrial settings, it is still fairly complex, as our experiments show. MECCANO was collected by 20 different participants in two countries (Italy and United Kingdom). We densely annotated the acquired videos with temporal labels to indicate the start and end times of each human-object interaction, and with bounding boxes around the active objects involved in the interactions. We hope that the proposed dataset will encourage research in this challenging domain. The dataset is publicly released at the following link: https://iplab.dmi.unict.it/MECCANO/.

We show that the proposed dataset can be used to study four fundamental tasks related to the understanding of human-object interactions: 1) Action Recognition, 2) Active Object Detection, 3) Active Object Recognition and 4) Egocentric Human-Object Interaction Detection. While past works have already investigated the tasks of action recognition [13, 37, 41, 56], active object detection [46], and active object recognition [13] in the context of egocentric vision, Human-Object Interaction (HOI) detection has been generally studied in the context of third person vision [31, 28, 47, 9, 67, 38, 63]. Since we believe that modelling actions both semantically and spatially is fundamental for egocentric vision applications, we instantiate the Human-Object Interaction detection task in the context of the proposed dataset.

HOI detection consists in detecting the occurrence of human-object interactions, localizing both the humans taking part in the action and the interacted objects. HOI detection also aims to understand the relationships between humans and objects, which is usually described with a verb. Possible examples of HOIs are "talk on the cell phone" or "hold a fresbee". HOI detection models mostly consider one single object involved in the interaction [31, 30, 28, 65, 9]. Hence, an interaction is generally defined as a triplet in the form *<human*, *verb*, *object>*, where the human is the subject of the action specified by a verb and an object. Sample images related to human-object interactions in a third-person scenario are shown in Figure 2-top. We define Egocentric Human-Object Interaction (EHOI) detection as the task of producing <verb, objects> pairs describing the interaction observed from the egocentric point of view. Note that in EHOI, the human interacting with the objects is always the camera wearer, while one or more objects can be involved simultaneously in the interaction. The goal of EHOI detection is to infer the verb and noun classes, and to localize each active object involved in the interaction. Figure 2-bottom reports some examples of Egocentric Human-Object Interactions.

We perform experiments with baseline approaches to tackle the four considered tasks. Results suggest that the proposed dataset is a challenging benchmark for understanding egocentric human-object interactions in industrial-like settings. In sum, the contributions of this work are as follows: 1) we present MECCANO, a new challenging egocentric dataset to study human-object interactions in an industrial-like domain; 2) we instantiate the HOI definition in the context of egocentric vision (EHOI); 3) we show that the current state-of-the-art approaches achieve limited performance, which suggests that the proposed dataset is an interesting benchmark for studying egocentric human-object interactions in industrial-like domains.

# 2. Related Work

**Datasets for Human Behavior Understanding** Previous works have proposed datasets to tackle the task of Human-Object Interaction (HOI) detection. Among the most notable datasets, we can mention V-COCO [31], which adds 26 verb labels to the 80 objects of COCO [40], HICO-Det [8], labeled with 117 verbs and 80 objects, HOI-A [38], which focuses on 10 verbs and 11 objects indicat-

<sup>&</sup>lt;sup>1</sup>Images in the first row were taken from the COCO dataset [40] while those in the second row belong to the MECCANO dataset.

Dataset	Settings	EGO?	Video?	Tasks	Year	Frames	Sequences	AVG. video duration	Action classes	Object classes	Object BBs	Participants
MECCANO	Industrial-like	~	~	EHOI, AR, AOD, AOR	2020	299,376	20	20.79 min	61	20	64,349	20
EPIC-KITCHENS [13]	Kitchens	$\checkmark$	~	AR, AOR	2018	11,5M	432	7.64 min	125	352	454,255	32
EGTEA Gaze+ [37]	Kitchens	$\checkmark$	$\checkmark$	AR	2018	2,4M	86	0.05 min	106	0	0	32
THU-READ [57]	Daily activties	$\checkmark$	$\checkmark$	AR	2017	343,626	1920	7.44 min	40	0	0	8
ADL [46]	Daily activities	$\checkmark$	$\checkmark$	AR, AOR	2012	1,0M	20	30.0 min	32	42	137,780	20
CMU [17]	Kitchens	$\checkmark$	$\checkmark$	AR	2009	200,000	16	15.0 min	31	0	0	16
Something-Something [29]	General	Х	$\checkmark$	AR, HOI	2017	5,2 M	108,499	0.07 min	174	N/A	318,572	N/A
Kinetics [6]	General	Х	$\checkmark$	AR	2017	N/A	455,000	0.17 min	700	0	0	N/A
ActivityNet [19]	Daily activites	Х	$\checkmark$	AR	2015	91,6 M	19,994	2.55 min	200	N/A	N/A	N/A
HOI-A [38]	General	Х	Х	HOI, AOR	2020	38,668	N/A	N/A	10	11	60,438	N/A
HICO-DET [8]	General	Х	Х	HOI, AOR	2018	47,776	N/A	N/A	117	80	256,672	N/A
V-COCO [31]	General	Х	Х	HOI, OD	2015	10,346	N/A	N/A	26	80	N/A	N/A

Table 1. Comparative overview of relevant datasets. HOI: HOI Detection. EHOI: EHOI Detection. AR: Action Recognition. AOD: Active Object Detection. AOR: Active Object Recognition. OD: Object Detection.

ing actions dangerous while driving. Other works have proposed datasets for action recognition from video. Among these, ActivityNet [19] is a large-scale dataset composed of videos depicting 203 activities that are relevant to how humans spend their time in their daily lives, Kinetics [33, 6] is a dataset containing 700 human action classes, Something-Something [29] includes low-level concepts to represent simple everyday aspects of the world. Previous works also proposed datasets of egocentric videos. Among these datasets, EPIC-Kitchens [13, 15, 14] focuses on unscripted activities in kitchens, EGTEA Gaze+ [37] contains videos paired with gaze information collected from participants cooking different recipes in a kitchen, CMU [17] is a multimodal dataset of egocentric videos including RGB, audio and motion capture information, ADL [46] contains egocentric videos of subjects performing daily activities, THU-READ [57] contains RGB-D videos of subjects performing daily-life actions in different scenarios. Table 1 compares the aforementioned datasets with respect to the proposed dataset. MECCANO is the first dataset of egocentric videos collected in an industrial-like domain and annotated to perform EHOI Detection. It is worth noting that previous egocentric datasets have considered scenarios related to kitchens, offices, and daily-life activities and that they have generally tackled the action recognition task rather than EHOI detection.

Action Recognition Action recognition from video has been thoroughly investigated especially in the third person vision domain. Classic works [36, 12] relied on motionbased features such as optical flow and space-time features. Early deep learning works fused processing of RGB and optical flow features with two-stream networks [54, 25, 62], 3D ConvNets are commonly used to encode both spatial and temporal dimensions [58, 59, 7], long-term filtering and pooling has focused on representing actions considering their full temporal extent [61, 25, 62, 66, 66, 39]. Other works separately factor convolutions into separate 2D spatial and 1D temporal filters [24, 60, 64, 48]. Among recent works, Slow-Fast networks [23] avoid using precomputed optical flow and encodes motion into a "fast" pathway (which operates at a high frame rate) and simultaneously a "slow" pathway which captures semantics (operating at a low frame rate). We asses the performance of state-of-the-art action recognition methods on the proposed dataset considering SlowFast networks [23], I3D [7] and 2D CNNs as baselines.

HOI Detection Previous works have investigated HOI detection mainly from a third person vision point of view. The authors of [31] proposed a method to detect people performing actions able to localize the objects involved in the interactions on still images. The authors of [28] proposed a human-centric approach based on a three-branch architecture (InteractNet) instantiated according to the definition of HOI in terms of a <human, verb, object> triplet. Some works [47, 9, 67] explored HOI detection using graph convolutional neural networks after detecting humans and objects in the scene. Recent works [38, 63] represented the relationship between both humans and objects as the intermediate point which connects the center of the human and object bounding boxes. The aforementioned works addressed the problem of HOI detection in the third person vision domain. In this work, we look at the task of HOI detection from an egocentric perspective considering the proposed MECCANO dataset.

EHOI Detection EHOI detection is understudied due to the limited availability of egocentric datasets labelled for this task. While some previous datasets such as EPIC-KITCHENS [13, 14] and ADL [46] have been labeled with bounding box annotations, these datasets have not been explicitly labeled for the EHOI detection task indicating relationships between labeled objects and actions, hence preventing the development of EHOI detection approaches. Some related studies have modeled the relations between entities for interaction recognition as object affordances [42, 43, 21]. Other works tackled tasks related to EHOI detection proposing hand-centric methods [5, 4, 52]. Despite these related works have considered human-object interaction from an egocentric point of view, the EHOI detection task has not yet been defined or studied systematically in past works. With this paper we aim at providing a definition of the task, a suitable benchmark dataset, as well as an initial evaluation of baseline approaches.



Figure 3. Examples of data acquired by the 20 different participants in two countries (Italy, United Kingdom).



Figure 4. Statistics of the 20 participants.

# 3. MECCANO Dataset

# 3.1. Data Collection

The MECCANO dataset has been acquired in an industrial-like scenario in which subjects built a toy model of a motorbike (see Figure 1). The motorbike is composed of 49 components with different shapes and sizes belonging to 19 classes. In our settings, the components *A054* and *A051* of Figure 1 have been grouped under the "screw" class, whereas *A053*, *A057* and *A077* have been grouped under the "washers" class. As a result, we have 16 component classes<sup>2</sup>. Note that multiple instances of each class are necessary to build the model. In addition, 2 tools, a *screwdriver* and a *wrench*, are available to facilitate the assembly of the toy model. The subjects can use the instruction booklet to understand how to build the toy model following the sequential instructions.

For the data collection, the 49 components related to the considered 16 classes, the 2 tools and the instruction booklet have been placed on a table to simulate an industrial-like environment. Objects of the same component class have been grouped and placed in a heap, and heaps have been placed randomly (see Figure 3). Other objects not related to the toy model were present in the scene (i.e., clutter background). We have considered two types of table: a light-colored table and a dark one. The dataset has been acquired by 20 different subjects in 2 countries (Italy and United Kingdom) be-



Figure 5. Long-tail distribution of verbs classes.

tween May 2019 and January 2020. Participants were from 8 different nationalities with ages between 18 and 55. Figure 4 reports some statistics about participants. We asked participants to sit and build the model of the motorbike. No other particular instruction was given to the participants, who were free to use all the objects placed in the table as well as the instruction booklet. Some examples of the captured data are reported in Figure 3.

Data was captured using an Intel RealSense SR300 device which has been mounted on the head of the participant with an headset. The headset was adjusted to control the point of view of the camera with respect to the different heights and postures of the participants, in order to have the hands located approximately in the middle of the scene to be acquired. Videos were recorded at a resolution of 1920x1080 pixels and with a framerate of 12*fps*. Each video corresponds to a complete assembly of the toy model starting from the 49 pieces placed on the table. The average duration of the captured videos is 21.14*min*, with the longest one being 35.45*min* and the shortest one being 9.23*min*.

### 3.2. Data Annotation

We annotated the MECCANO dataset in two stages. In the first stage, we temporally annotated the occurrences of all human-object interactions indicating their start and end times, as well as a verb describing the interaction. In the second stage, we annotated the *active objects* with bounding boxes for each temporal segment.

Stage 1: Temporal Annotations We considered 12 dif-

<sup>&</sup>lt;sup>2</sup>See supplementary material for more details.



Figure 6. Example of two overlapping temporal annotations along with the associated verbs.



Figure 7. Example of bounding box annotations for active objects (first row) and occluded active objects (second row).

ferent verbs which describe the actions performed by the participants: take, put, check, browse, plug, pull, align, screw, unscrew, tighten, loosen and fit. As shown in Figure 5, the distribution of verb classes of the labeled samples in our dataset follows a long-tail distribution, which suggests that the taxonomy captures the complexity of the considered scenario. Since a participant can perform multiple actions simultaneously, we allowed the annotated segments to overlap (see Figure 6). In particular, in the MECCANO dataset there are 1401 segments (15.82 %) which overlap with at least another segment. We consider the start time of a segment as the timestamp in which the hand touches an object, changing its state from *passive* to *active*. The only exception is for the verb *check*, in which case the user doesn't need to touch an object to perform an interaction. In this case, we annotated the start time when it is obvious from the video sequence that the user is looking at the object (see Figure 6). With this procedure, we annotated 8857 video segments.

**Stage 2:** Active Object Bounding Box Annotations We considered 20 object classes which include the 16 classes categorizing the 49 components, the two tools (*screwdriver* and *wrench*), the instructions booklet and a *partial\_model* class. The latter object class represents assembled components of the toy model which are not yet complete (e.g., a *screw* and a *bolt* fixed on a *bar* which have not yet been assembled with the rest of the model<sup>3</sup>). For each temporal segment, we annotated the *active* objects in frames sampled every 0.2 seconds. Each active object annotation consists in a (*class, x, y, w, h*) tuple, where *class* represents the class of the object and (x, y, w, h) defines a bounding box around the object. We annotated multiple objects when they were active simultaneously (see Figure 7 - first row). Moreover, if an active object is occluded, even just in a few frames, we annotated it with a (*class, x, y*) tuple, specifying the class of the object and its estimated 2D position. An example of occluded active object annotation is reported in the second row of Figure 7. With this procedure, we labeled a total of 64349 frames.

Action Annotations Starting from the temporal annotations, we defined 61 action classes<sup>4</sup>. Each action is composed by a verb and one or more objects, for example "*align screwdriver to screw*" in which the verb is *align* and the objects are *screwdriver* and *screw*. Depending on the verb and objects involved in the interaction, each temporal segment has been associated to one of the 61 considered action classes. Figure 8 shows the list of the 61 action classes, which follow a long-tail distribution.

**EHOI Annotations** Let  $O = \{o_1, o_2, ..., o_n\}$  and  $V = \{v_1, v_2, ..., v_m\}$  be the sets of objects and verbs respectively. We define an Egocentric Human-Object Interaction e as:

$$e = (v_h, \{o_1, o_2, ..., o_i\})$$
(1)

where  $v_h \in V$  is the verb characterizing the interac-

<sup>&</sup>lt;sup>3</sup>See the supplementary material for examples of partial model.

<sup>&</sup>lt;sup>4</sup>See the supplementary material for details on action class selection.

- <b>T</b>			/						
Train	11	236.47	55%	5057	37386	6/5	6/5		
Val	2	46.57	10%	977	6983	1/1	1/1		
Test	7	134.93	35%	2824	19980	4/3	4/3		
1050	,	Table 2 (	Statistics	of the three or	Lita Train Vali	dation and	Teat	—	
		Table 2. S	Statistics	of the three sp	onts: Train, vano	dation and	i Test.		
	Act	tion instanc	es		ID Acti	ion	ID Action	ID	Action
					0 check_booklet		20 put_screwdriver	40	take_red_perforated_junction_bar
					1 align_screwdriver_	to_screw	21 put_red_perforated_junction_bar	41	fit_rim_tire
					2 take_partial_mode	1	22 put_gray_angled_perforated_bar	42	take_rim
					3 plug_rod		23 take_red_perforated_bar	43	take_red_4_perforated_junction_bar
					4 screw_screw_with	_screwdriver	24 take_gray_perforated_bar	44	put_screw
					5 take_bolt		25 take_red_angled_perforated_bar	45	put_rod
					6 align_objects		26 tighten_nut_with_hands	46	put_washer
					7 take_washer		27 take_white_angled_perforated_bar	47	unscrew_screw_with_screwdriver
					8 take_screw		28 take_rod	48	put_red_perforated_bar
					9 put_white_angled_	_perforated_bar	29 put_tire	49	put_wrench
					10 unscrew_screw_wi	ith_hands	30 put_roller	50	put_bolt
					11 take_screwdriver		31 pull_partial_model	51	take_wheels_axle
					12 plug_handlebar		32 pull_screw	52	put_wheels_axle
					13 plug_screw		33 take_gray_angled_perforated_bar	53	put_red_angled_perforated_bar
1					14 tighten_nut_with_	wrench	34 take_tire	54	put_red_4_perforated_junction_bar
					15 put_gray_perforate	ed_bar	35 pull_rod	55	take_objects
					16 align_wrench_to_b	oolt	36 take_wrench	56	put_objects
HILL.					17 put_partial_model		37 browse_booklet	57	loosen_bolt_with_hands
HIIIII III IIIIIIIIIIIIIIIIIIIIIIIIIII					18 screw_screw_with	_hands	38 take_roller	58	put_booklet
					19 take_booklet		39 take_handlebar	59	put_rim
	1111111							60	put_handlebar
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 1	7 18 19 20 21 22 23 24	25 26 27 28 29 30 31 32 33 3	34 35 36 37 38 39 40 4	1 42 43 44 45 46 47 48 49 50 51 5	53 54 55 56 57 58 59 60				

#EHOIs Segments Bounding Boxes Country (ILK/Italy) Table (Light/Dark)

Figure 8. Distribution of action instances in the MECCANO dataset.

tion and  $(o_1, o_2, ..., o_i) \subseteq O$  represent the active objects involved in the interaction. Given the previous definition, we considered all the observed combinations of verbs and objects to represent EHOIs performed by the participants during the acquisition (see examples in Figure 2-bottom). Each EHOI annotation is hence composed of a verb annotation and the *active* object bounding boxes. The MECCANO dataset is the first dataset of egocentric videos explicitly annotated for the EHOI detection task.

Split | #Videos Duration (min)

2000

150

1000

500

%

# 4. Benchmarks and Baseline Results

The MECCANO dataset is suitable to study a variety of tasks, considering the challenging industrial-like scenario in which it was acquired. In this paper, we consider four tasks for which we provide baseline results: 1) Action Recognition, 2) Active Object Detection, 3) Active Object Recognition and 4) Egocentric Human-Object Interaction (EHOI) Detection. While some of these tasks have been considered in previous works, none of them has been studied in industrial scenarios from the egocentric perspective. Moreover, it is worth noting that the EHOI Detection task has never been treated in previous works. We split the dataset into three subsets (Training, Validation and Test) designed to balance the different types of desks (light, dark) and countries in which the videos have been acquired (IT, U.K.). Table 2 reports some statistics about the three splits, such as the number of videos, the total duration (in seconds), the number of temporally annotated EHOIs and the number of bounding box annotations.

#### 4.1. Action Recognition

**Task:** Action Recognition consists in determining the action performed by the camera wearer from an egocentric video segment. Specifically, let  $C_a = \{c_1, c_2, ..., c_n\}$  be the

set of action classes and let  $A_i = [t_{s_i}, t_{e_i}]$  be a video segment, where  $t_{s_i}$  and  $t_{e_i}$  are the start and the end times of the action respectively. The aim is to assign the correct action class  $c_i \in C_a$  to the segment  $A_i$ .

**Evaluation Measures:** We evaluate action recognition using Top-1 and Top-5 accuracy computed on the whole test set. As class-aware measures, we report class-mean precision, recall and  $F_1$ -score.

**Baselines:** We considered 2D CNNs as implemented in the PySlowFast library [20] (C2D), I3D [7] and SlowFast [23] networks, which are state-of-the-art methods for action recognition. In particular, for all baselines we used the PySlowFast implementation based on a ResNet-50 [32] backbone pre-trained on Kinetics [33]. See supplementary material for implementation details.

**Results:** Table 3 reports the results obtained by the baselines for the action recognition task. All baselines obtained similar performance in terms of Top-1 and Top-5 accuracy with SlowFast networks achieving slightly better performance. Interestingly, performance gaps are more consistent when we consider precision, recall and  $F_1$  scores, which is particularly relevant given the long-tailed distribution of actions in the proposed dataset (see Figure 8). Note that, in our benchmark, SlowFast obtained the best results with a Top-1 accuracy of 47.82 and an  $F_1$ -score of 41.05. See supplementary material for qualitative results. In general, the results suggest that action recognition with the MECCANO dataset is challenging and offers a new scenario to compare action recognition algorithms.

### 4.2. Active Object Detection

**Task:** The aim of the Active Object Detection task is to detect all the *active* objects involved in EHOIs. Let  $O_{act} = \{o_1, o_2, ..., o_n\}$  be the set of *active* objects in the

	Top-1 Accuracy	Top-5 Accuracy	Avg Class Precision	Avg Class Recall	Avg Class F_1-score
C2D [20]	41.92	71.95	37.6	38.76	36.49
I3D [7]	42.51	72.35	40.04	40.42	38.88

	T-1-1-2 D		fandles antisa		1-
[23]	42.85	72.47	42.11	41.48	41.05
	42.51	72.35	40.04	40.42	38.88
	41.92	71.95	37.6	38.76	36.49

Table 3. Baseline results for the action recognition task.

Method	AP (IoU >0.5)
Hand Object Detector [52]	11.17%
Hand Object Detector [52] (Avg dist.)	11.10%
Hand Object Detector [52] (All dist)	11.34%
Hand Object Detector [52] + Objs re-training	20.18%
Hand Object Detector [52] + Objs re-training (Avg dist.)	33.33%
Hand Object Detector [52] + Objs re-training (All dist.)	38.14%

SlowFast

Table 4. Baseline results for the active object detection task.

image. The goal is to detect with a bounding box each ac*tive* object  $o_i \in O_{act}$ .

Evaluation Measures: As evaluation measure, we use Average Precision (AP), which is used in standard object detection benchmarks. We set the IoU threshold equal to 0.5in our experiments.

Baseline: We considered the Hand-Object Detector proposed in [52]. The model has been designed to detect hands and objects when they are in contact. This architecture is based on Faster-RCNN [51] and predicts a box around the visible human hands, as well as boxes around the objects the hands are in contact with and a link between them. We used the Hand-Object Detector [52] pretrained on EPIC-Kitchens [13], EGTEA [37] and CharadesEGO [53] as provided by authors [52]. The model has been trained to recognize hands and to detect the active objects regardless their class. Hence, it should generalize to others domains. With default parameters, the Hand-Object Detector can find at most two active objects in contact with hands. Since our dataset tends to contain more active objects in a single EHOI (up to 7), we consider two variants of this model by changing the threshold on the distance between hands and detected objects. In the first variant, the threshold is set to the average distance between hands and active objects on the MECCANO dataset. We named this variant "Avg distance". In the second variant, we removed the thresholding operation and considered all detected objects as active objects. We named this variant "All objects". We further adapted the Hand-Object Detector [52] re-training the Faster-RCNN component to detect all active objects of the MECCANO dataset. See supplementary material for implementation details.

**Results:** Table 4 shows the results obtained by the *active* object detection task baselines. The results highlight that the Hand-Object Detector [52] is not able to generalize to a domain different than the one on which it was trained. All the three variants of the Hand-Object Detector using the original object detector obtained an AP approximately equal to 11% (first three rows of Table 4). Re-training the object detector on the MECCANO dataset allowed to improve performance by significant margins. In particular, using the standard distance threshold value, we obtained an AP of 20.18%. If we consider the average distance as the threshold to discriminate active and passive objects, we obtain an AP of 33.33%. Removing the distance threshold (last row of Table 4), allows to outperform all the previous results obtaining an AP equal to 38.14%. This suggests that adapting the general object detector to the challenging domain of the proposed dataset is key to performance. Indeed, training the object detector to detect only active objects in the scene already allows to obtain reasonable results, while there still space for improvement.

### 4.3. Active Object Recognition

Task: The task consists in detecting and recognizing the active objects involved in EHOIs considering the 20 object classes of the MECCANO dataset. Formally, let  $O_{act} = \{o_1, o_2, ..., o_n\}$  be the set of *active* objects in the image and let  $C_o = \{c_1, c_2, ..., c_m\}$  be the set of object classes. The task consists in detecting objects  $o_i \in O_{act}$ and assigning them the correct class label  $c \in C_o$ .

Evaluation Measures: We use mAP [18] with threshold on IoU equal to 0.5 for the evaluations.

Baseline: As a baseline, we used a standard Faster-RCNN [51] object detector. For each image the object detector predicts (x, y, w, h, class) tuples which represent the object bounding boxes and the associated classes. See supplementary material for implementation details.

**Results:** Table 5 reports the results obtained with the baseline in the Active Object Recognition task. We report the AP values for each class considering all the videos belonging to the test set of the MECCANO dataset. The last column shows the average of the AP values for each class and the last row reports the mAP value for the test set. The mAP was computed as the average of the mAP values obtained in each test video. AP values in the last column show that large objects are easier to recognize (e.g. instruction booklet: 46.48%; screwdriver: 60.50%; tire: 58.91%; rim: 50.35%). Performance suggests that the proposed dataset is challenging due to the presence of small objects. We leave the investigation of more specific approaches to active object detection to future studies.

### 4.4. EHOI Detection

Task: The goal is to determine egocentric human-object interactions (EHOI) in each image. Given the definition of EHOIs as <verb, objects> pairs (see Equation 1), methods should detect and recognize all the active objects in the scene, as well as the verb describing the action performed

ID	Class	AP (per class)
0	instruction booklet	46.18%
1	gray_angled_perforated_bar	09.79%
2	partial_model	36.40%
3	white_angled_perforated_bar	30.48%
4	wrench	10.77%
5	screwdriver	60.50%
6	gray_perforated_bar	30.83%
7	wheels_axle	10.86%
8	red_angled_perforated_bar	07.57%
9	red_perforated_bar	22.74%
10	rod	15.98%
11	handlebar	32.67%
12	screw	38.96%
13	tire	58.91%
14	rim	50.35%
15	washer	30.92%
16	red_perforated_junction_bar	19.80%
17	red_4_perforated_junction_bar	40.82%
18	bolt	23.44%
19	roller	16.02%
	m A D	20.20%
	шАГ	50.39%

Table 5. Baseline results for the active object recognition task.

#### by the human.

**Evaluation Measures:** Following [31, 28], we use the "*role AP*" as an evaluation measure. Formally, a detected EHOI is considered as a true positive if 1) the predicted object bounding box has a IoU of 0.5 or higher with respect to a ground truth annotation and 2) the predicted verb matches with the ground truth. Note that only the *active* object bounding box location (not the correct class) is considered in this measure. Moreover, we used different values of IoU (e.g., 0.5, 0.3 and 0.1) to compute the "*role AP*".

**Baseline:** We adopted three baselines for the EHOI detection task. The first one is based on InteractNet [28], which is composed by three branches: 1) the "human-branch" to detect the humans in the scene, 2) the "object-branch" to detect the objects and 3) the "interaction-branch" which predicts the verb of the interaction focusing on the humans and objects appearance. The second one is an extension of InteractNet which also uses context features derived from the whole input frame to help the "interaction-branch" in verb prediction. The last baseline is based on the combination of a SlowFast network [20] trained to predict the verb of the EHOI considering the spatial and temporal dimensions, and Faster-RCNN [51] which detects and recognizes all *active* objects in the frame. See supplementary material for implementation details.

**Results:** Table 6 reports the results obtained by the baselines on the test set for the EHOI detection task. The InteractNet method obtains low performance on this task with a mAP role of 4.92%. Its extension with context features, slightly improves the performance with a mAP role of 8.45%, whereas SlowFast network with Faster-RCNN achieved best results with a mAP equal to 25.93%. The results highlight that current state-of-the-art approaches developed for the analysis of still images in third person scenarios are unable to detect EHOIs in the proposed dataset, which is likely due to the presence of multiple tiny ob-

	mAP role	
$IoU \ge 0.5$	$IoU \geq 0.3$	$IoU \geq 0.1$
04.92%	05.30%	05.72%
08.45%	09.01%	09.45%
25.93%	28.04%	29.65%
	$\begin{array}{c} \textbf{IoU} \geq \textbf{0.5} \\ 04.92\% \\ 08.45\% \\ \textbf{25.93\%} \end{array}$	$\begin{tabular}{ c c c c c } \hline mAP \ role \\ \hline IoU \ge 0.5 & IoU \ge 0.3 \\ \hline 04.92\% & 05.30\% \\ \hline 08.45\% & 09.01\% \\ \hline 25.93\% & 28.04\% \\ \hline \end{tabular}$



Figure 9. Qualitative results for the EHOI detection task.

jects involved simultaneously in the EHOI and to the actions performed. On the contrary, adding the ability to process video clips with SlowFast allows for significant performance boosts. Figure 9 shows qualitative results obtained with the SlowFast+Faster-RCNN baseline. Note that in the second example the method correctly predicted all the objects involved simultaneously in the EHOI. Despite promising performance of the suggested baseline, the proposed EHOI detection task needs more investigation due to the challenging nature of the considered industrial-like domain.

# 5. Conclusion

We proposed MECCANO, the first dataset to study egocentric human-object interactions (EHOIs) in an industriallike scenario. We publicly release the dataset with both temporal (action segments) and spatial (active object bounding boxes) annotations considering a taxonomy of 12 verbs, 20 nouns and 61 unique actions. In addition, we defined the Egocentric Human-Object Interaction (EHOI) detection task and performed baseline experiments to show the potential of the proposed dataset on four challenging tasks: action recognition, *active* object detection, *active* object recognition and EHOI detection. Future works will explore approaches for improved performance on this challenging data.

### 6. Acknowledgments

This research has been supported by MIUR PON PON R&I 2014-2020 - Dottorati innovativi con caratterizzazione industriale, by MIUR AIM - Attrazione e Mobilita Internazionale Linea 1 - AIM1893589 - CUP: E64118002540007, and by MISE - PON I&C 2014-2020 - Progetto ENIGMA - Prog n. F/190050/02/X44 - CUP: B61B19000520008.

# References

- [1] Epson moverio bt 300. https://www.epson. eu/products/see-through-mobile-viewer/ moverio-bt-300.
- [2] Microsoft hololens 2. https://www.microsoft. com/en-us/hololens.
- [3] Vuzix blade. https://www.vuzix.com/products/ blade-smart-glasses.
- [4] S. Bambach, S. Lee, D. J. Crandall, and C. Yu. Lending a hand: Detecting hands and recognizing activities in complex egocentric interactions. In *ICCV*, pages 1949–1957, 2015.
- [5] M. Cai, K. M. Kitani, and Y. Sato. Understanding handobject manipulation with grasp types and object attributes. In *Robotics: Science and Systems*, 2016.
- [6] J. Carreira, E. Noland, C. Hillier, and A. Zisserman. A short note on the kinetics-700 human action dataset. *ArXiv*, abs/1907.06987, 2019.
- [7] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. *CVPR*, pages 4724– 4733, 2017.
- [8] Y. Chao, Z. Wang, Y. He, J. Wang, and J. Deng. Hico: A benchmark for recognizing human-object interactions in images. In *ICCV*, pages 1017–1025, 2015.
- [9] Y. W. Chao, Y. Liu, X. Liu, H. Zeng, and J. Deng. Learning to detect human-object interactions. WACV, pages 381–389, 2018.
- [10] S. Colombo, Y. Lim, and F. Casalegno. Deep vision shield: Assessing the use of hmd and wearable sensors in a smart safety device. In ACM PETRA, 2019.
- [11] R. Cucchiara and A. Del Bimbo. Visions for augmented cultural heritage experience. *IEEE MultiMedia*, 21(1):74–82, 2014.
- [12] N. Dalal, B. Triggs, and C. Schmid. Human detection using oriented histograms of flow and appearance. In *ECCV*, page 428–441, 2006.
- [13] D. Damen, H. Doughty, G. M. Farinella, S. Fidler, A. Furnari, E. Kazakos, D. Moltisanti, J. Munro, T. Perrett, W. Price, and M. Wray. Scaling egocentric vision: The epic-kitchens dataset. In *ECCV*, 2018.
- [14] D. Damen, H. Doughty, G. M. Farinella, S. Fidler, A. Furnari, E. Kazakos, D. Moltisanti, J. Munro, T. Perrett, W. Price, and M. Wray. The epic-kitchens dataset: Collection, challenges and baselines. *IEEE TPAMI*, 2020.
- [15] D. Damen, H. Doughty, G. M. Farinella, A. Furnari, J. Ma, E. Kazakos, D. Moltisanti, J. Munro, T. Perrett, W. Price, and M. Wray. Rescaling egocentric vision. *CoRR*, abs/2006.13256, 2020.
- [16] D. Damen, T. Leelasawassuk, O. Haines, A. Calway, and W. Mayol-Cuevas. You-do, i-learn: Discovering task relevant objects and their modes of interaction from multi-user egocentric video. In *BMVC*, 2014.
- [17] F. de la Torre, J. K. Hodgins, J. Montano, and S. Valcarcel. Detailed human data acquisition of kitchen activities: the cmu-multimodal activity database (cmu-mmac). In CHI 2009 Workshop. Developing Shared Home Behavior Datasets to Advance HCI and Ubiquitous Computing Research, 2009.

- [18] M. Everingham, S. M. Eslami, L. Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge: A retrospective. *Int. J. Comput. Vision*, 111(1):98–136, Jan. 2015.
- [19] B. Ghanem F. C. Heilbron, V. Escorcia and J. C. Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, pages 961–970, 2015.
- [20] H. Fan, Y. Li, B. Xiong, W. Y. Lo, and C. Feichtenhofer. Pyslowfast. https://github.com/ facebookresearch/slowfast, 2020.
- [21] K. Fang, T. Wu, D. Yang, S. Savarese, and J. J. Lim. Demo2vec: Reasoning object affordances from online videos. In CVPR, pages 2139–2147, 2018.
- [22] G. M. Farinella, G. Signorello, S. Battiato, A. Furnari, F. Ragusa, R. Leonardi, E. Ragusa, E. Scuderi, A. Lopes, L. Santo, and M. Samarotto. Vedi: Vision exploitation for data interpretation. In *ICIAP*, 2019.
- [23] C. Feichtenhofer, H. Fan, J. Malik, and K. He. Slowfast networks for video recognition. In *ICCV*, pages 6202–6211, 2018.
- [24] C. Feichtenhofer, A. Pinz, and R. P. Wildes. Spatiotemporal residual networks for video action recognition. In *NeurIPS*, NIPS'16, page 3476–3484, Red Hook, NY, USA, 2016. Curran Associates Inc.
- [25] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional two-stream network fusion for video action recognition. In *CVPR*, 2016.
- [26] R. Girshick. Fast R-CNN. In ICCV, 2015.
- [27] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.
- [28] G. Gkioxari, R. B. Girshick, P. Dollár, and K. He. Detecting and recognizing human-object interactions. *CVPR*, pages 8359–8367, 2018.
- [29] R. Goyal, S. E. Kahou, V. Michalski, J. Materzynska, S. Westphal, H. Kim, V. Haenel, I. Fruend, P. Yianilos, M. Mueller-Freitag, F. Hoppe, C. Thurau, I. Bax, and R. Memisevic. The "something something" video database for learning and evaluating visual common sense. In *ICCV*, pages 5843–5851, 2017.
- [30] A. Gupta, A. Kembhavi, and L. S. Davis. Observing humanobject interactions: Using spatial and functional compatibility for recognition. *IEEE TPAMI*, 31(10), 2009.
- [31] S. Gupta and J. Malik. Visual semantic role labeling. ArXiv, abs/1505.04474, 2015.
- [32] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. arXiv preprint arXiv:1512.03385, 2015.
- [33] W. Kay, J. Carreira, K. Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, F. Viola, T. Green, T. Back, A. Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset. *ArXiv*, abs/1705.06950, 2017.
- [34] E. Kazakos, A. Nagrani, A. Zisserman, and D. Damen. Epicfusion: Audio-visual temporal binding for egocentric action recognition. In *ICCV*, 2019.

- [35] A. Krizhevsky, S. Ilya, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *NeurIPS*, pages 1097–1105. 2012.
- [36] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld. Learning realistic human actions from movies. In *CVPR*. IEEE Computer Society, 2008.
- [37] Y. Li, M. Liu, and J. M. Rehg. In the eye of beholder: Joint learning of gaze and actions in first person video. In ECCV, 2018.
- [38] Yue Liao, Si Liu, Fei Wang, Yanjie Chen, Chen Qian, and Jiashi Feng. PPDM: Parallel point detection and matching for real-time human-object interaction detection. In *CVPR*, 2020.
- [39] J. Lin, C. Gan, and S. Han. Tsm: Temporal shift module for efficient video understanding. In *ICCV*, pages 7082–7092, 2019.
- [40] T. Y. Lin, M. Maire, S. Belongie, L. Bourdev, R. Girshick, J. Hays, P. Perona, D. Ramanan, C. L. Zitnick, and P. Dollár. Microsoft coco: Common objects in context, 2014. cite arxiv:1405.0312Comment: 1) updated annotation pipeline description and figures; 2) added new section describing datasets splits; 3) updated author list.
- [41] M. Ma, H. Fan, and K. M. Kitani. Going deeper into firstperson activity recognition. In CVPR, 2016.
- [42] T. Nagarajan, C. Feichtenhofer, and K. Grauman. Grounded human-object interaction hotspots from video. In *ICCV*, pages 8687–8696, 2019.
- [43] T. Nagarajan, Y. Li, C. Feichtenhofer, and K. Grauman. Ego-topo: Environment affordances from egocentric video. *ArXiv*, abs/2001.04583, 2020.
- [44] A. Ortis, G. Farinella, V. D'Amico, Luca Addesso, Giovanni Torrisi, and S. Battiato. Organizing egocentric videos of daily living activities. *Pattern Recognition*, 72, 2017.
- [45] H. Pirsiavash and D. Ramanan. Detecting activities of daily living in first-person camera views. In CVPR, 2012.
- [46] Hamed Pirsiavash and Deva Ramanan. Detecting activities of daily living in first-person camera views. In *CVPR*, pages 2847–2854. IEEE Computer Society, 2012.
- [47] S. Qi, W. Wang, B. Jia, J. Shen, and S. C. Zhu. Learning human-object interactions by graph parsing neural networks. *ArXiv*, abs/1808.07962, 2018.
- [48] Z. Qiu, T. Yao, and T. Mei. Learning spatio-temporal representation with pseudo-3d residual networks. pages 5534– 5542, 10 2017.
- [49] F. Ragusa, A. Furnari, S. Battiato, G. Signorello, and G. M. Farinella. EGO-CH: Dataset and fundamental tasks for visitors behavioral understanding using egocentric vision. *Pattern Recognition Letters*, 2020.
- [50] J. Redmon and A. Farhadi. Yolov3: An incremental improvement. CoRR, abs/1804.02767, 2018.
- [51] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks.
- [52] D. Shan, J. Geng, M. Shu, and D. Fouhey. Understanding human hands in contact at internet scale. 2020.

- [53] G. A. Sigurdsson, A. Gupta, C. Schmid, A. Farhadi, and A. Karteek. Actor and observer: Joint modeling of first and third-person videos. *CVPR*, pages 7396–7404, 2018.
- [54] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *NeurIPS*, 2014.
- [55] B. Soran, A. Farhadi, and L. Shapiro. Generating notifications for missing actions: Don't forget to turn the lights off! pages 4669–4677, 12 2015.
- [56] S. Sudhakaran, S. Escalera, and O. Lanz. Lsta: Long shortterm attention for egocentric action recognition. In *CVPR*, June 2019.
- [57] Y. Tang, Y. Tian, J. Lu, J. Feng, and J. Zhou. Action recognition in rgb-d egocentric videos. In *ICIP*, pages 3410–3414, 2017.
- [58] G. W. Taylor, R. Fergus, Y. LeCun, and C. Bregler. Convolutional learning of spatio-temporal features. In *ECCV*, page 140–153, 2010.
- [59] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, pages 4489–4497, 2015.
- [60] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, pages 6450–6459, 2018.
- [61] G. Varol, I. Laptev, and C. Schmid. Long-term temporal convolutions for action recognition. *IEEE TPAMI*, 40(6):1510– 1517, 2018.
- [62] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool. Temporal segment networks: Towards good practices for deep action recognition. volume 9912, 10 2016.
- [63] T. Wang, T. Yang, M. Danelljan, F. S. Khan, X. Zhang, and J. Sun. Learning human-object interaction detection using interaction points. In *CVPR*, June 2020.
- [64] S. Xie, C. Sun, J. Huang, Z. Tu, and K. Murphy. Rethinking spatiotemporal feature learning for video understanding. *CoRR*, abs/1712.04851, 2017.
- [65] B. Yao and L. Fei-Fei. Recognizing human-object interactions in still images by modeling the mutual context of objects and human poses. *IEEE TPAMI*, 34(9):1691–1703, 2012.
- [66] B. Zhou, A. Andonian, and A. Torralba. Temporal relational reasoning in videos. *ArXiv*, abs/1711.08496, 2018.
- [67] P. Zhou and M. Chi. Relation parsing neural network for human-object interaction detection. In *ICCV*, pages 843– 851, 2019.