Egocentric Vision: Exploring User-Centric Perspectives

Francesco Ragusa

LIVE Group @ UNICT - https://iplab.dmi.unict.it/live/

Next Vision - http://www.nextvisionlab.it/

Department of Mathematics and Computer Science - University of Catania

francesco.ragusa@unict.it - https://francescoragusa.github.io/

Agenda

1) Part I: History and motivations [14.15 - 15.45]
   a) Agenda of the tutorial;
   b) Perception and Egocentric Vision;
   c) Seminal works in Egocentric Vision;
   d) Differences between Third Person and First Person Vision;
   e) First Person Vision datasets;
   f) Wearable devices to acquire/process first person visual data;
   g) Main research trends in First Person (Egocentric) Vision;
   h) What's next?

Coffee Break [15.45 – 16.00]

Keynote presentation: Julien Pettrè [16.00 – 17.00]

1) Part II: Fundamental tasks for First Person Vision systems [17.15 – 18.30]
   a) Localization;
   b) Hand/Object Detection;
   c) Action/Activity Recognition;
   d) Human-Object Interaction;
   e) Anticipation;
   f) Industrial Applications;
   g) Conclusion.

Before we begin...

The slides of this tutorial are available online at:
https://francescoragusa.github.io/visigrapp2025

Part II

Fundamental Tasks for First Person Vision Systems

Data Acquisition

Four things to pay attention to when collecting first person visual data

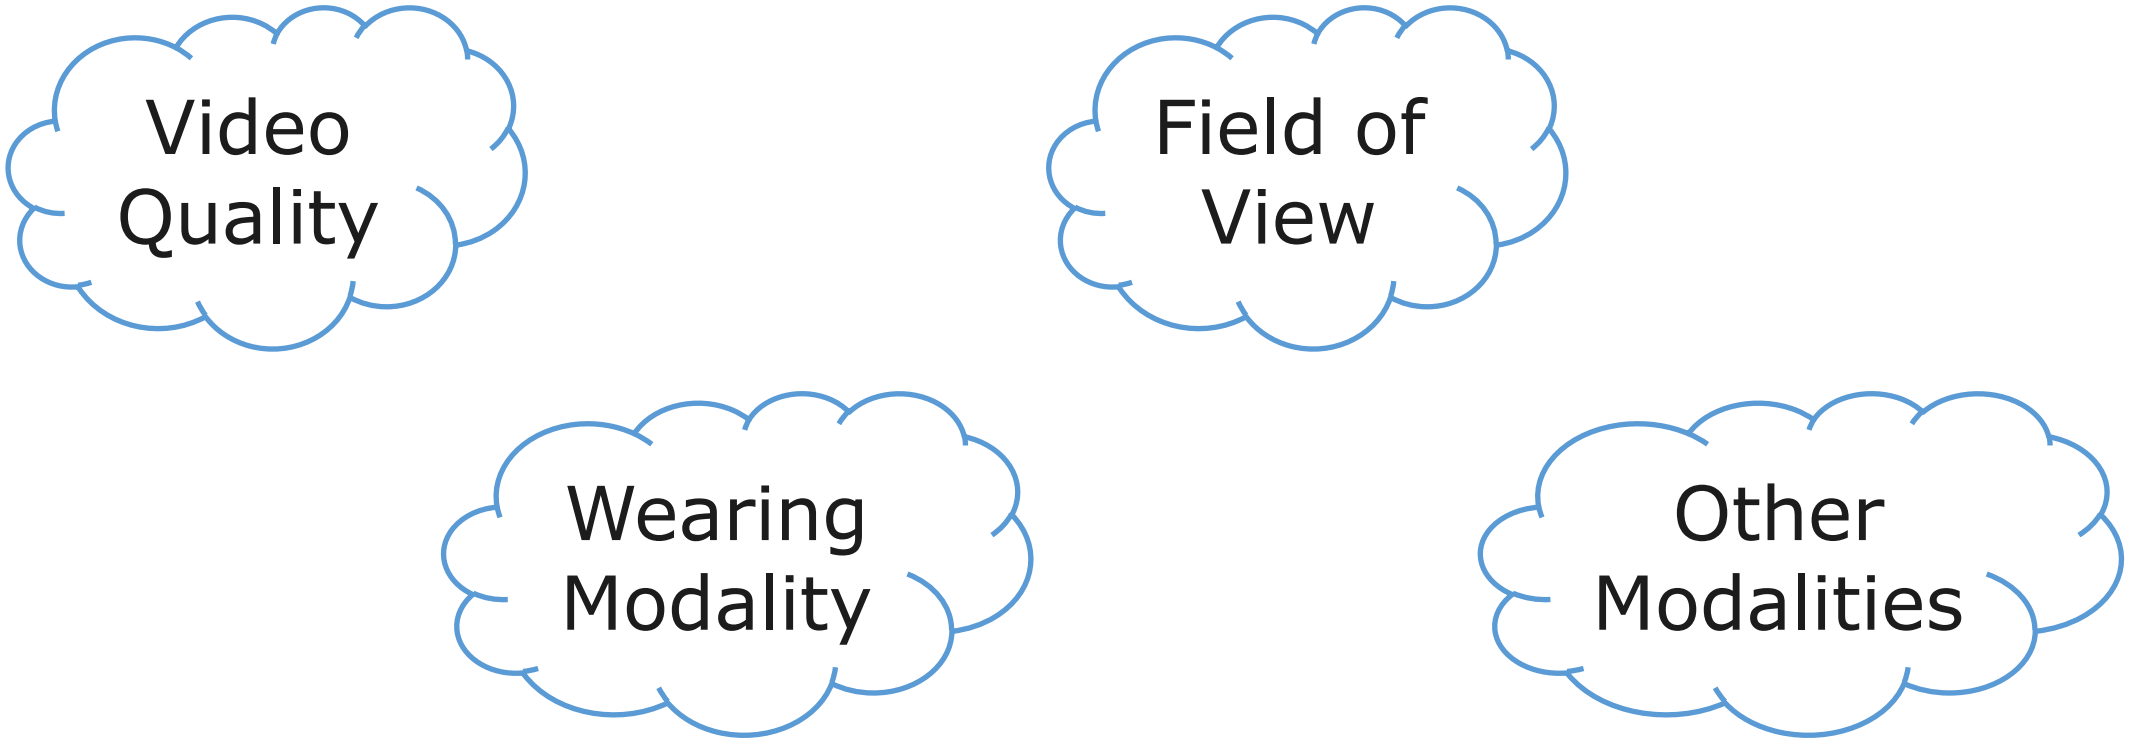

Data Acquisition – Video Quality

- Try to get a high quality camera to get high quality images!
- Egocentric video is subject to motion blur and exposure issues.

High Quality Video Obtained with a GoPro

Average Quality Video

Data Acquisition – Camera Wearing Modalities

A,B: head mounted, D: chest mounted

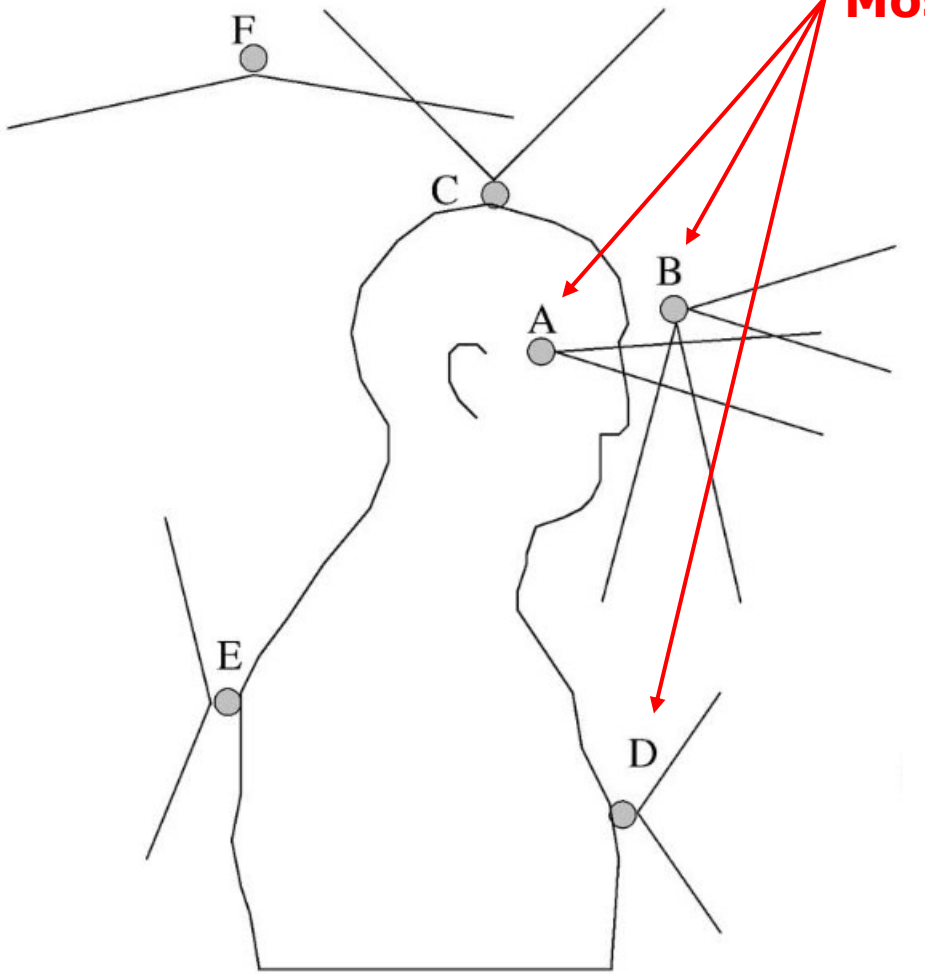

Most Common Wearing Modalities

A

B (frontward)

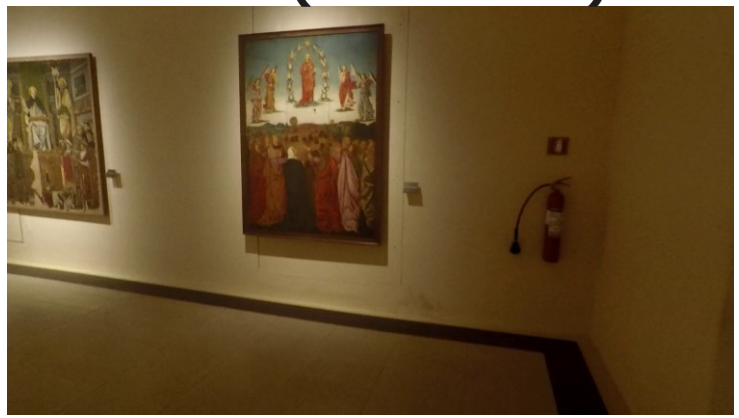

B (downward)

D

Mayol-Cuevas, W. W., Tordoff, B. J., & Murray, D. W. (2009). On the choice and placement of wearable vision sensors. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, *39*(2), 414-425.

Data Acquisition – Camera Wearing Modalities

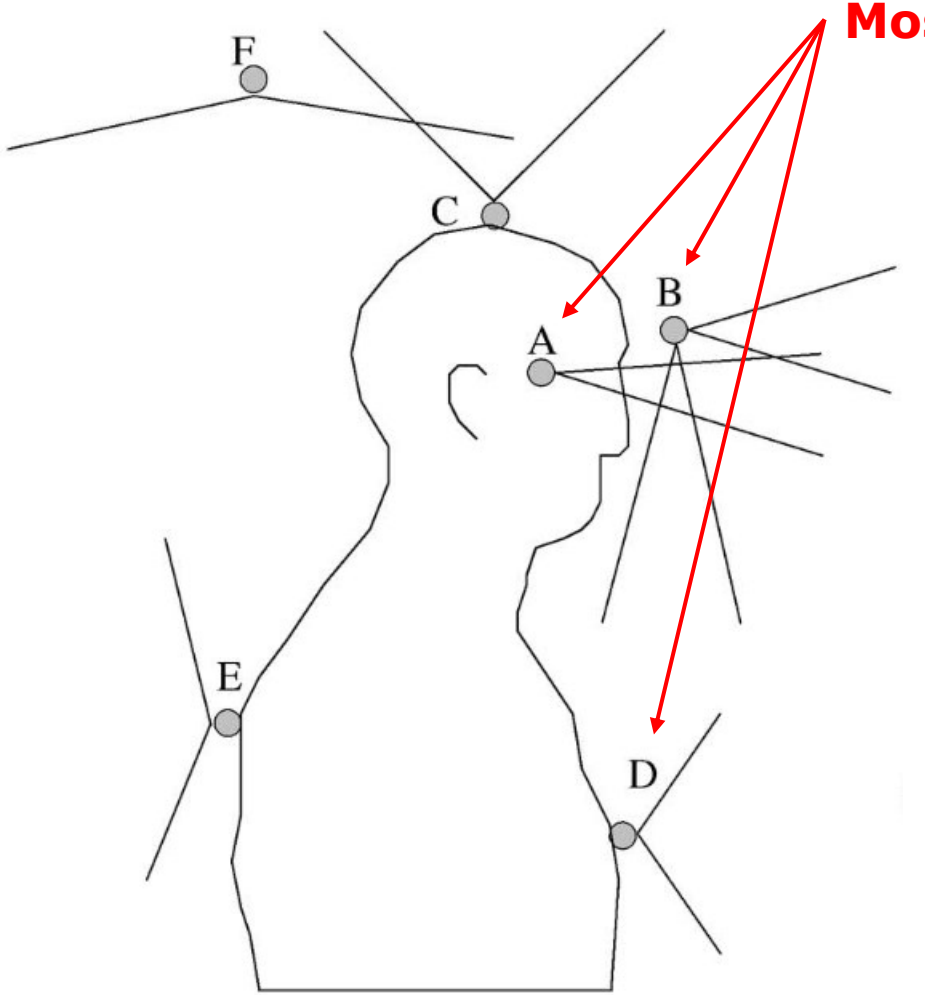

Most Common Wearing Modalities

- A-B are best to capture objects:
  - A, B (frontward) to capture objects in front of the subjects (e.g., paintings in a museum);
  - B (downward) to capture objects manipulated with hands (e.g., kitchen);
- Chest-mounted cameras (D) are less obtrusive and give stable video, but they may miss details on what the user is looking at;

Mayol-Cuevas, W. W., Tordoff, B. J., & Murray, D. W. (2009). On the choice and placement of wearable vision sensors. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, *39*(2), 414-425.

Data Acquisition – Field of View (FOV)

A wide FOV allows to capture more scene but it may introduce distortion

Narrow Angle

Wide Angle

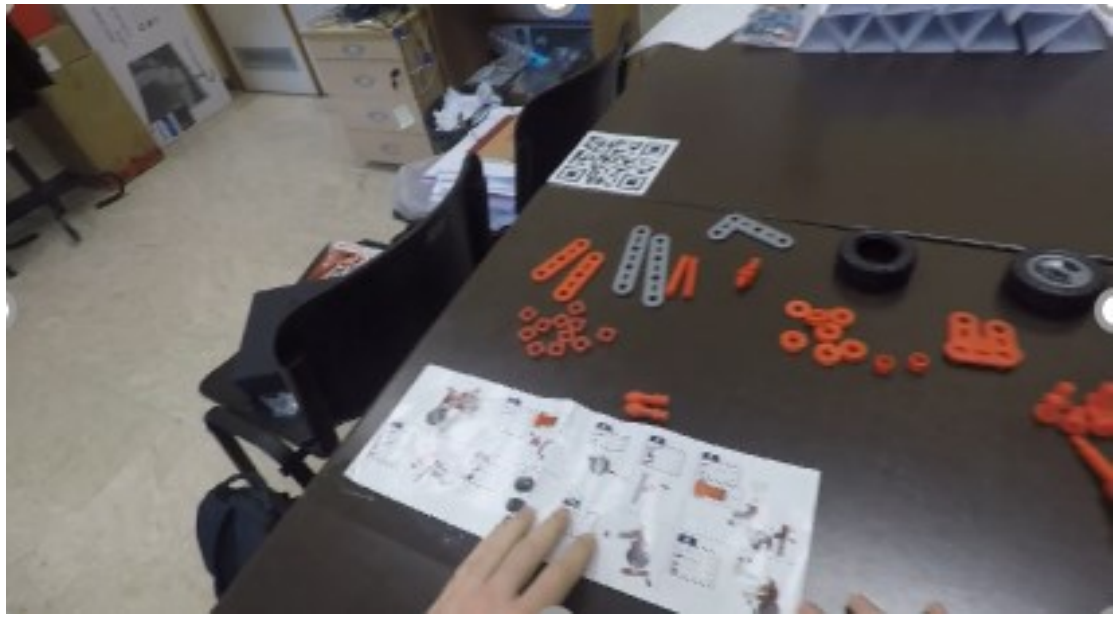

Data Acquisition – Other Modalities – Depth

- Depth can improve scene understanding by highlighting the position of objects and hands;

Wan, S., & Aggarwal, J. K. (2015). Mining discriminative states of hands and objects to recognize egocentric actions with a wearable RGBD camera. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops* (pp. 36-43).

Data Acquisition – Other Modalities – Depth (2)

https://github.com/microsoft/HoloLensForCV

- Microsoft HoloLens has a «Research Mode» which allows to access:
  - short-range depth
  - long-range depth;
  - IR reflectivity;

Microsoft HoloLens Research Mode

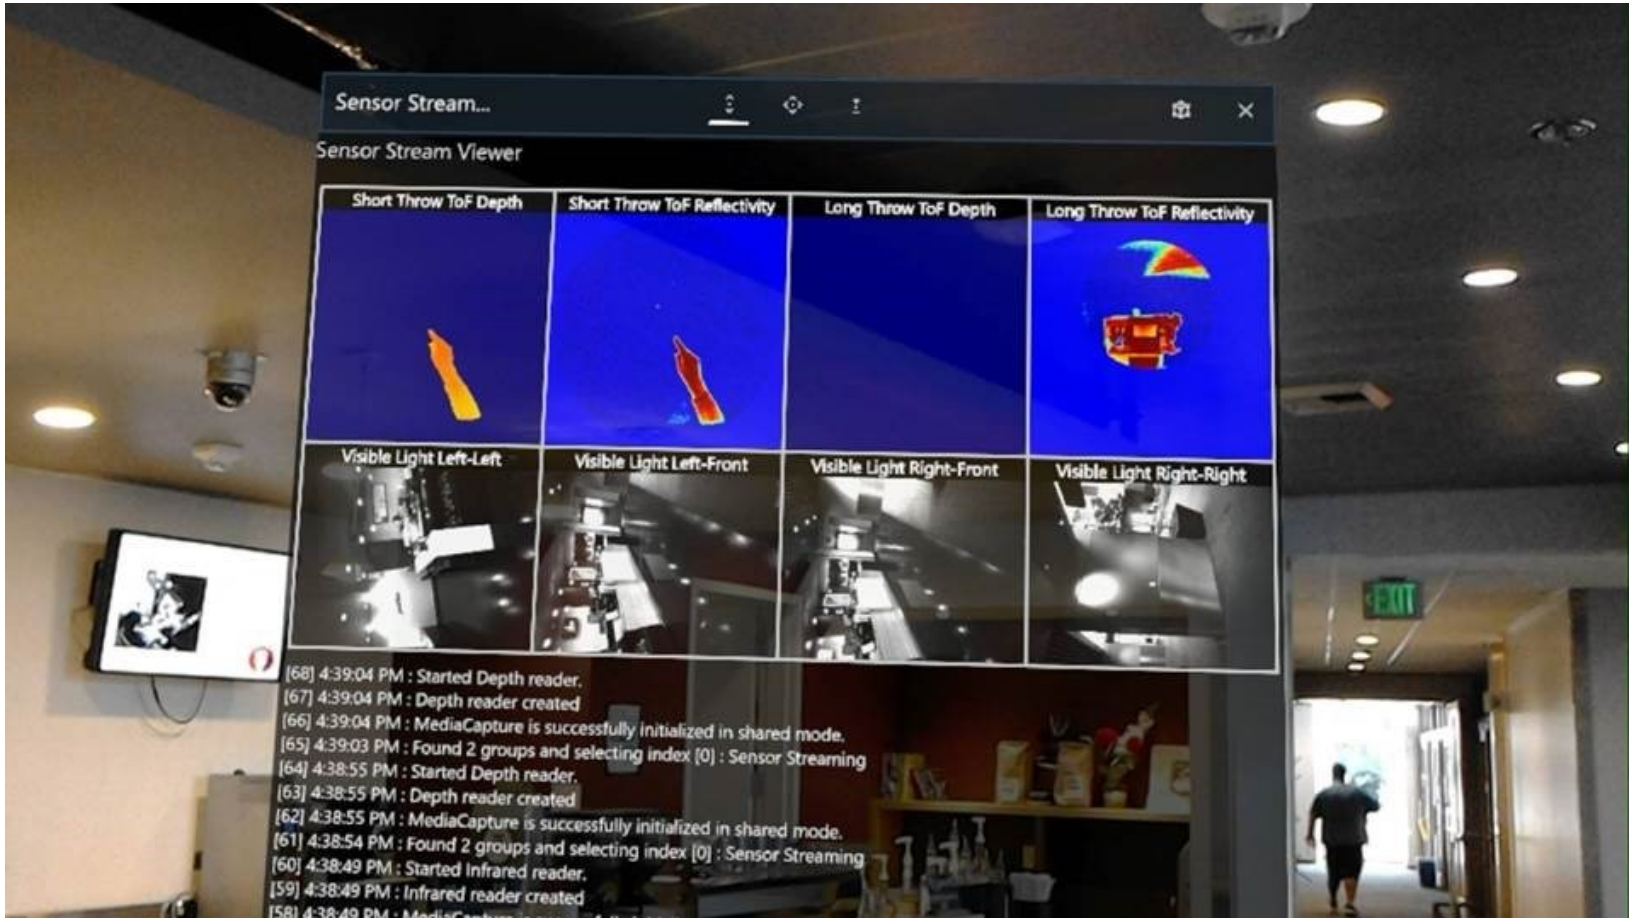

https://docs.microsoft.com/en-us/windows/mixed-reality/research-mode

Data Acquisition – Other Modalities – Gaze

Gaze can give information on what the user is paying attention to.

However, gaze trackers generally require a calibration process (and some expertise).

F. Ragusa, A. Furnari, S. Livatino, G. M. Farinella. The MECCANO Dataset: Understanding Human-Object Interactions from Egocentric Videos in an Industrial-like Domain. WACV 2021 (ORAL) (https://arxiv.org/abs/2010.05654).

Datasets

MECCANO

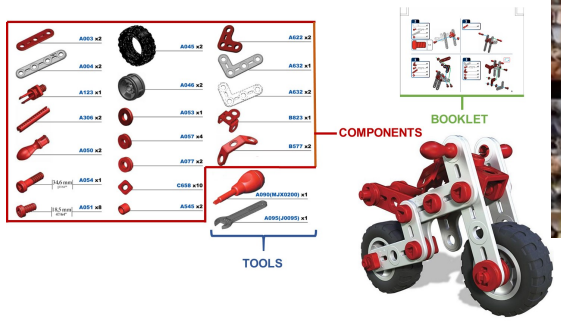

EPIC KITCHENS

HOI4D: A 4D Egocentric Dataset for Category-Level Human-Object Interaction

EGO 4D

ASSEMBLY101

ADL

EGO-CH

EGTEA Gaze+

Charades-Ego

Plizzari, C., Goletto, G., Furnari, A., Bansal, S., Ragusa, F., Farinella, G. M., Damen., D. & Tommasi, T. (2023). An Outlook into the Future of Egocentric Vision. *IJCV, 2024*.

And More

Datasets (non-exhaustive)

Dataset	URL	Settings	Annotations	Goal
EGO-EXO4D	https://ego-exo4d-data.org/	839 participants performing procedural and physical activities.	Natural language descriptions, segmentation masks, temporal segments of keysteps, task-graphs, profiency labels, 3D human pose	Keystep Recognition, Proficiency Estimation, Relation, Pose Estimation
EGO4D	https://ego4d-data.org/	931 participants performing different activities in different domains.	Different temporal and spatial annotations related to 5 benchmarks	Episodic Memory, Hand-Object Interaction, Audio-Visual Diarization, Social Interactions, Forecasting
EPIC-KITCHENS-100	https://epic-kitchens.github.io/2020-100	Subjects performing unscripted actions in their native kitchens.	Temporal segments	Action recognition, detection, anticipation, retrieval.
MECCANO	https://iplab.dmi.unict.it/MECCANO/	20 subjects assembling a toy motorbike.	Temporal segments, active objects, human-object interactions	Action recognition, Active object detection, Egocentric Human-Object Interaction Detection
ASSEMBLY101	https://assembly-101.github.io/	53 subjects assembling in a cage settings 101 children's toys.	Temporal segments, 3D hand poses	Action recognition, Action Anticipation, Temporal Segmentation

Datasets (non-exhaustive)

Dataset	URL	Settings	Annotations	Goal
ENIGMA-51	https://iplab.dmi.unict.it/ENIGMA-51/	Participants performing procedural activities in the industrial domain.	Textual procedures, Hand and Object annotations, human-object interactions, next-object interactions	Untrimmed temporal annotations of human-object interactions, Egocentric Human-object interactions, short-term object interaction anticipation, NLU of intents and entities
HOLOASSIST	https://holoassist.github.io/	350 instructor-performer pairs which collaboratively complete physical manipulation tasks.	Action and conversational annotations	Action recognition and anticipation, mistake detection, inervention type prediction, 3D hand pose forecasting
ARIA Digital Twin	https://www.projectaria.com/datasets/adt/			
IndustReal	https://timschoonbeek.github.io/industreal.html	Participants performing procedural activities building a toy model of a car	Step and mistake annotations	Action recognition, assembly state detection, procedure step recognition

Datasets (non-exhaustive)

Dataset	URL	Settings	Annotations	Goal
EPIC-KITCHENS 2018	https://epic-kitchens.github.io/2018	32 subjects performing unscripted actions in their native environments	action segments, object annotations	Action recognition, Action Anticipation, Object Detection
Charade-Ego	https://allenai.org/plato/charades/	paired first-third person videos	action classes	Action recognition
EGTEA Gaze+	http://ai.stanford.edu/~alireza/GTEA/	32 subjects, 86 sessions, 28 hours	action segments, gaze, hand masks	Understading daily activities, action recognition
ADL	https://www.csee.umbc.edu/~hpirsiav/papers/ADLdataset/	20 subjects performing daily activities in their native environments	activity segments, objects	Detecting activities of daily living
CMU kitchen	http://www.cs.cmu.edu/~espriggs/cmu-mmac/annotations/	multimodal, 18 subjects cooking 5 different recipes: brownies, eggs, pizza, salad, sandwiche	action segments	Understading daily activities
EgoSeg	http://www.vision.huji.ac.il/egoseg/	Long term actions (walking, running, driving, etc.)	long term activity	Temporal Segmentation, Indexing

Datasets (non-exhaustive)

Dataset	URL	Settings	Annotations	Goal
First-Person Social Interactions	http://ai.stanford.edu/~alireza/Disney/	8 subjects at disneyworld	Activities: walking, waiting, gathering, sitting, buying something, eating, etc.	Recognizing social interactions
UEC Dataset	http://www.cs.cmu.edu/~kkitani/datasets/	two choreographed datasets with different egoactions (walk, jump, climb, etc.) + 6 youtube sports videos	activities	Unsupervised activity recognition
JPL	http://michaelryoo.com/jpl-interaction.html	interaction with a robot	activities performed on the robot + pose	Interaction recognition/prediction
Multimodal Egocentric Activity Dataset	http://people.sutd.edu.sg/~1000892/dataset	15 seconds clips of 20 activities	activity (walking, elevator, etc.)	Life-logging
LENA: An egocentric video database of visual lifelog	http://people.sutd.edu.sg/~1000892/dataset	13 activities performed by 10 subjects (Google Glass)	activity (walking, elevator, etc.)	Life-logging

Datasets (non-exhaustive)

Dataset	URL	Settings	Annotations	Goal
FPPA	http://tamaraberg.com/prediction/Prediction.html	Five subjects performing 5 daily actions	activity (drinking water, putting on clothes, etc.)	Temporal prediction
UT Egocentric	http://vision.cs.utexas.edu/projects/egocentric/index.html	3-5 hours long videos capturing a person's day	important regions	Summarization
VINST/ Visual Diaries	http://www.csc.kth.se/cvap/vinst/NovEgoMotion.html	31 videos capturing the visual experience of a subject walkin from metro station to work	location id, novel egomotion	Novelty detection
Bristol Egocentric Object Interaction (BEOID)	https://www.cs.bris.ac.uk/~damen/BEOID/	8 subjects, six locations. Interaction with objects and environment	gaze, objects, mode of interaction (pick, plug, etc.)	Provide assistance on object usage
Object Search Dataset	https://github.com/Mengmi/deepfuturegaze_gan	57 sequences of 55 subjects on search and retrieval tasks	gaze	gaze prediction

Datasets (non-exhaustive)

Dataset	URL	Settings	Annotations	Goal
UNICT-VEDI	http://iplab.dmi.unict.it/VEDI/	different subjects visiting a museum	location, observed objects	localizing visitors of a museum and estimating their attention
UNICT-VEDI-POI	http://iplab.dmi.unict.it/VEDI_POIs/	different subjects visiting a museum	object bounding boxes annotations, observed objects	recognizing points of interest observed by the visitors
Simulated Egocentric Navigations	http://iplab.dmi.unict.it/SimulatedEgocentricNavigations/	simulated navigations of a virtual agent within a large building	3-DOF pose of the agent in each image	egocentric localization
EgoCart	http://iplab.dmi.unict.it/EgocentricShoppingCartLocalization/	egocentric images collected by a shopping cart in a retail store	3-DOF pose of the shopping cart in each image	egocentric localization
Unsupervised Segmentation of Daily Livign Activities	http://iplab.dmi.unict.it/dailylivingactivities	egocentric videos of daily activities	activities	unsupervised segmentation with respect to the activities

Datasets (non-exhaustive)

Dataset	URL	Settings	Annotations	Goal
Visual Market Basket Analysis	http://iplab.dmi.unict.it/vmba/	egocentric images colelcted by a shopping cart in a retail store	class-location of each image	egocentric localization
Location Based Segmentation of Egocentric Videos	http://iplab.dmi.unict.it/PersonalLocationSegmentation/	egocentric videos of daily activities	location classes	egocentric localization, video indexing
Recognition of Personal Locations from Egocentric Videos	http://iplab.dmi.unict.it/PersonalLocations/	egocentric videos clips of daily activities	location classes	recognizing personal locations
EgoGesture	http://www.nlpr.ia.ac.cn/iva/yfzhang/datasets/egogesture.html	2k videos from 50 subjects performing 83 gestures	Gesture labels, depth	Gesture recognition
EgoHands	http://vision.soic.indiana.edu/projects/egohands/	48 videos of interactions between two people	Hand segmentation masks	Egocentric hand segmentation
DoMSEV	http://www.verlab.dcc.ufmg.br/semantic-hyperlapse/cvpr2018-dataset/	80 hours/different activities	Scene/Action labels with IMU, GPS mad depth	Summarization

Datasets (non-exhaustive)

Dataset	URL	Settings	Annotations	Goal
EGO-HPE	http://imagelab.ing.unimore.it/imagelab2015/researchactivity.asp?idAttivita=23	Egocentric videos for head pose estimation	Head pose of the subjects	Head-pose estimation
EGO-GROUP	http://imagelab.ing.unimore.it/imagelab2015/researchactivity.asp?idAttivita=23	18 videos of people engaging social relationships	Social relationships	Understanding social relationships
DR(eye)VE	http://aimagelab.ing.unimore.it/dreyeve	74 videos of people driving	Eye fixations	Autonomous and assisted driving
THU-READ	http://ivg.au.tsinghua.edu.cn/dataset/THU_READ.php	8 subjects performing 40 actions with a head-mounted RGBD camera	Action segments	RGBD egocentric action recognition
EGO-CH	https://iplab.dmi.unict.it/EGO-CH/	70 subjects visiting two cultural sites in Sicily, Italy.	Temporal segments, room-based localization, objects	Room-basd localization, Object detection, Behavioral analysis

Fundamental Tasks

12 Egocentric Vision Research Tasks

1. Localisation
2. 3D Scene Understanding
3. Anticipation
4. Action Recognition
5. Gaze Understanding and Prediction
6. Social Behaviour Understanding
7. Full Body Pose Estimation
8. Hand and Hand-Object Interactions
9. Person Identification
10. Privacy
11. Summarisation
12. Visual Question Answering

Plizzari, C., Goletto, G., Furnari, A., Bansal, S., Ragusa, F., Farinella, G. M., Damen., D. & Tommasi, T. (2023). An Outlook into the Future of Egocentric Vision. *IJCV, 2024*.

Fundamental Tasks

12 Egocentric Vision Research Tasks

1. **Localisation**
2. 3D Scene Understanding
3. **Anticipation**
4. **Action Recognition**
5. Gaze Understanding and Prediction
6. Social Behaviour Understanding
7. Full Body Pose Estimation
8. **Hand and Hand-Object Interactions**
9. Person Identification
10. Privacy
11. Summarisation
12. Visual Question Answering

Plizzari, C., Goletto, G., Furnari, A., Bansal, S., Ragusa, F., Farinella, G. M., Damen., D. & Tommasi, T. (2023). An Outlook into the Future of Egocentric Vision. *IJCV, 2024*.

Localization

Localization – Levels of Granularity

SCENE RECOGNITION

off-the-shelf detectors

CAMERA POSE-ESTIMATION

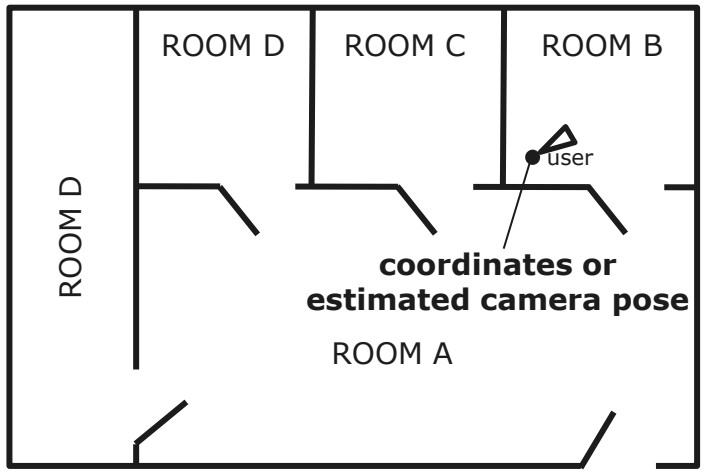

3D reconstruction of the building

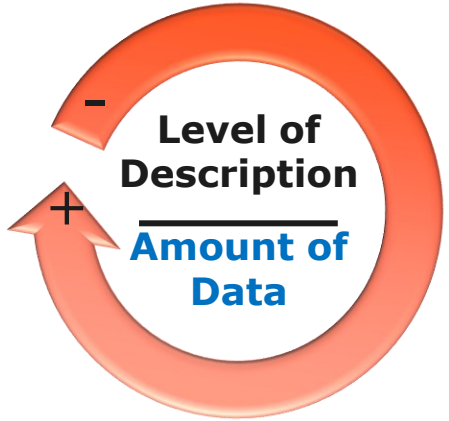

ROOM-LEVEL RECOGNITION

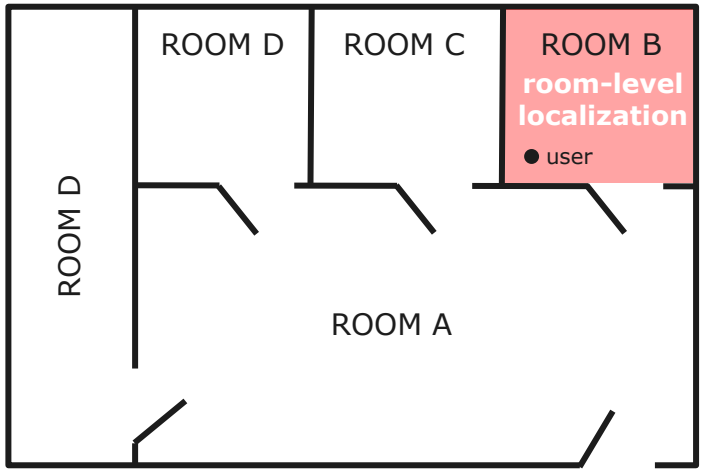

moderate amount of training data

Scene Recognition

- The most basic form of localization;
- Tells what kind of scene the user is in;
- Useful to distinguish between (even for unseen places) :
  - indoor/outdoor
  - natural/artificial
  - conf. room
  - Office
- Can use off-the-shelf detections.

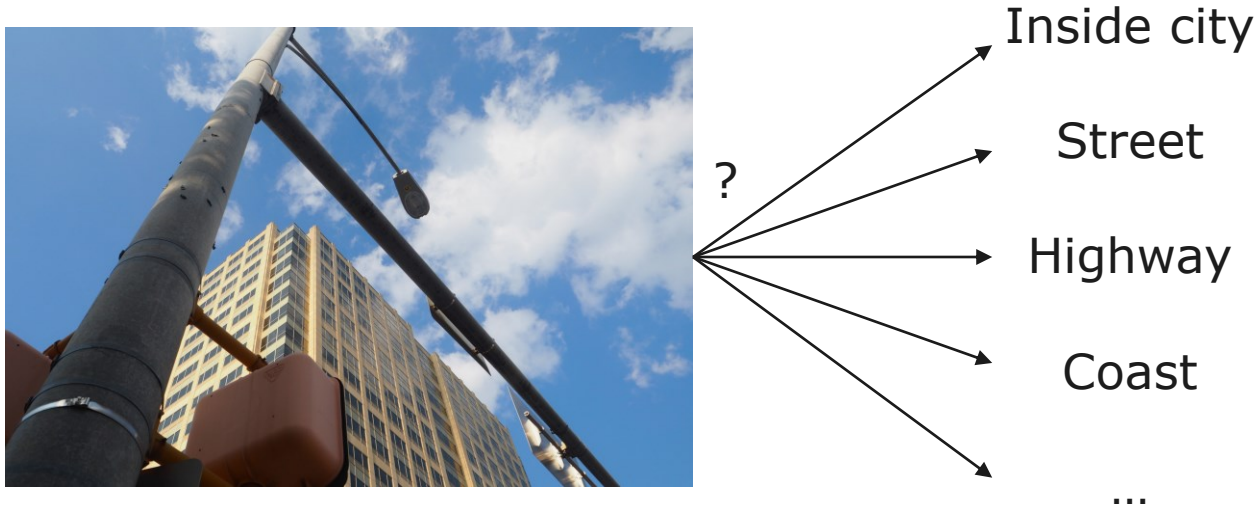

Scene Recognition – Places

DATA & CODE HERE -> http://places2.csail.mit.edu/

GT: cafeteria
top-1: cafeteria (0.179)
top-2: restaurant (0.167)
top-3: dining hall (0.091)
top-4: coffee shop (0.086)
top-5: restaurant patio (0.080)

- Places is a large (10M images – 400+ classes) dataset for scene recognition;
- CNN models trained to recognize 365 scene classes available for download;
- Can be used off-the-shelf!

A 10 million Image Database for Scene Recognition B. Zhou, A. Lapedriza, A. Khosla, A. Oliva, and A. Torralba *IEEE Transactions on Pattern Analysis and Machine Intelligence, 2017*

Localization – Levels of Granularity

SCENE RECOGNITION

off-the-shelf detectors

CAMERA POSE-ESTIMATION

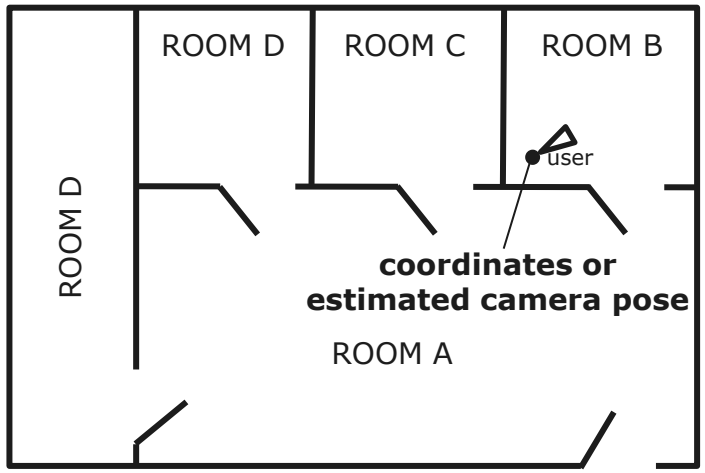

3D reconstruction of the building

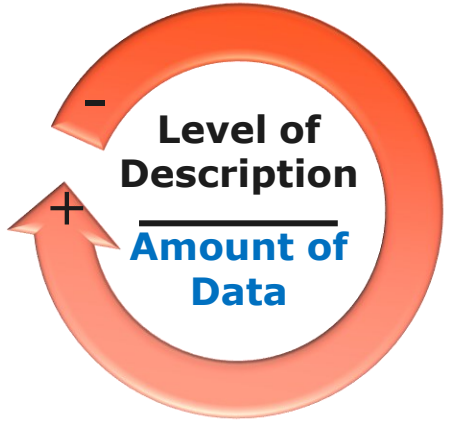

ROOM-LEVEL RECOGNITION

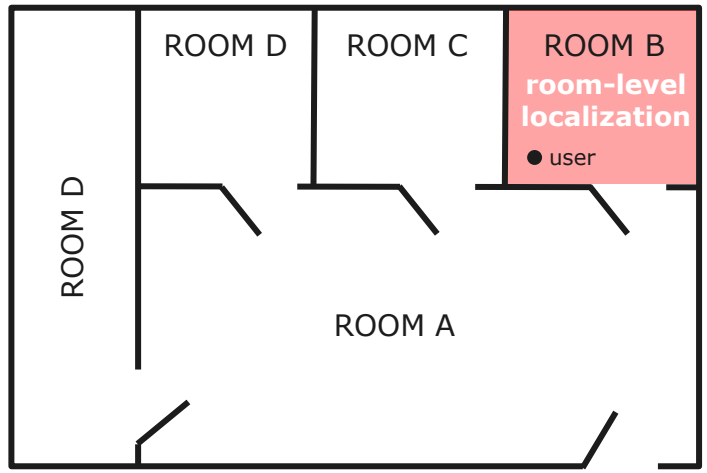

moderate amount of training data

Room-Level Localization – Museums

Cultural Site (e.g., museum) divided into contexts (e.g., rooms)

(videos acquired in the different contexts)

Training Set

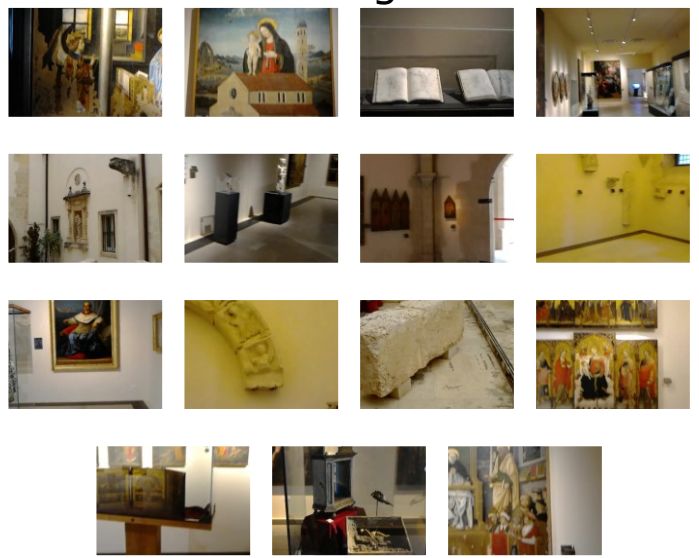

(frames extracted from videos acquired in the different contexts)

F. Ragusa, A. Furnari, S. Battiato, G. Signorello, G. M. Farinella (2020). EGO-CH: Dataset and Fundamental Tasks for Visitors Behavioral Understanding using Egocentric Vision . Pattern Recognition Letters.

Room-Level Localization – Full Model

CODE HERE -> https://iplab.dmi.unict.it/VEDI/

https://iplab.dmi.unict.it/PersonalLocationSegmentation/

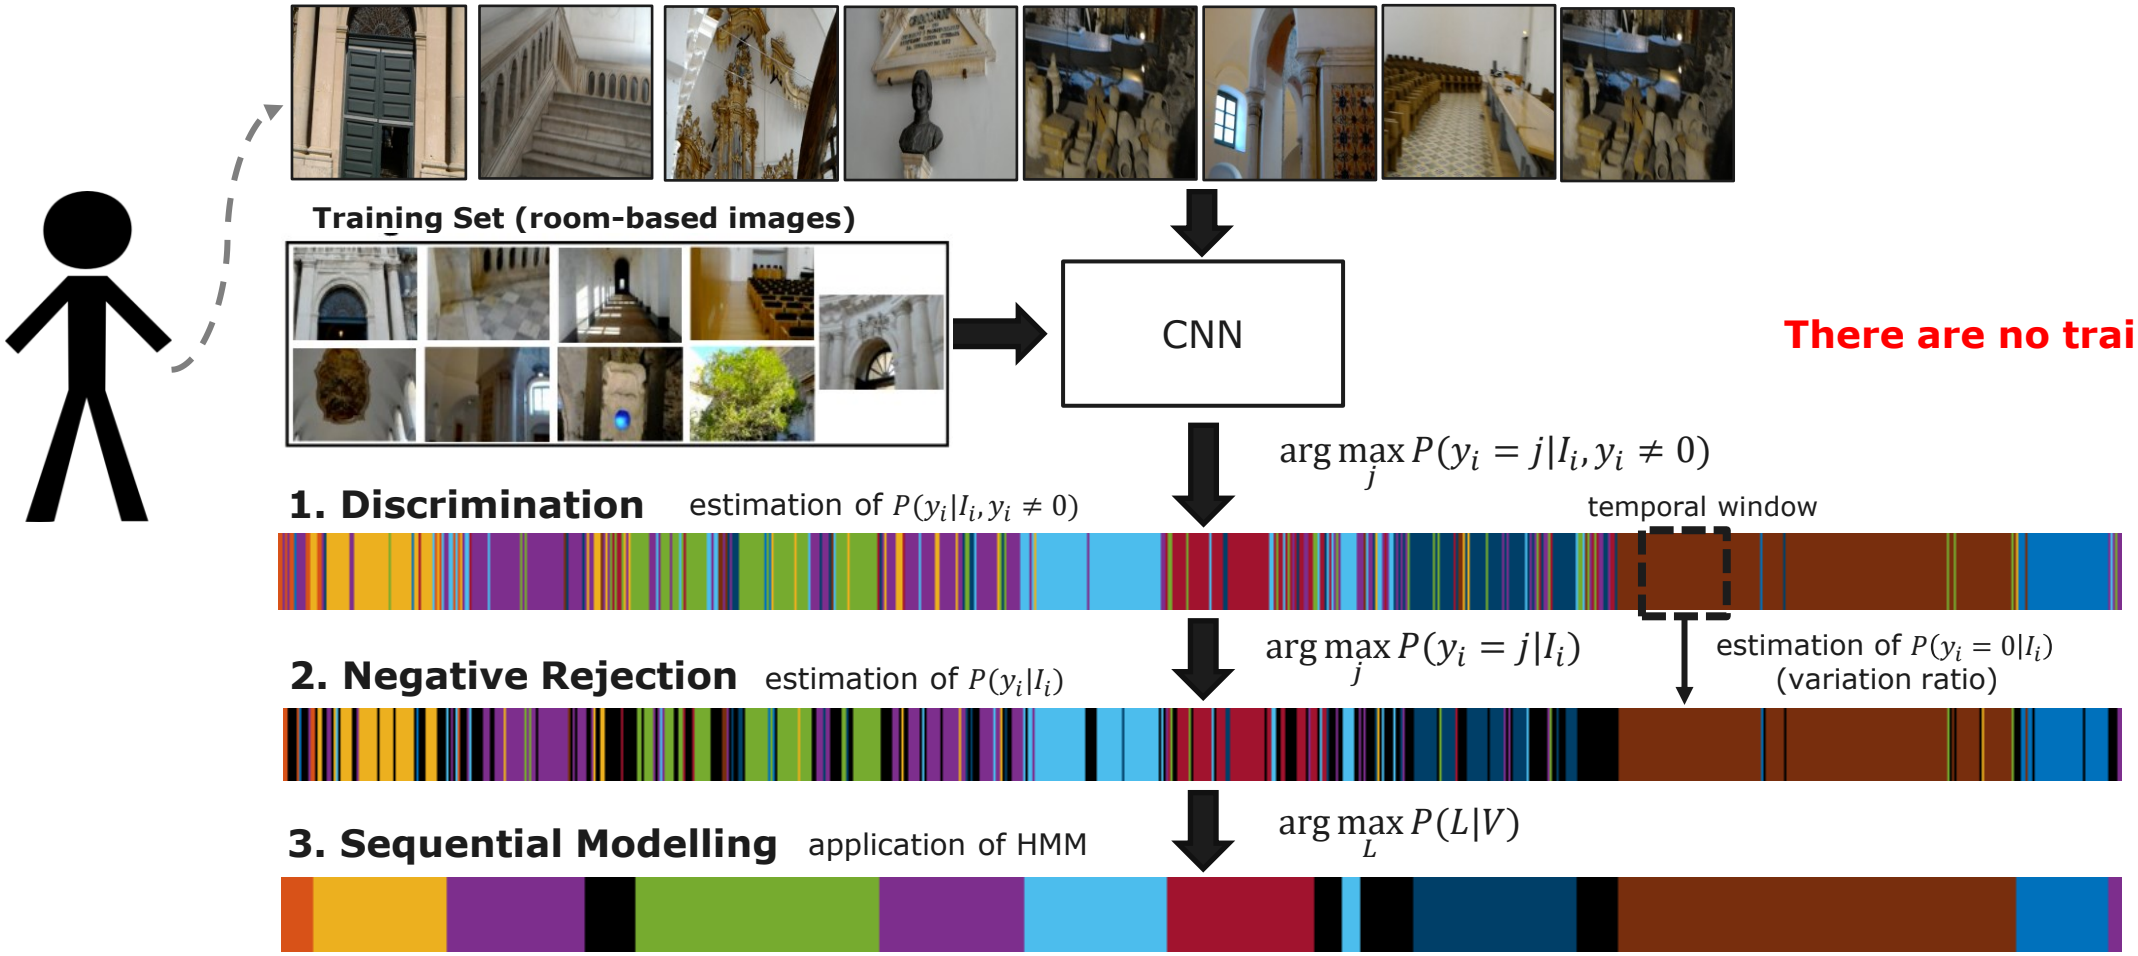

A. Furnari, G. M. Farinella, S. Battiato, Personal-Location-Based Temporal Segmentation of Egocentric Video for Lifelogging Applications, Journal of Visual Communication and Image Representation, 2017.

F. Ragusa, A. Furnari, S. Battiato, G. Signorello, G. M. Farinella. Egocentric Visitors Localization in Cultural Sites. In Journal on Computing and Cultural Heritage (JOCCH), 2019.

Room-Level Localization – Demo

Cutgana
CENTRO UNIVERSITARIO PER LA TUTELA E LA GESTIONE DEGLI AMBIENTI NATURALI E DEGLI AGRO-ECOSISTEMI

VEDI – Vision Exploitation for Data Interpration, PON MISE Horizon 2020

F. Ragusa, A. Furnari, S. Battiato, G. Signorello, G. M. Farinella

Time Spent at Location

LOC	EST	GT
G. Novizi	00:00	00:00
Cortile	00:03	00:03
Scalone	00:00	00:00
Corridoi	00:00	00:00
C. Notte	00:00	00:00
Antiref.	00:00	00:00
S. Mazz.	00:00	00:00
Cucina	00:00	00:00
Ventre	00:00	00:00
Negative	00:00	00:00

Prop.

GT

Detected Shots for Storyboard Summary

Estimated Probabilities	Predicted Class	GT Class
Giardino dei Novizi		
Cortile	●	●
Scalone Monumentale		
Corridoi		
Coro di Notte		
Antirefettorio		
Aula Santo Mazzarino		
Cucina		
Ventre		
Negative		

Localization-based analysis

G. M. Farinella, G. Signorello, S. Battiato, A. Furnari, F. Ragusa, R. Leonardi, E. Ragusa, E. Scuderi, A. Lopes, L. Santo, M. Samarotto. VEDI: Vision Exploitation for Data Interpretation. In 20th International Conference on Image Analysis and Processing (ICIAP), 2019

Localization – Levels of Granularity

SCENE RECOGNITION

off-the-shelf detectors

CAMERA POSE-ESTIMATION

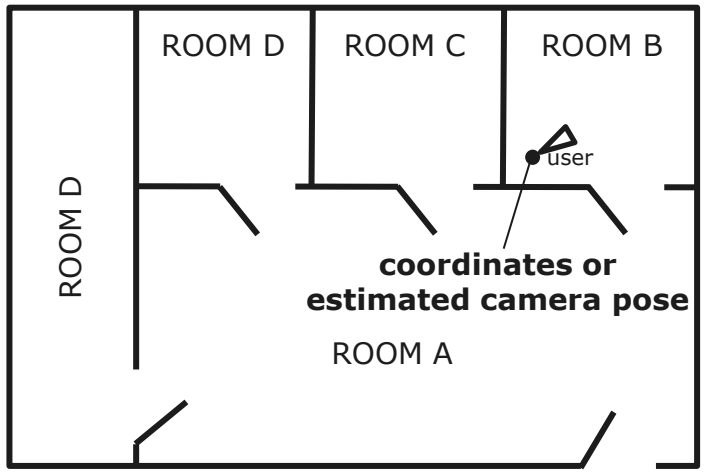

3D reconstruction of the building

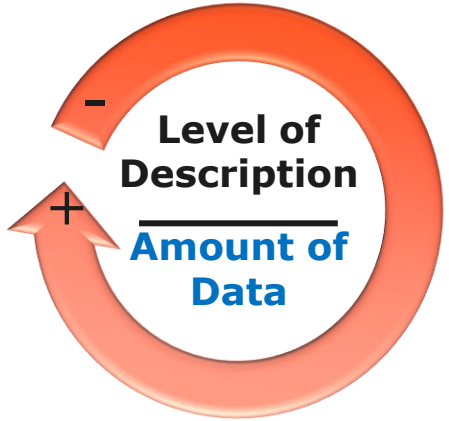

ROOM-LEVEL RECOGNITION

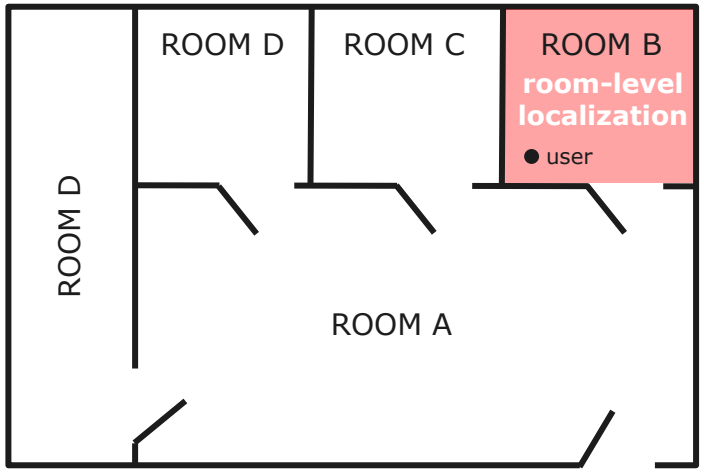

moderate amount of training data

Camera Pose Estimation – Dataset Creation

Images

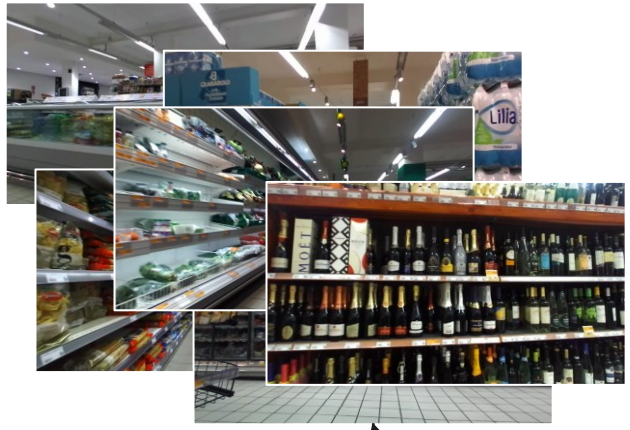

Structure from Motion (SfM)

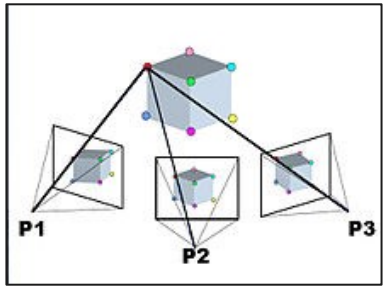

(P,Q)

3D Model

Attach estimated 6DOF pose to each image

Arbitrary Coordinate System (pose/scale)

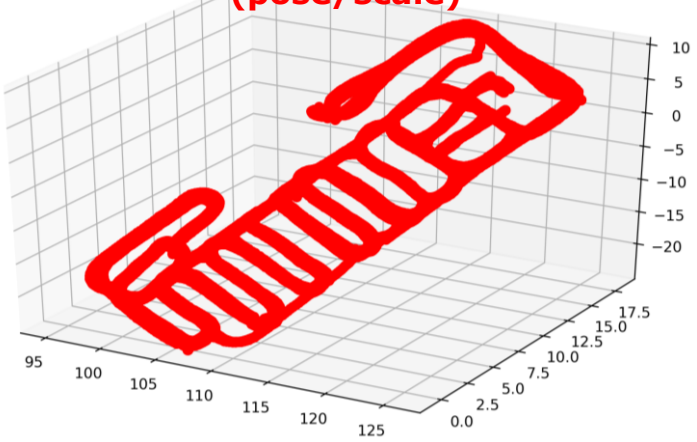

camera poses

PCA

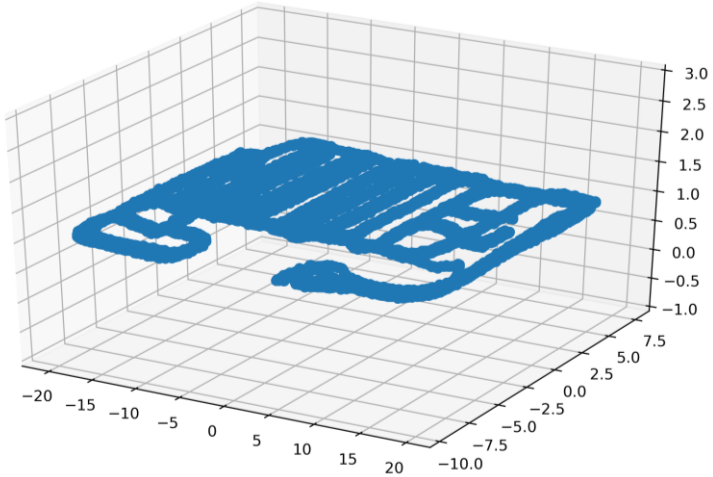

rotated poses

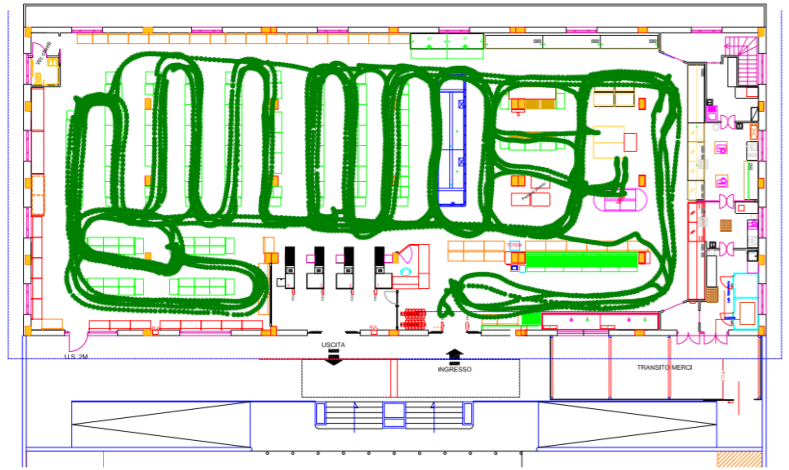

scaled/aligned poses

Camera Pose Estimation – Labeling

Structure from Motion attaches every input image to a 3D model.

Many options available:

COLMAP (free)

https://colmap.github.io/

Visual SFM (free)

http://ccwu.me/vsfm/

3D Zephir (paid)

https://www.3dflow.net/it/3df-zephyr-pro-3d-models-from-photos/

Camera Pose Estimation – Retrieval Approach

Use deep metric learning to learn a representation function φ which maps close to each other images of nearby locations

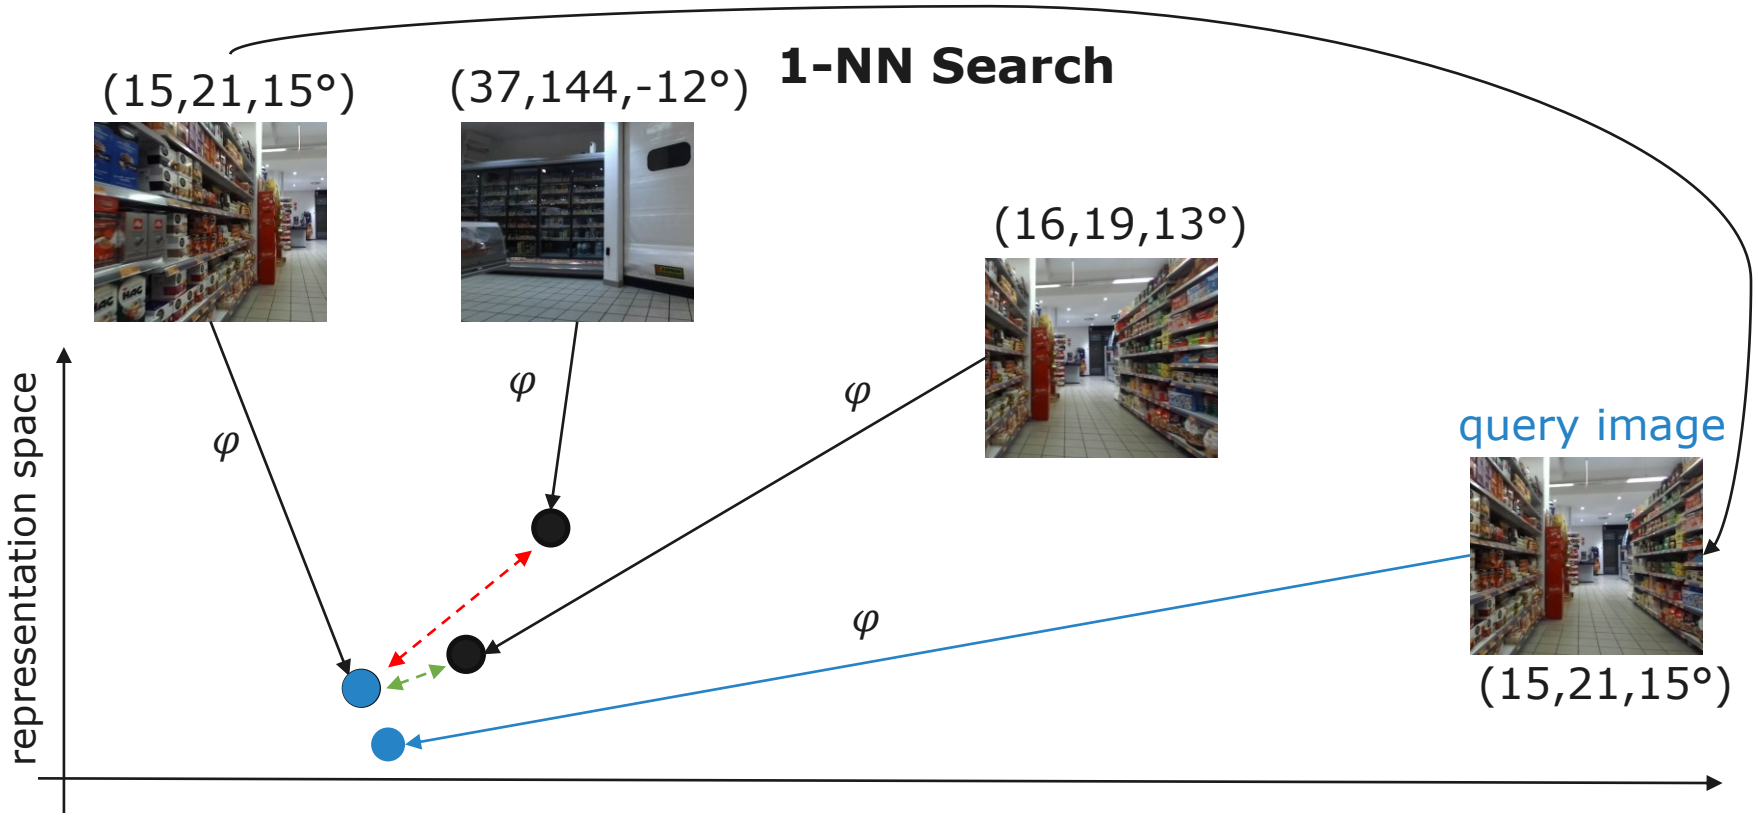

E. Spera, A. Furnari, S. Battiato, G. M. Farinella, Egocentric Shopping Cart Localization, International Conference on Pattern Recognition (ICPR), 2018
S. A. Orlando, A. Furnari, S. Battiato, G. M. Farinella. Image-Based Localization with Simulated Egocentric Navigations. VISAPP 2019

Visual Localization Literature is Rich!

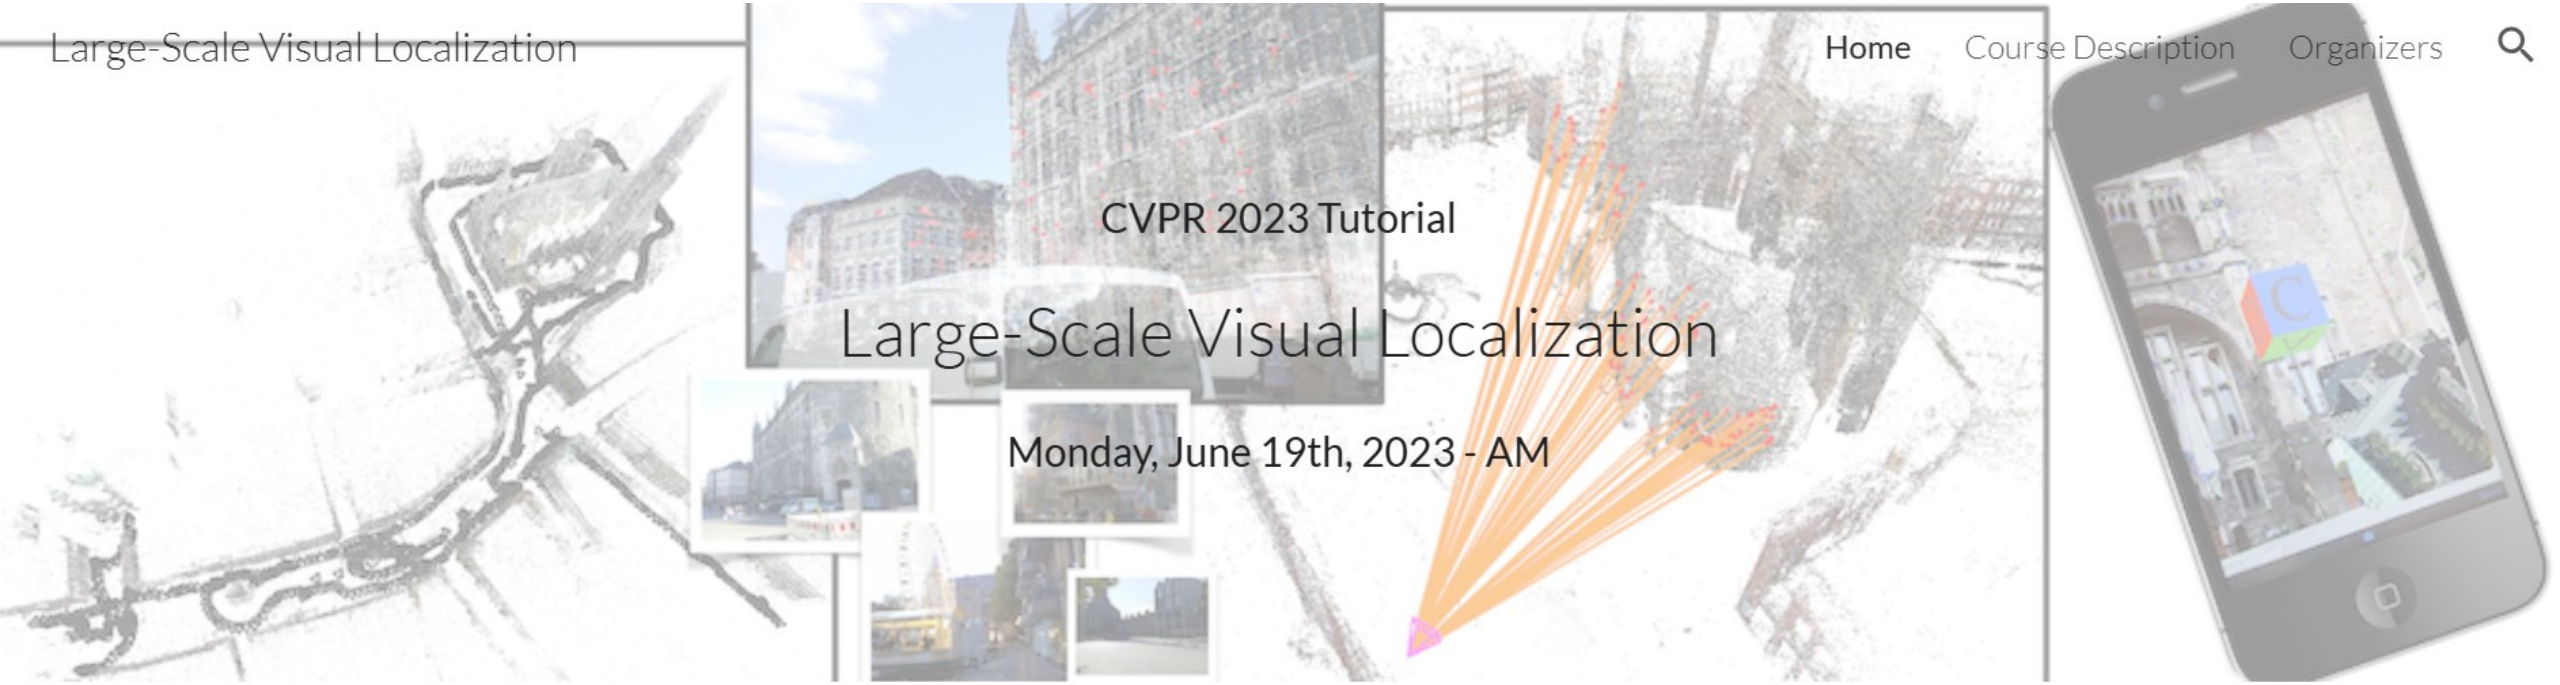

Course Information

- **When:** Monday, June 19th, 2023
- **Where:** East 2
- **Time:** 8:30 AM - 12:15 PM (local time)
- **Schedule**
  - **Introduction** [5 min] [8:30 - 8:35]
  - **Part I: Instance retrieval for coarse localization** [45 min](Giorgos) [8:35 - 9:20] [slides]

https://sites.google.com/view/lsvl2023/home

Object Detection/Interaction

Object Detection

Objects and Actions are tight!

Useful to know what is in the scene

Useful to know what actions can be performed

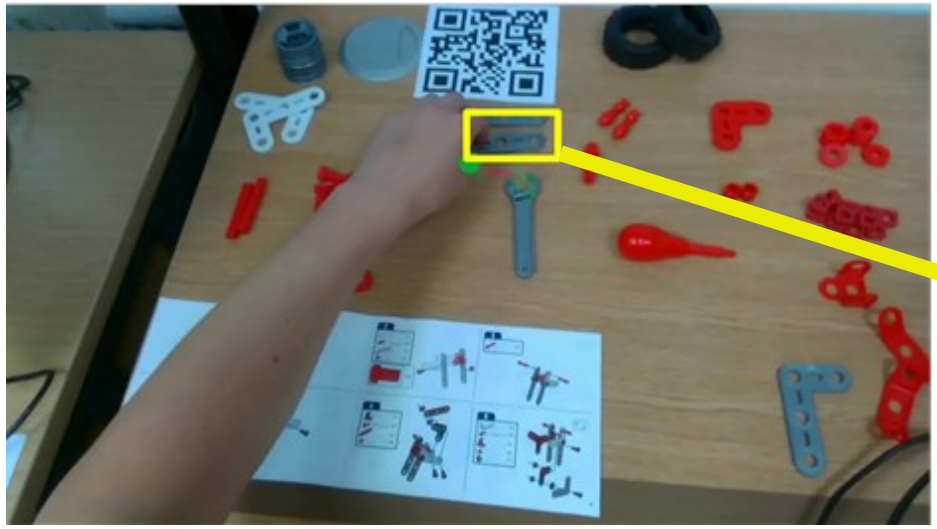

ID	Class
0	instruction booklet
1	gray_angled_perforated_bar
2	partial_model
3	white_angled_perforated_bar
4	wrench
5	screwdriver
6	gray_perforated_bar
7	wheels_axle
8	red_angled_perforated_bar
9	red_perforated_bar
10	rod
11	handlebar
12	screw
13	tire
14	rim
15	washer
16	red_perforated_junction_bar
17	red_4_perforated_junction_bar
18	bolt
19	roller

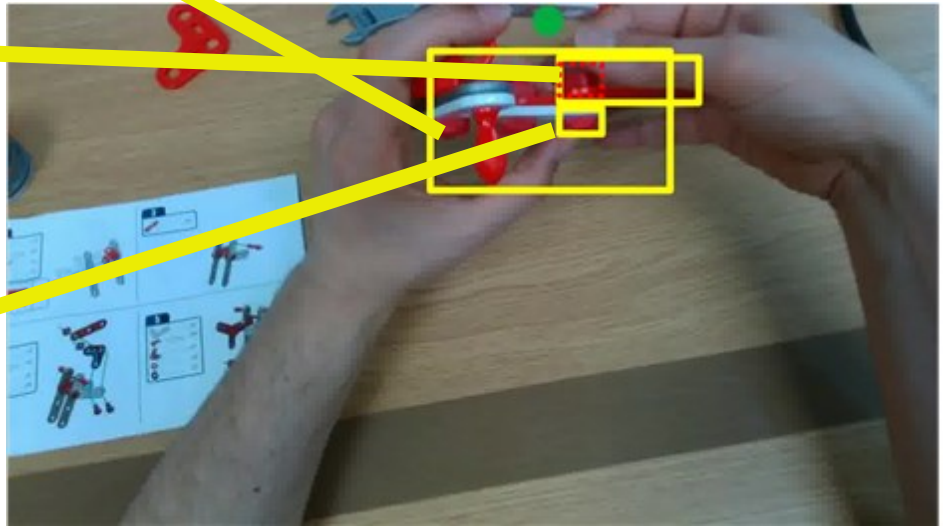

F. Ragusa, A. Furnari, G. M. Farinella. MECCANO: A Multimodal Egocentric Dataset for Humans Behavior Understanding in the Industrial-like Domain. Computer Vision and Image Understanding (CVIU), 2023 (https://arxiv.org/abs/2209.08691).

Off-the-shelf object detectors

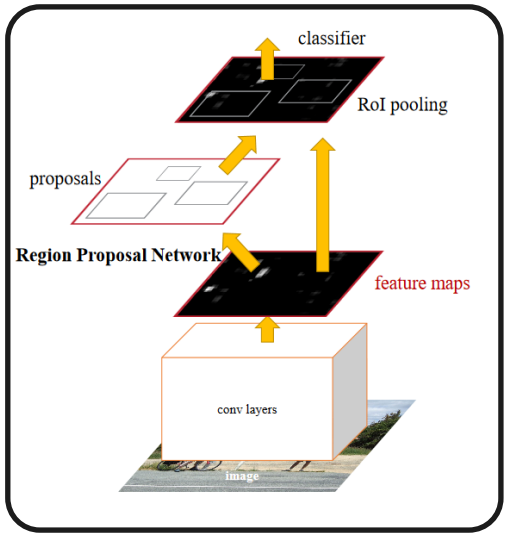

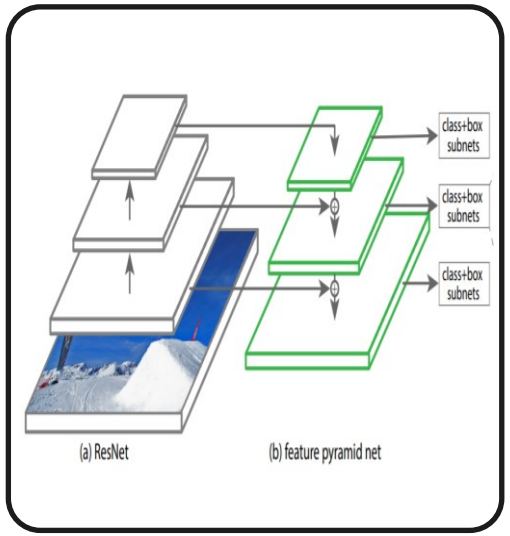

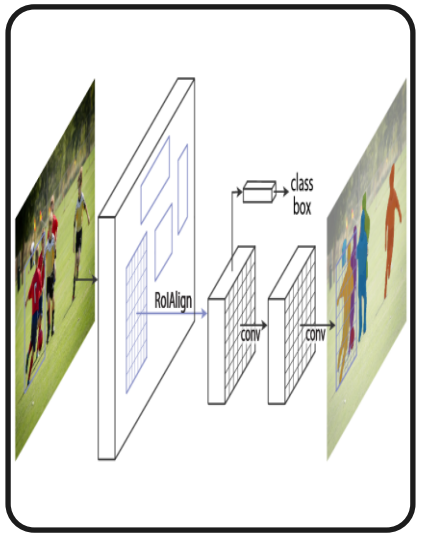

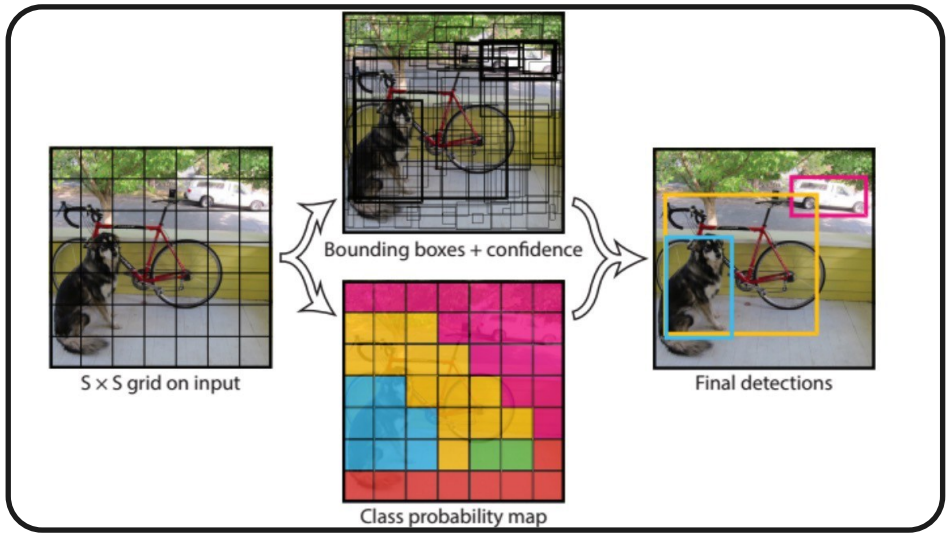

Faster-RCNN
(bounding boxes)

RetinaNet
(bounding boxes - faster)

Mask-RCNN
(boxes + segments)

YOLO
(much faster, but less accurate)

https://github.com/facebookresearch/detectron2

https://pjreddie.com/darknet/yolo/

Transformer-Based Detectors: https://github.com/IDEA-Research/awesome-detection-transformer

Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*.
Joseph Redmon, Ali Farhadi, YOLO9000: Better, Faster, Stronger, The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017
He, K., Gkioxari, G., Dollár, P., & Girshick, R. (2017, October). Mask r-cnn. In *Computer Vision (ICCV), 2017* (pp. 2980-2988). IEEE.

Off-the-shelf detectors on Ego-Objects

Depending on the scenario, off-the-shelf detectors can be a starting point, but they are not always accurate.

368 categories!

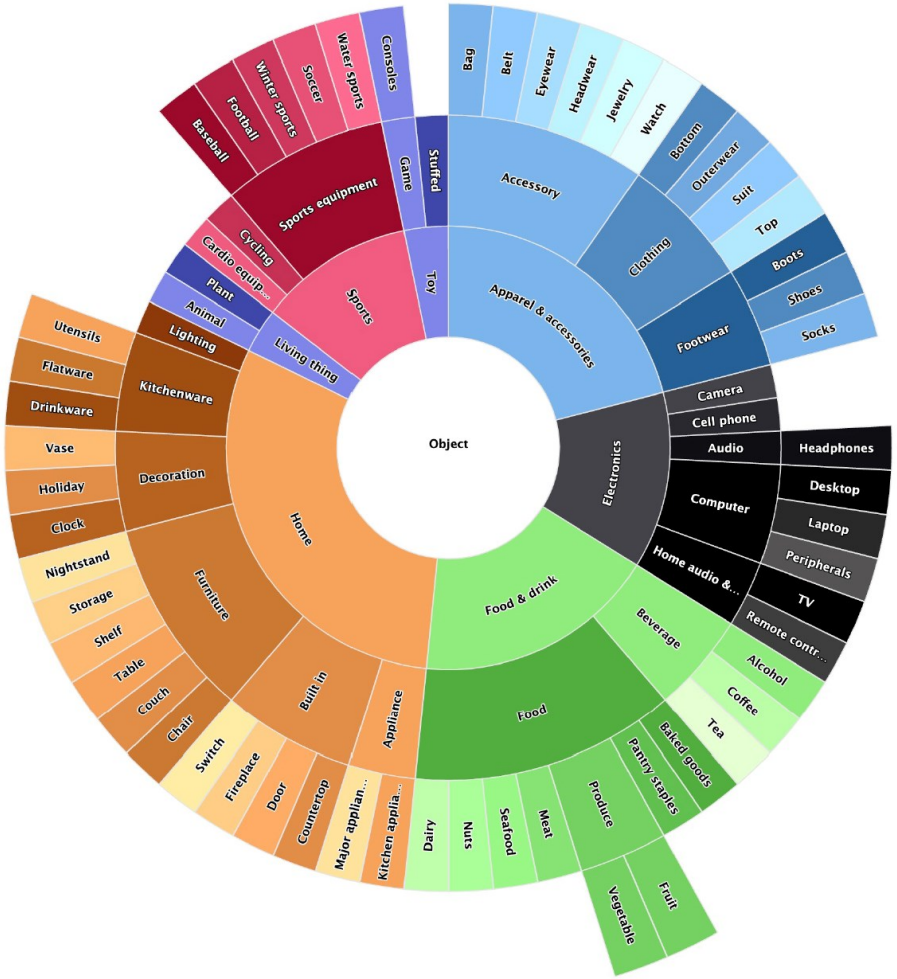

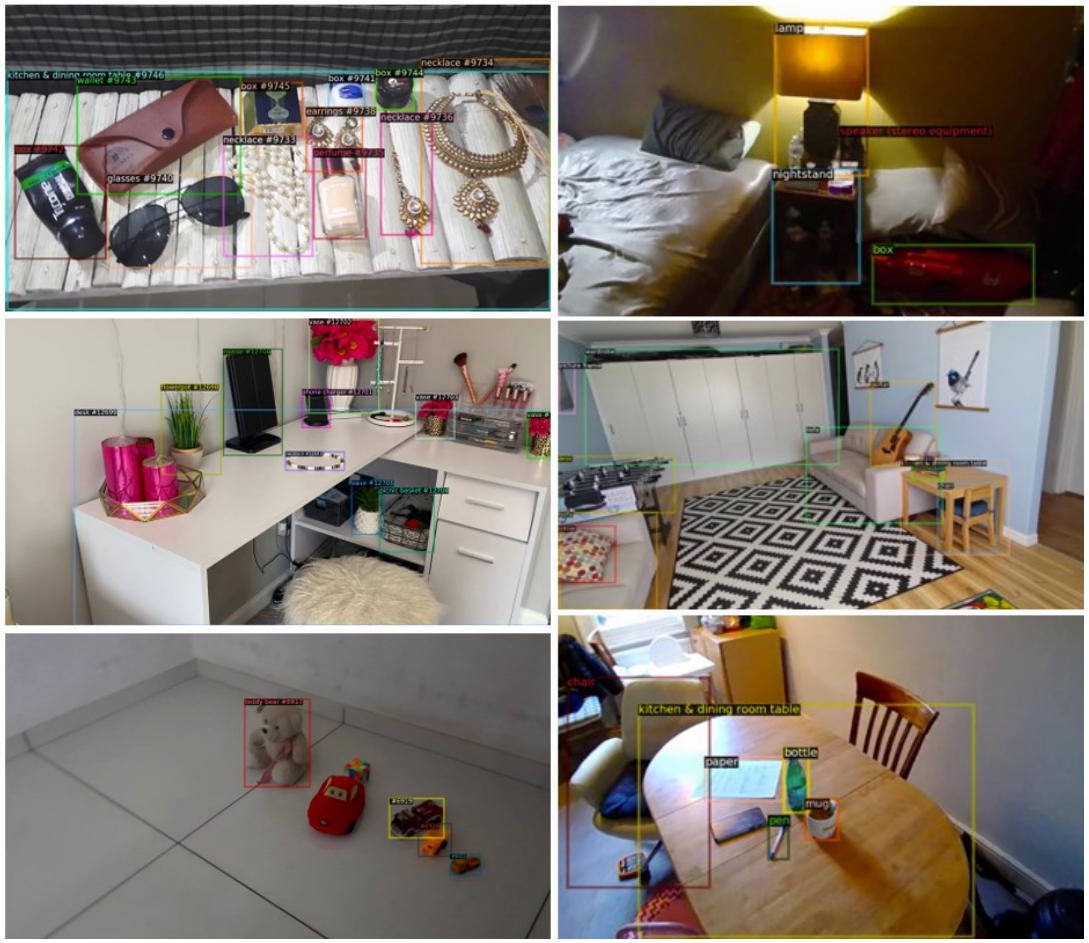

EgoObjects

https://github.com/facebookresearch/EgoObjects

C. Zhu et al., "EgoObjects: A Large-Scale Egocentric Dataset for Fine-Grained Object Understanding“ in IEEE/CVF International Conference on Computer Vision (ICCV), 2023

Train/Finetune your own object detector

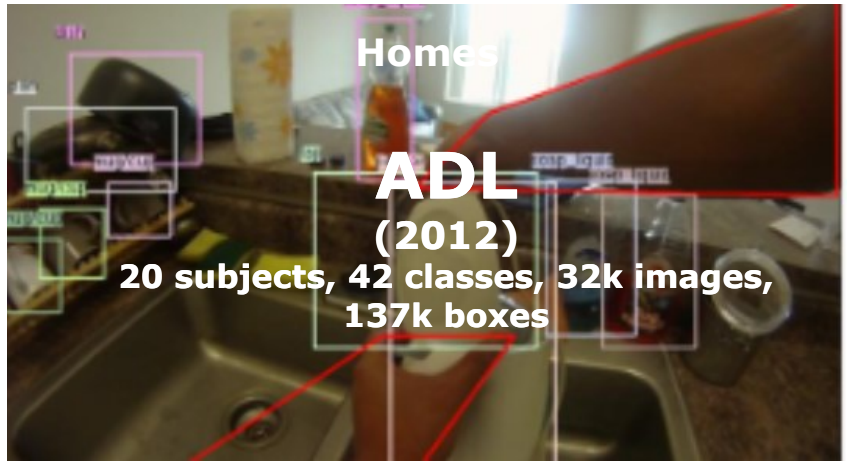

https://www.csee.umbc.edu/~hpirsiav/papers/ADLdataset/

http://epic-kitchens.github.io/

https://iplab.dmi.unict.it/EGO-CH/

https://iplab.dmi.unict.it/MECCANO/

- In some scenarios, it could be necessary to fine-tune an object-detector with application-specific data.
- Main egocentric datasets providing bounding box annotations.
- EGO4D is multi- domain annotated with 295K bounding boxes.

Understanding human-object interactions (HOI)

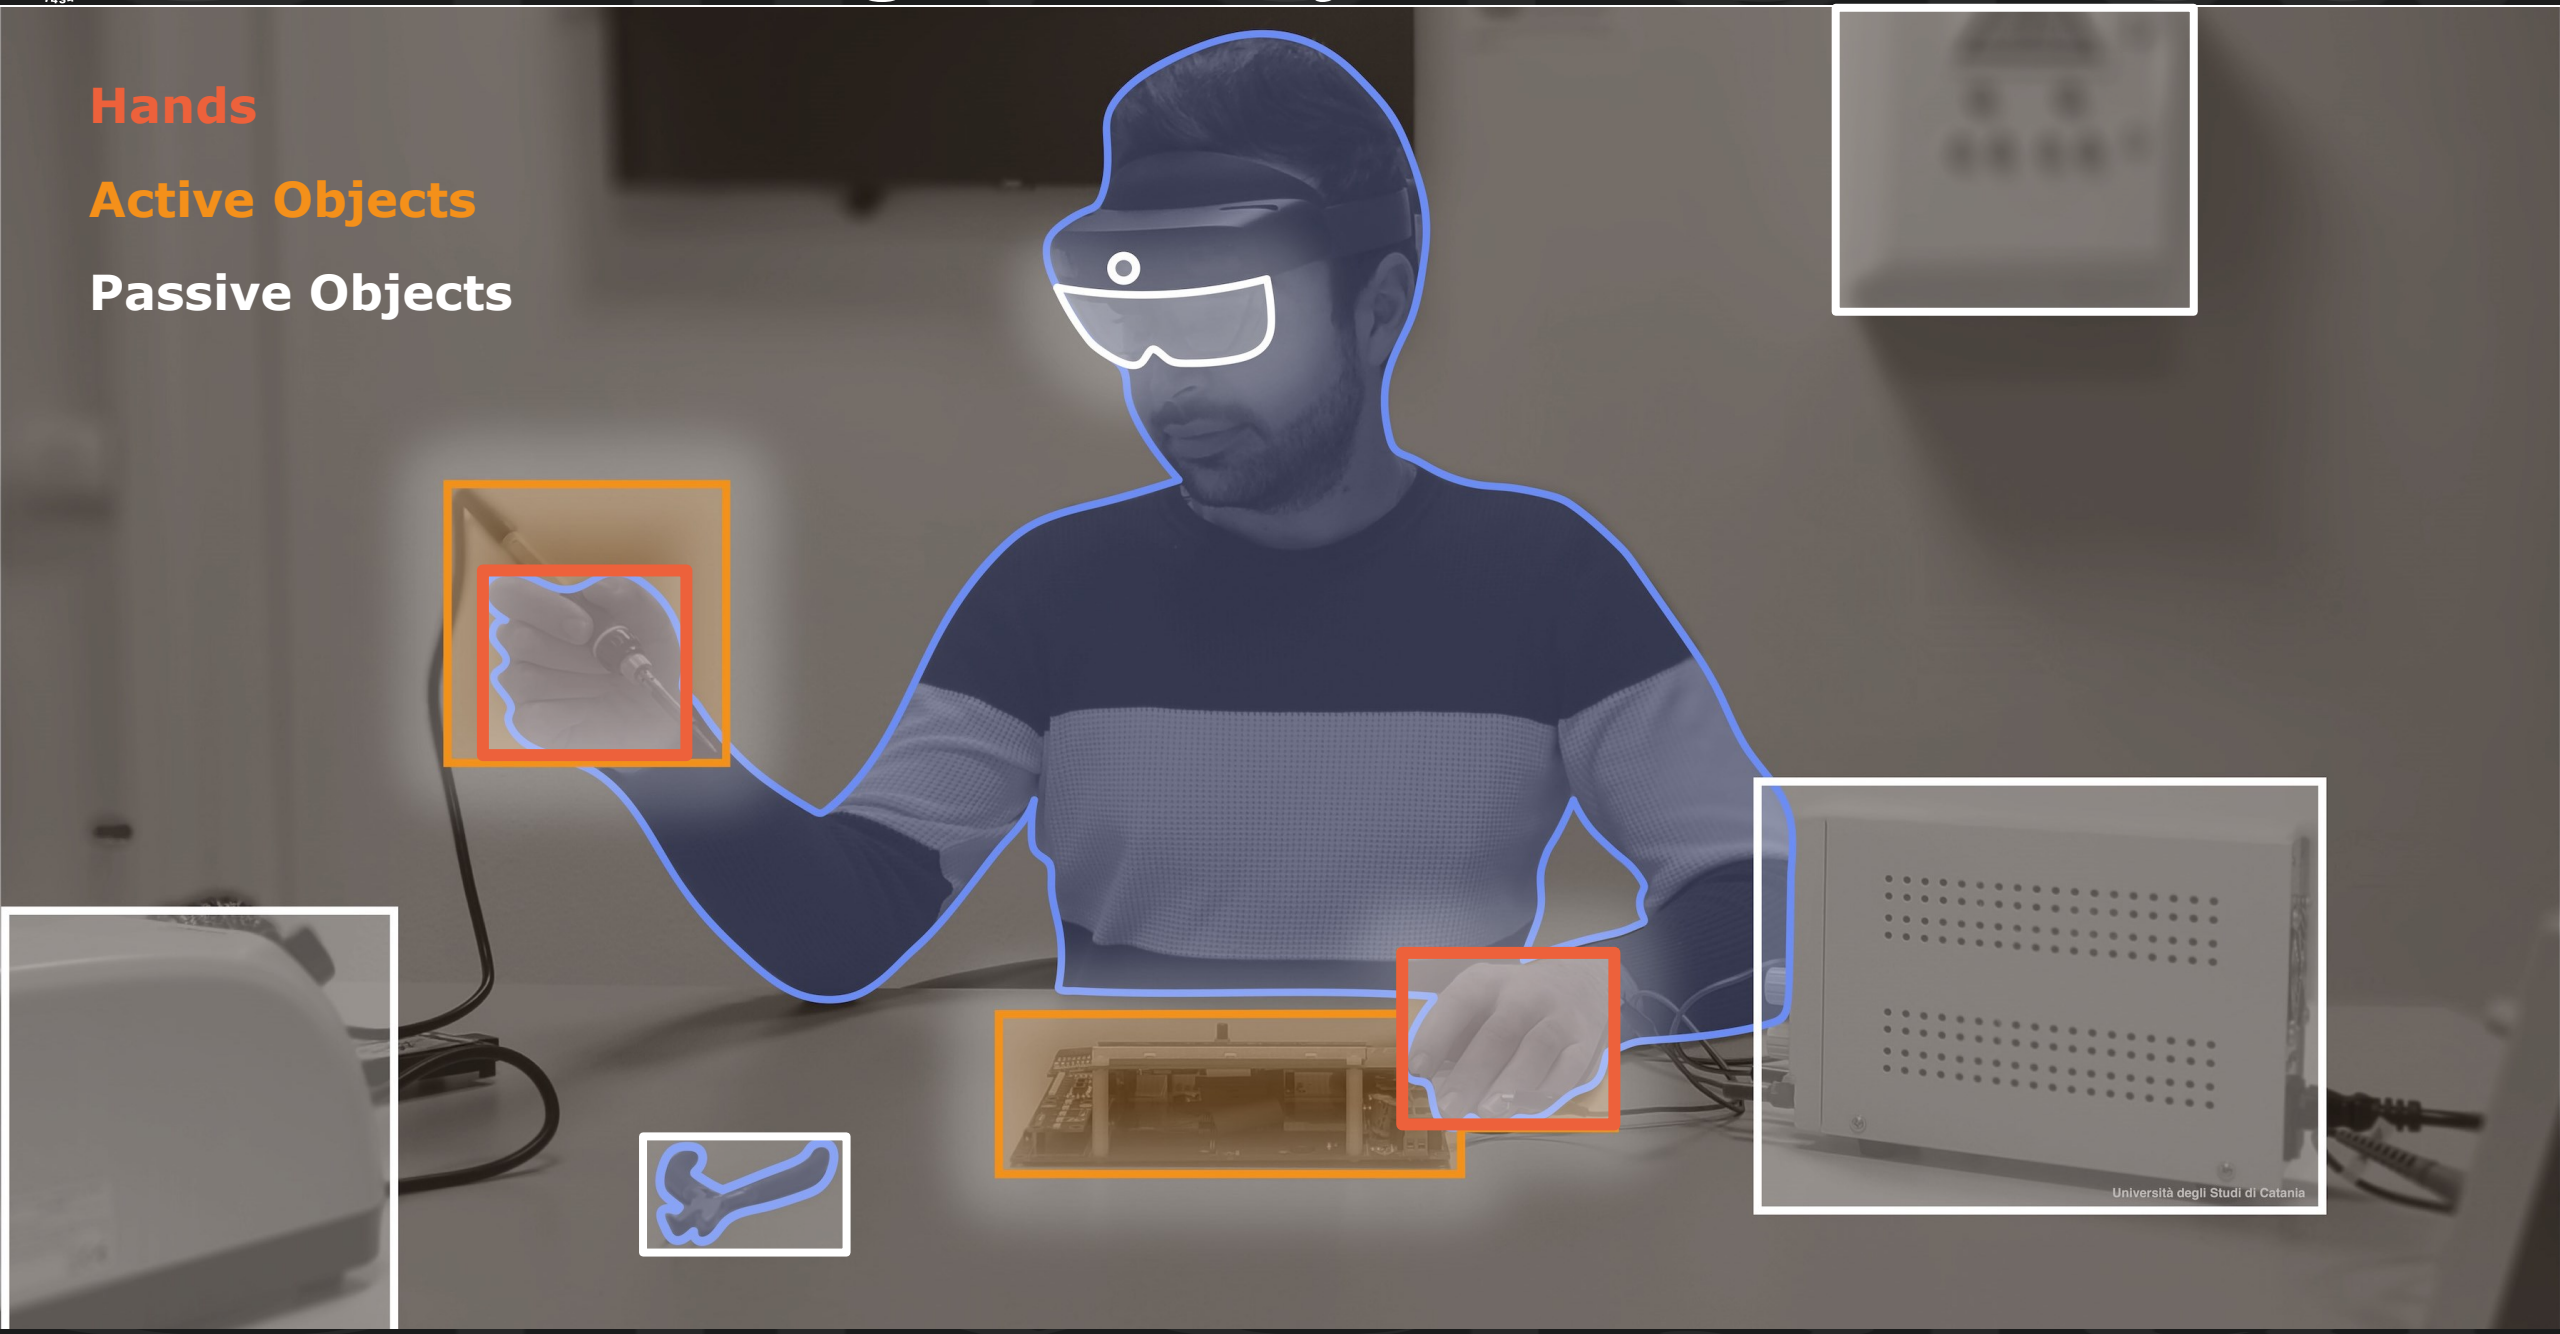

Hands in Contact – Hands + Objects

CODE & DATA HERE -> https://fouheylab.eecs.umich.edu/~dandans/projects/100DOH/

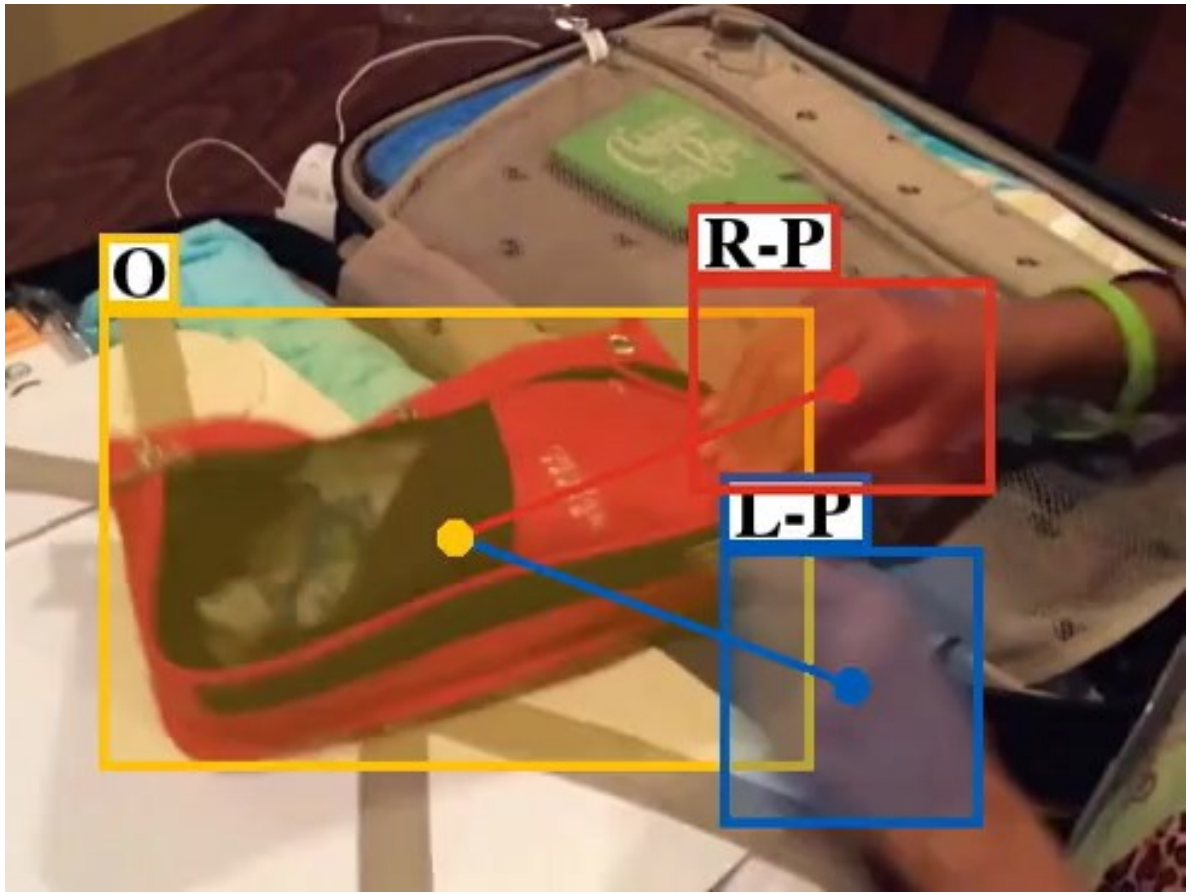

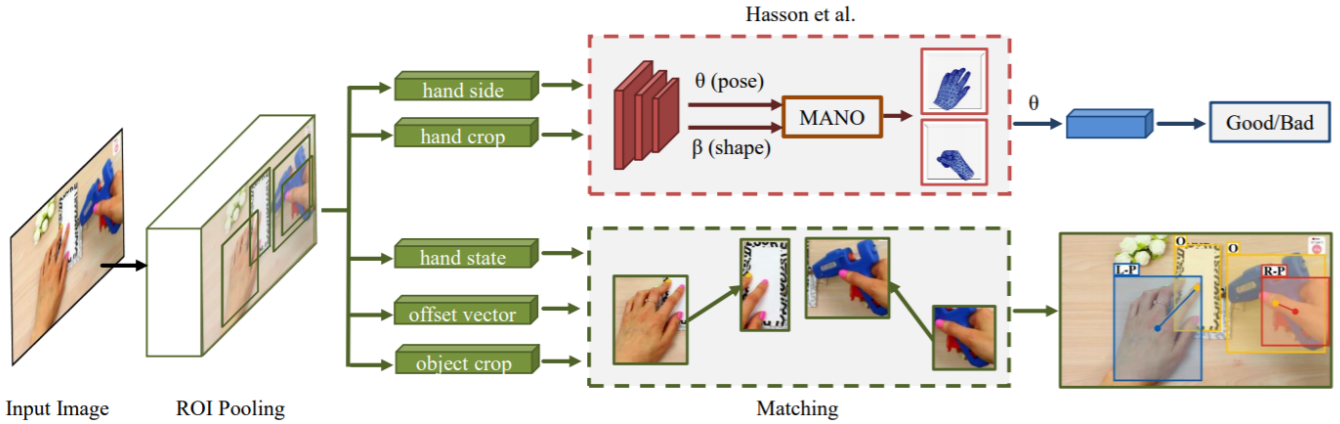

An «augmented» detector which recognizes:

- The left hand;
- The right hand;
- The interacted object.

VISOR DATASET

Darkhalil, Ahmad, et al. "Epic-kitchens visor benchmark: Video segmentations and object relations." *Advances in Neural Information Processing Systems* 35 (2022): 13745-13758.

Shan, D., Geng, J., Shu, M., & Fouhey, D. F. (2020). Understanding human hands in contact at internet scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 9869-9878).

Hands23 – Hands + Tool + Object

CODE & DATA HERE -> https://fouheylab.eecs.umich.edu/~dandans/projects/hands23/

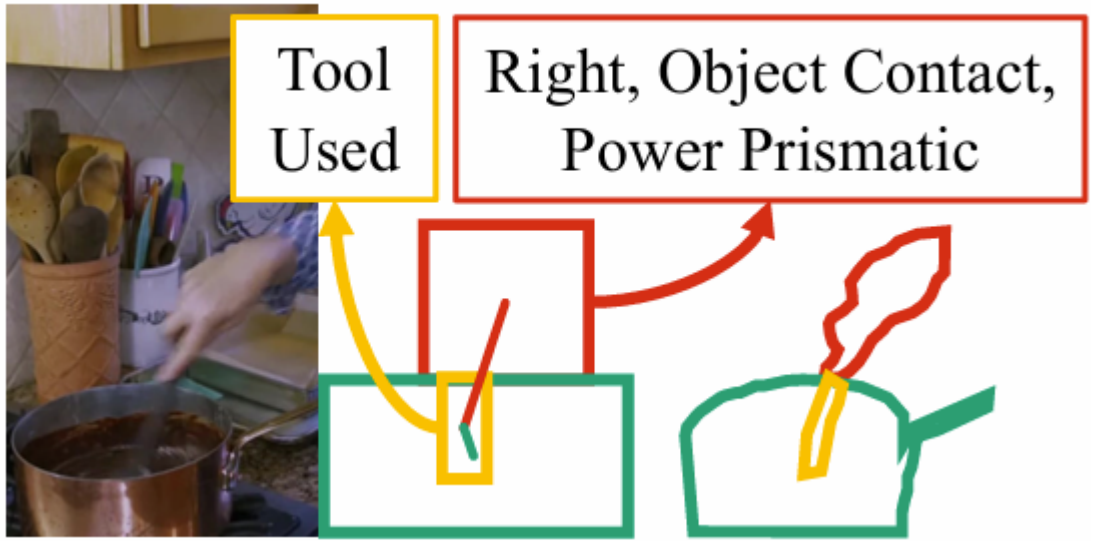

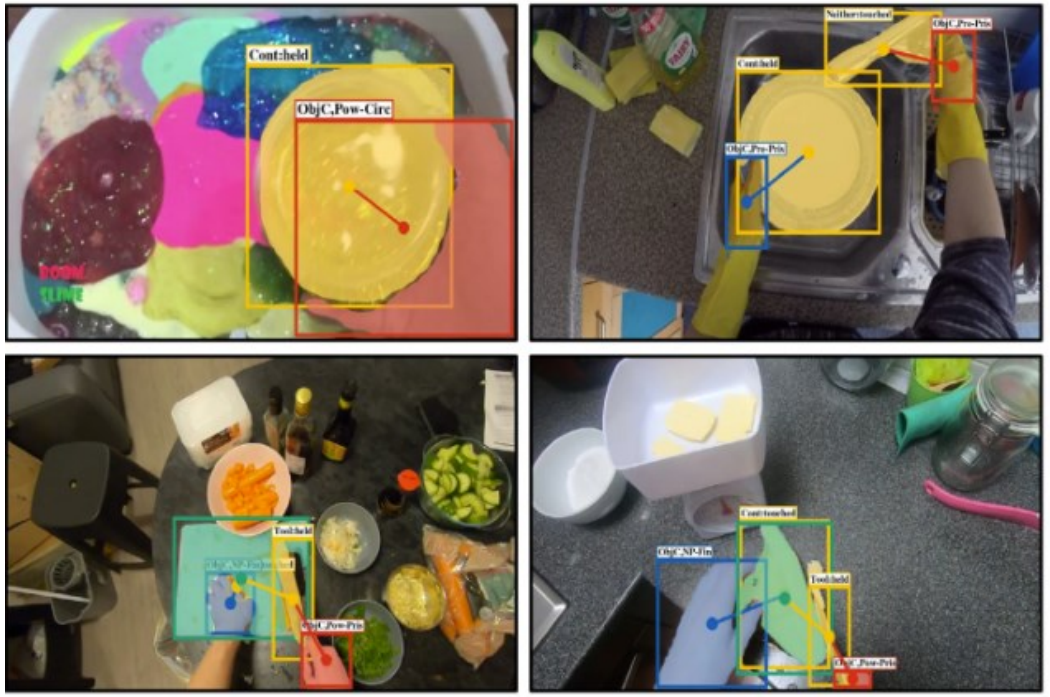

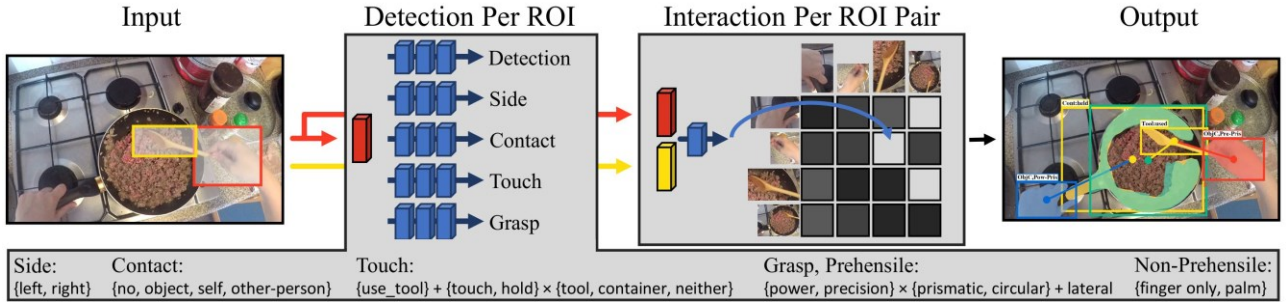

An «augmented» detector which recognizes:

- The left hand;
- The right hand;
- The interacted tool;
- The interacted object.

T. Cheng, D. Shan, A. Sultan, R. E. L. Higgins, and D. F. Fouhey. 2023. Towards a richer 2D understanding of hands at scale. In Proceedings of the 37th International Conference on Neural Information Processing Systems (NIPS '23).

Standard approach – relies on huge quantities of real data

Standard approach:

- Collect a lot of images and videos of construction sites;
- Label the data with domain-specific annotations;
- Train and test deep learning algorithms.

Learning in simulation

What if we could learn the «real thing» in simulation?

DATA HERE -> https://iplab.dmi.unict.it/HOI-Synth/

Are Synthetic Data Useful for Egocentric Hand-Object Interaction Detection?

Rosario Leonardi, Antonino Furnari, Francesco Ragusa, Giovanni Maria Farinella (2024). Are Synthetic Data Useful for Egocentric Hand-Object Interaction Detection?. In European Conference on Computer Vision (ECCV).

HOI-Synth benchmark

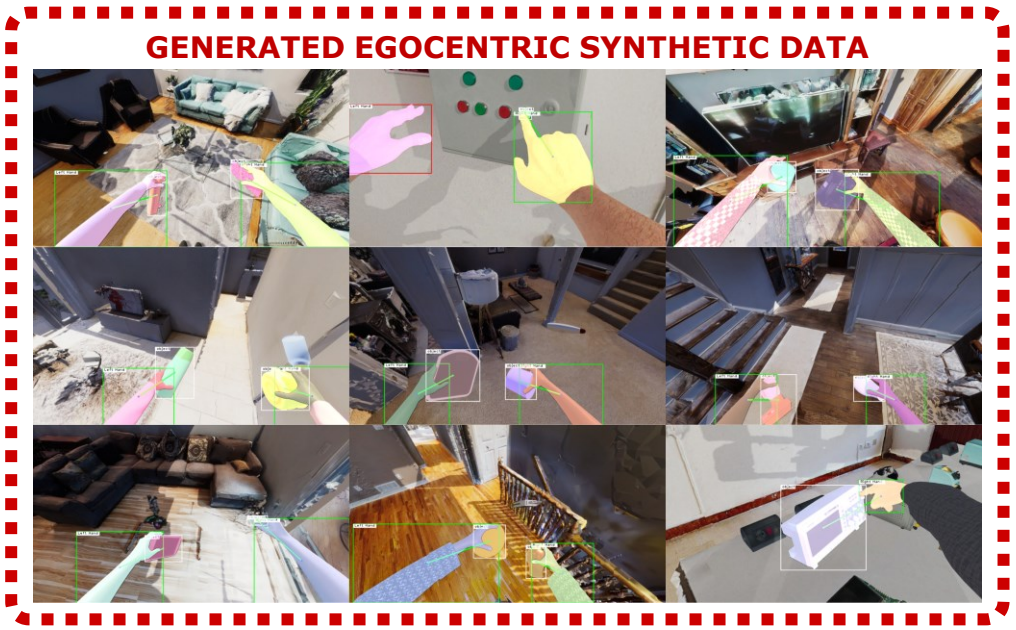

- **Epic-Kitchens VISOR**
  - 32,857 real + 30,259 synthetic images
- **EgoHOS**
  - 8,107 real + 8,107 synthetic images
- **ENIGMA-51**
  - 3,479 real images
  - In-Domain 16,773 + Out-domain 20,321 synthetic images

Rosario Leonardi, Antonino Furnari, Francesco Ragusa, Giovanni Maria Farinella (2024). Are Synthetic Data Useful for Egocentric Hand-Object Interaction Detection?. In European Conference on Computer Vision (ECCV).

Quantitative Results: Epic-Kitchens VISOR

a) Unsupervised Setting

% Real Labeled Data	Approach	Overall	H	H+S	H+C	O
0%	Synthetic-Only	09.88	28.41	24.89	08.64	01.23
	UDA	**33.33**	**80.16**	**65.98**	**33.47**	**08.35**
Absolute Improvement		**+23.45**	+51.75	+41.09	+24.83	+7.12

b) Semi-supervised Setting

% Real Labeled Data	Approach	Overall	H	H+S	H+C	O
25% (8,215 images)	Real-Only	37.90	90.14	**85.66**	53.99	17.85
	Synthetic+Real	38.19	89.98	84.67	**55.88**	18.49
	SSDA	**45.55**	**90.37**	84.42	52.59	**22.15**
Absolute Improvement		**+7.65**	+0.23	−0.99	+1.89	+4.30

C) Fully-supervised Setting

% Real Labeled Data	Approach	Overall	H	H+S	H+C	O
100% (32,857 images)	Real-Only	45.33	**92.25**	88.54	**59.24**	24.23
	Synthetic+Real	44.52	91.45	**88.94**	56.55	**27.77**
	FSDA	**46.48**	91.83	87.65	57.63	24.03
Absolute Improvement		**+1.15**	−0.42	+0.40	−1.61	+3.54

Rosario Leonardi, Antonino Furnari, Francesco Ragusa, Giovanni Maria Farinella (2024). Are Synthetic Data Useful for Egocentric Hand-Object Interaction Detection?. In European Conference on Computer Vision (ECCV).

Action Recognition

Action Recognition: Task

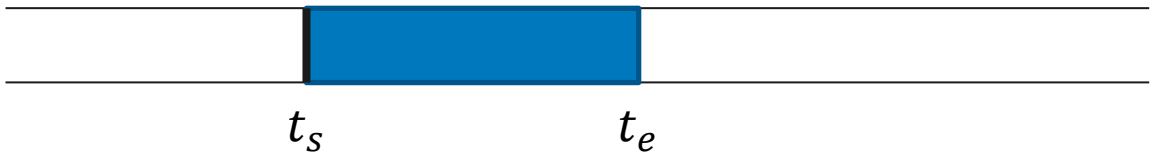

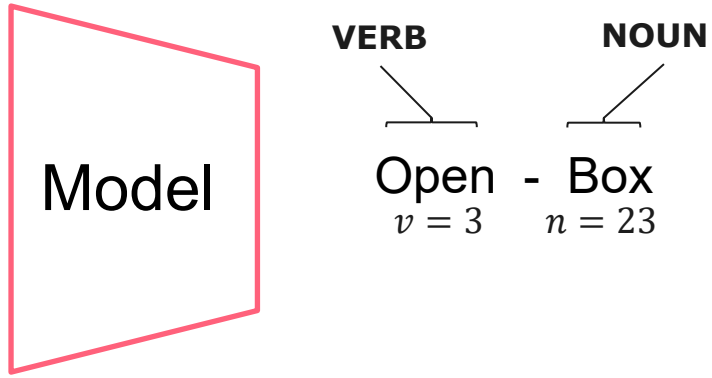

"observe a trimmed segment denoted by start and end time and classify the action present in the clip"

As defined in EPIC-KITCHENS-2020

Relation between Action and Interaction

TAKE SCREWDRIVER

F. Ragusa, A. Furnari, G. M. Farinella. MECCANO: A Multimodal Egocentric Dataset for Humans Behavior Understanding in the Industrial-like Domain. Computer Vision and Image Understanding (CVIU), 2023.

Relation between Action and Interaction

TAKE SCREWDRIVER

Start Action

Start Interaction (H-O)

Frame of Contact

Relation between Action and Interaction

TAKE SCREWDRIVER

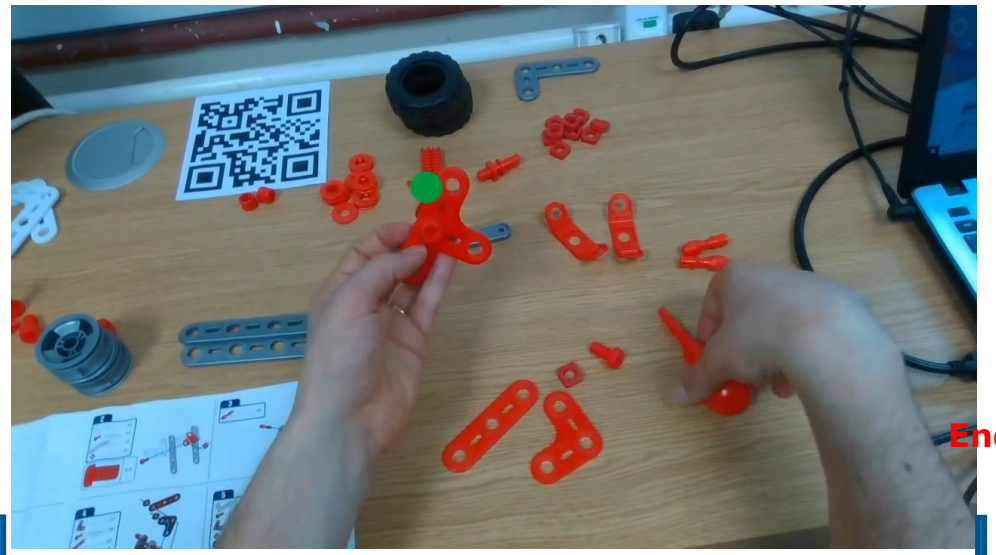

End Interaction

Start Action

Start Interaction (H-O)

End Action

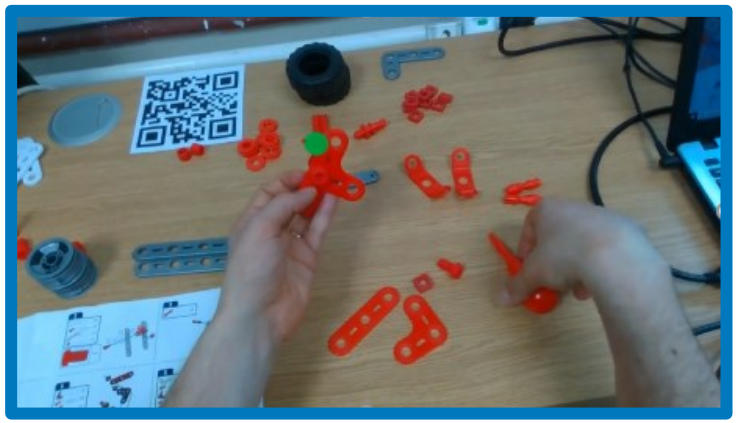

Frame of Contact

Frame of Decontact

Relation between Action and Interaction

F. Ragusa, A. Furnari, G. M. Farinella. MECCANO: A Multimodal Egocentric Dataset for Humans Behavior Understanding in the Industrial-like Domain. Computer Vision and Image Understanding (CVIU), 2023.

Relation	Verbs	MECCANO verbs
	pat, hit, kick	//
	pick up	take, fit, align, plug, pull
	close, open, turn on, press, push	browse
	walk, jump, run	//
	wring out, wash, cut, mix	pull
	throw, leave, place	put
	move	browse
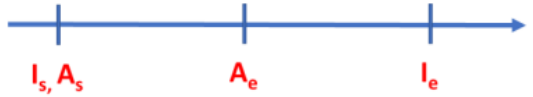	twist, rip	screw, unscrew, tighten, loosen
	stretch, knead, write, watch	check

SlowFast Networks for Video Recognition

CODE HERE -> https://github.com/facebookresearch/SlowFast

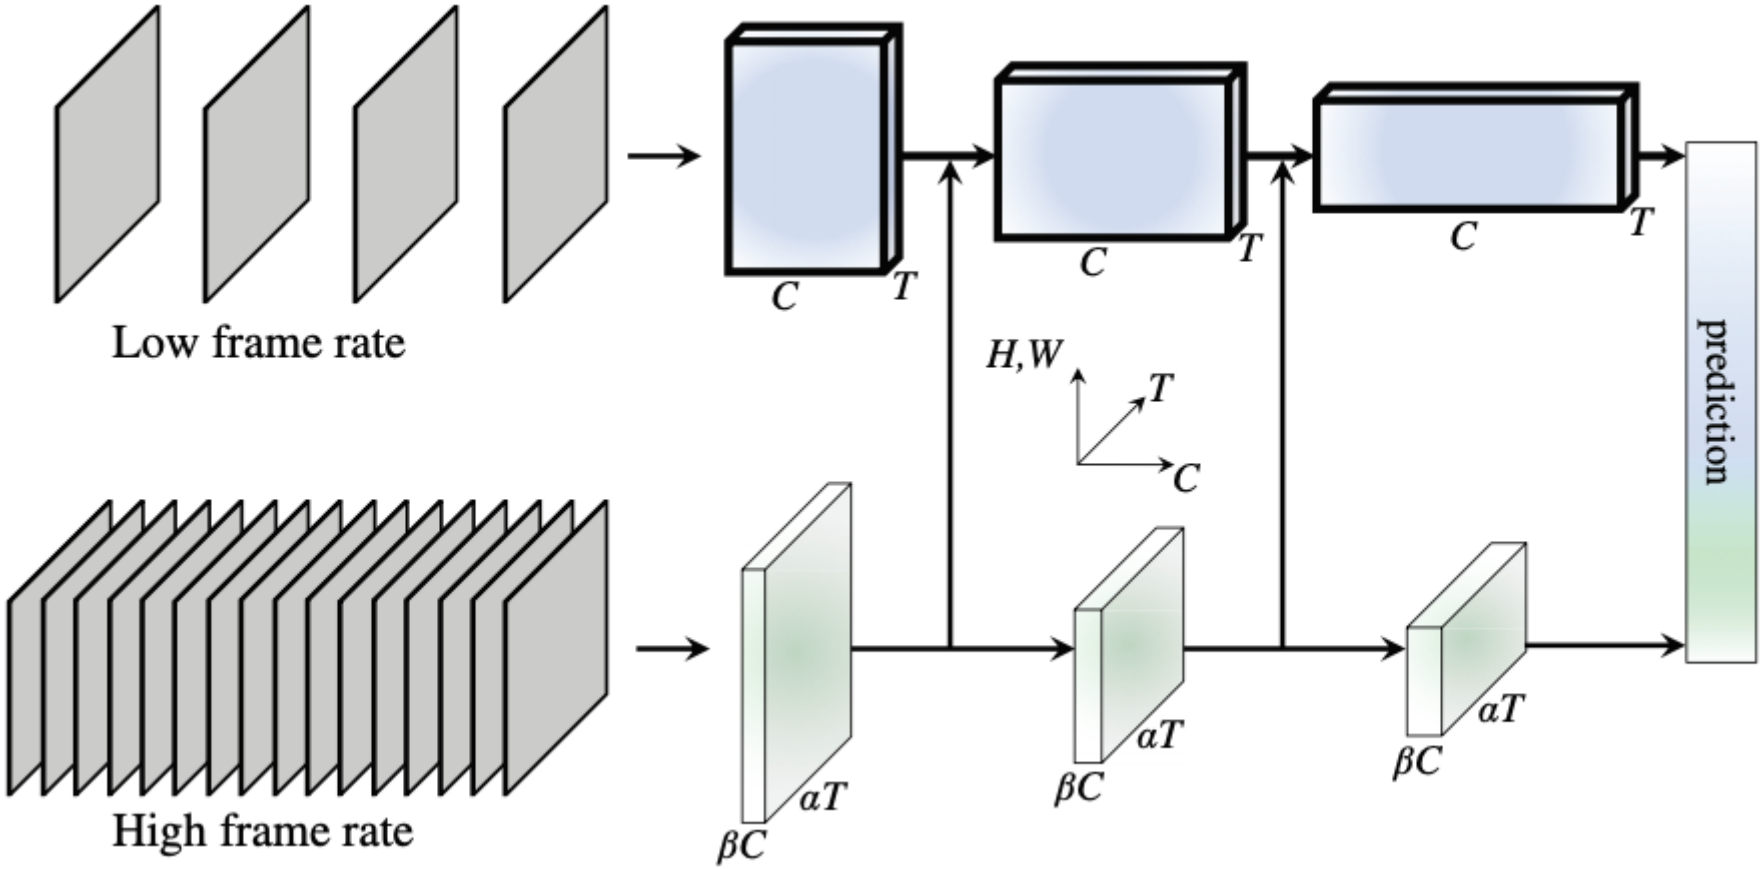

Feichtenhofer, C., Fan, H., Malik, J., & He, K. (2019). Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 6202-6211).

Vision Transformers for Video

CODE HERE -> https://github.com/facebookresearch/SlowFast

Factorized attention: Attend over space / time

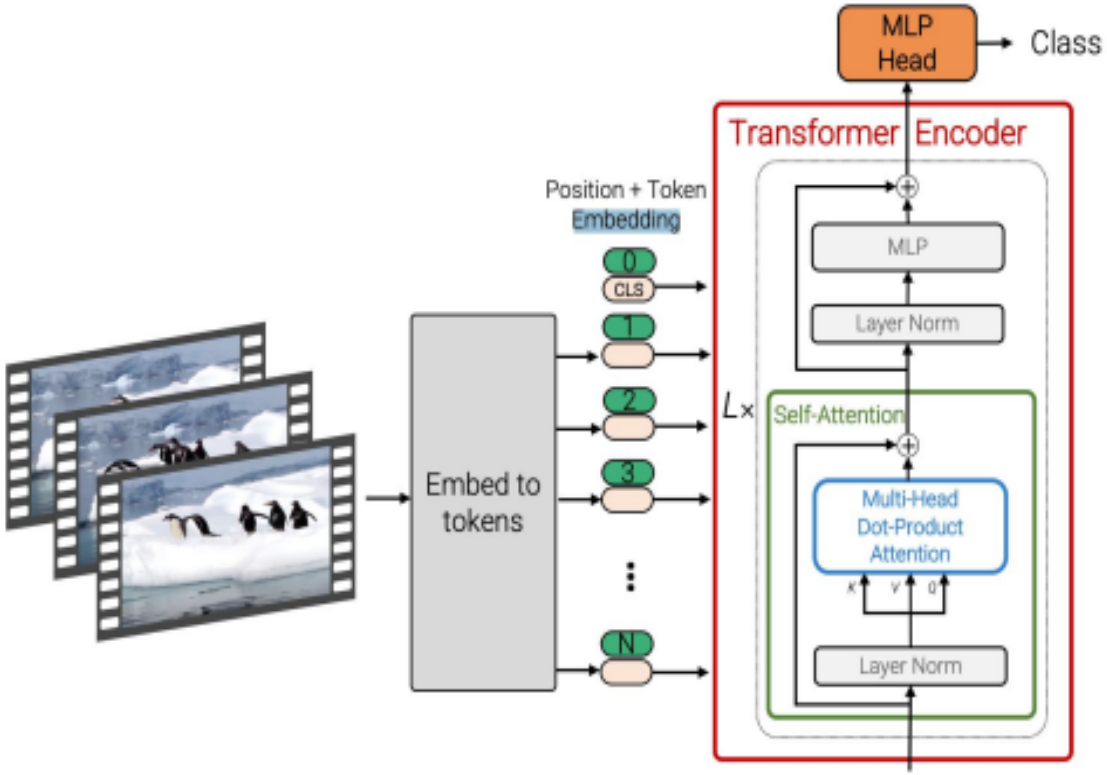

Bertasius et al, "Is Space-Time Attention All You Need for Video Understanding?", ICML 2021

Arnab et al, "ViViT: A Video Vision Transformer", ICCV 2021

Neimark et al, "Video Transformer Network", ICCV 2021

Pooling module: Reduce number of tokens

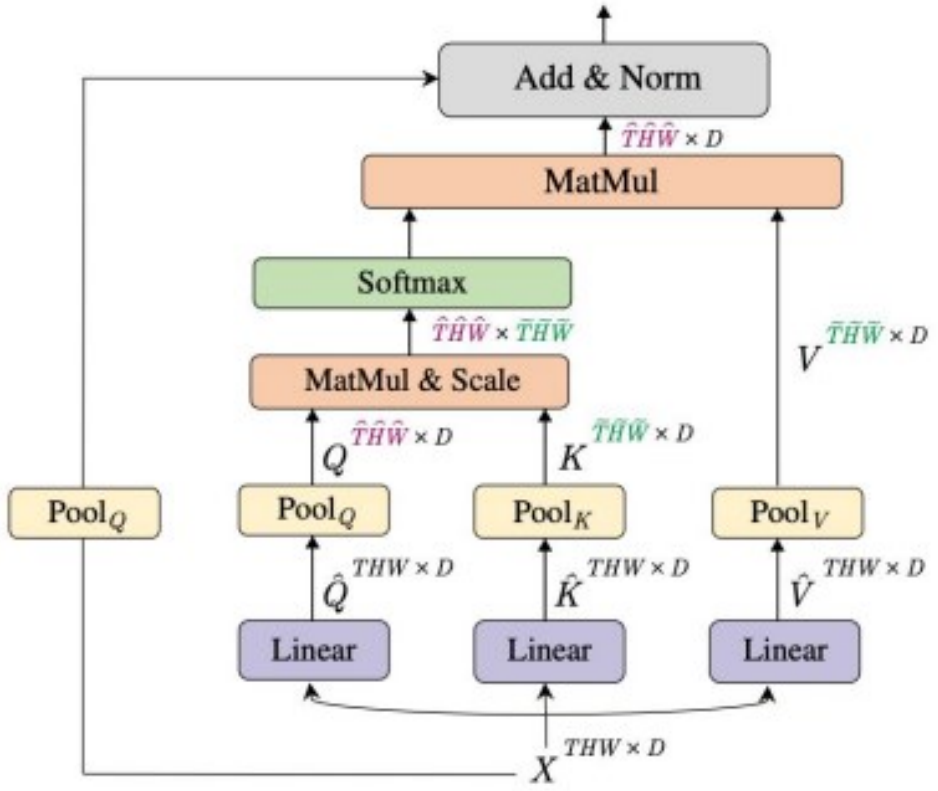

Fan et al, "Multiscale Vision Transformers", ICCV 2021

Li et al, "MViTv2: Improved Multiscale Vision Transformers for Classification and Detection", CVPR 2022

README.md

PySlowFast

PySlowFast is an open source video understanding codebase from FAIR that provides state-of-the-art video classification models with efficient training. This repository includes implementations of the following methods:

- SlowFast Networks for Video Recognition
- Non-local Neural Networks
- A Multigrid Method for Efficiently Training Video Models
- X3D: Progressive Network Expansion for Efficient Video Recognition
- Multiscale Vision Transformers
- A Large-Scale Study on Unsupervised Spatiotemporal Representation Learning
- MViTv2: Improved Multiscale Vision Transformers for Classification and Detection
- Masked Feature Prediction for Self-Supervised Visual Pre-Training
- Masked Autoencoders As Spatiotemporal Learners
- Reversible Vision Transformers

https://github.com/facebookresearch/SlowFast

Anticipation

Personal assistants and Future Predictions

Intelligent assistants should be able to understand what are the user's goals and what is going to happen in the future.

Next-active-object: **LOCKER**
Next action: **OPEN LOCKER**

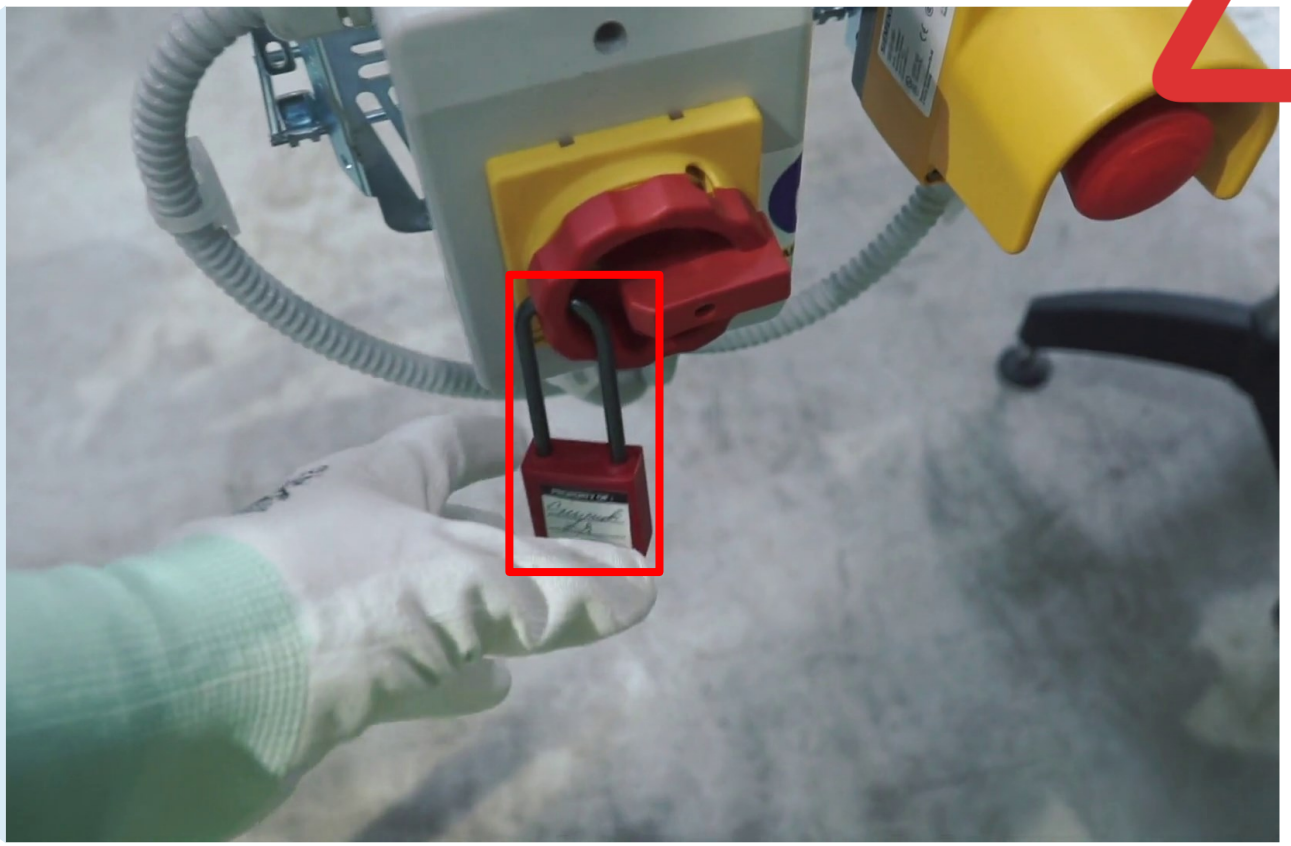

Future Predictions in Egocentric Vision

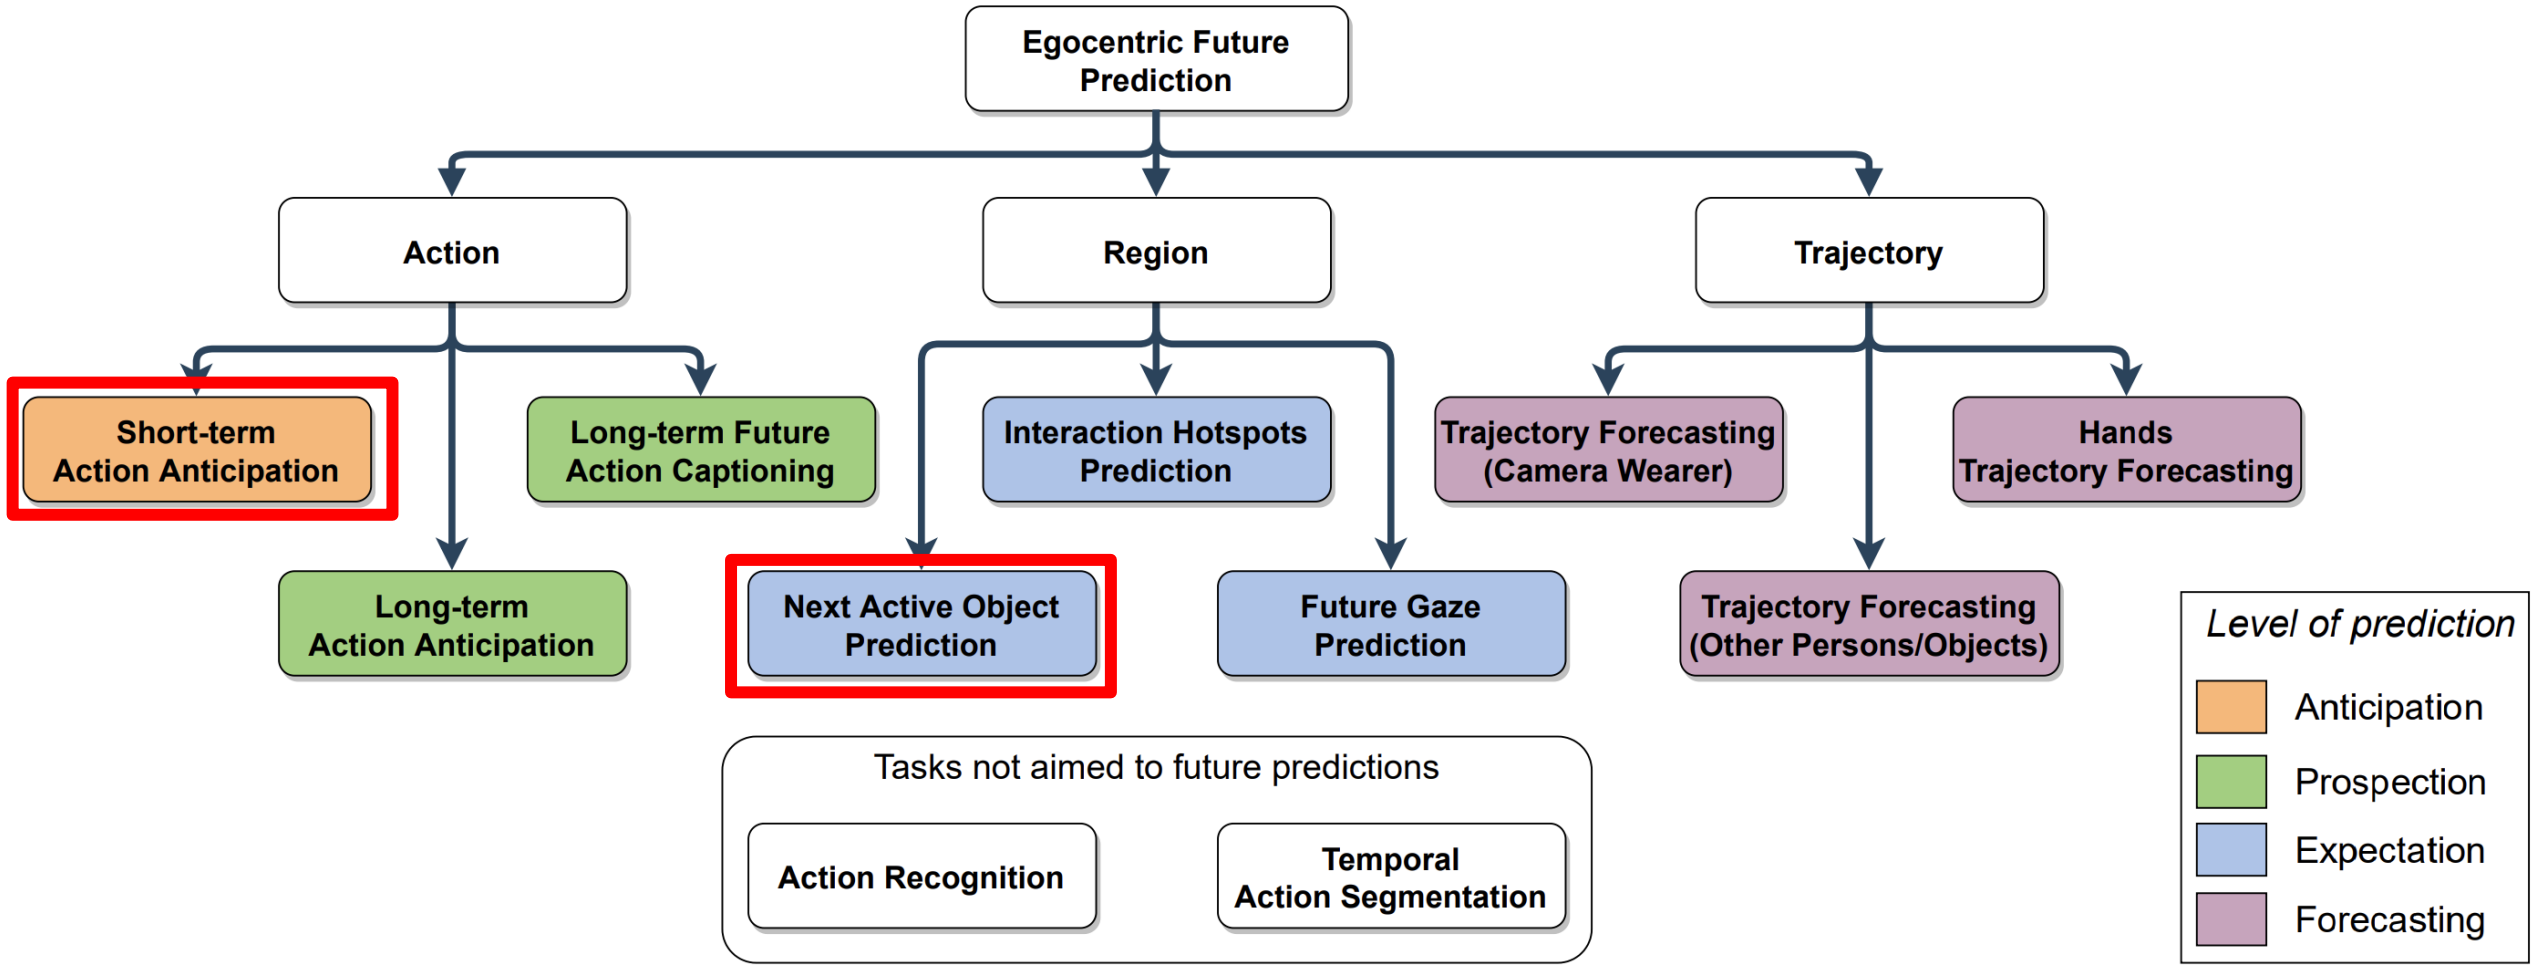

Ivan Rodin, Antonino Furnari, Dimitrios Mavroedis, Giovanni Maria Farinella (2021). Predicting the Future from First Person (Egocentric) Vision: A Survey. Computer Vision and Image Understanding, 211, pp. 103252.

Egocentric Action Anticipation Task

(observed video)

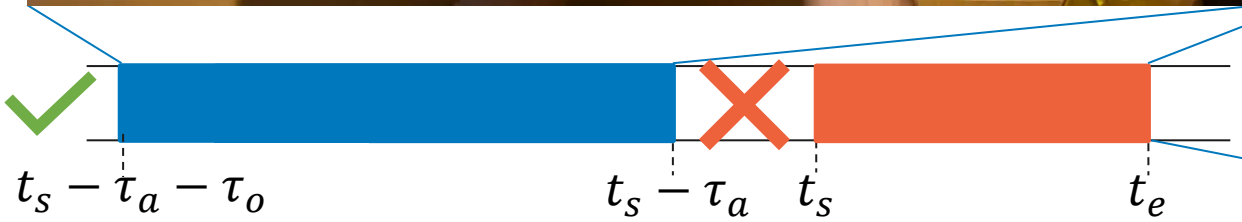

Model

Take - Plate

(unobserved)

τ_o arbitrary

$\tau_a = 1s$;

Damen, Dima, et al. "Scaling egocentric vision: The epic-kitchens dataset." *Proceedings of the European Conference on Computer Vision (ECCV)*. 2018.

Dima Damen et al. Rescaling Egocentric Vision . International Journal on Computer Vision (IJCV). 2021

Sequence to Sequence Models

Rolling-LSTM **ENCODING**

Unrolling-LSTM **INFERENCE**

We take inspiration from sequence to sequence models.

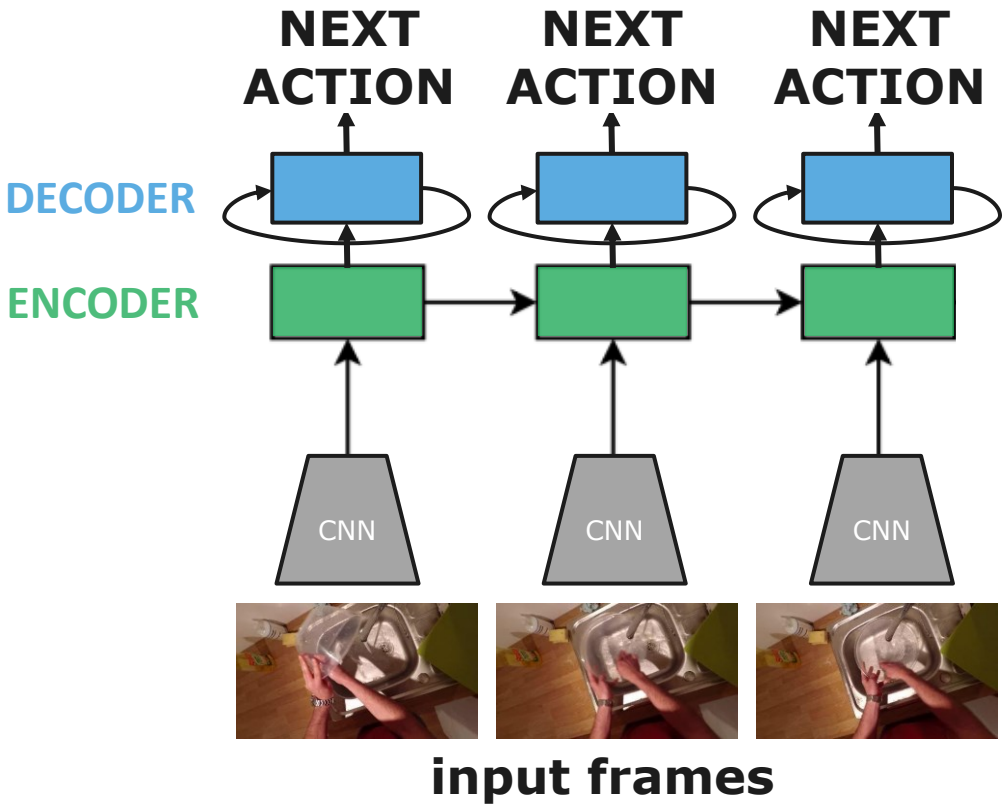

A. Furnari, G. M. Farinella, What Would You Expect? Anticipating Egocentric Actions with Rolling-Unrolling LSTMs and Modality Attention. ICCV 2019 (ORAL).

A. Furnari, G. M. Farinella. Rolling-Unrolling LSTMs for Action Anticipation from First-Person Video. TPAMI 2020. http://iplab.dmi.unict.it/rulstm

Demo Video: Egocentric Action Anticipation

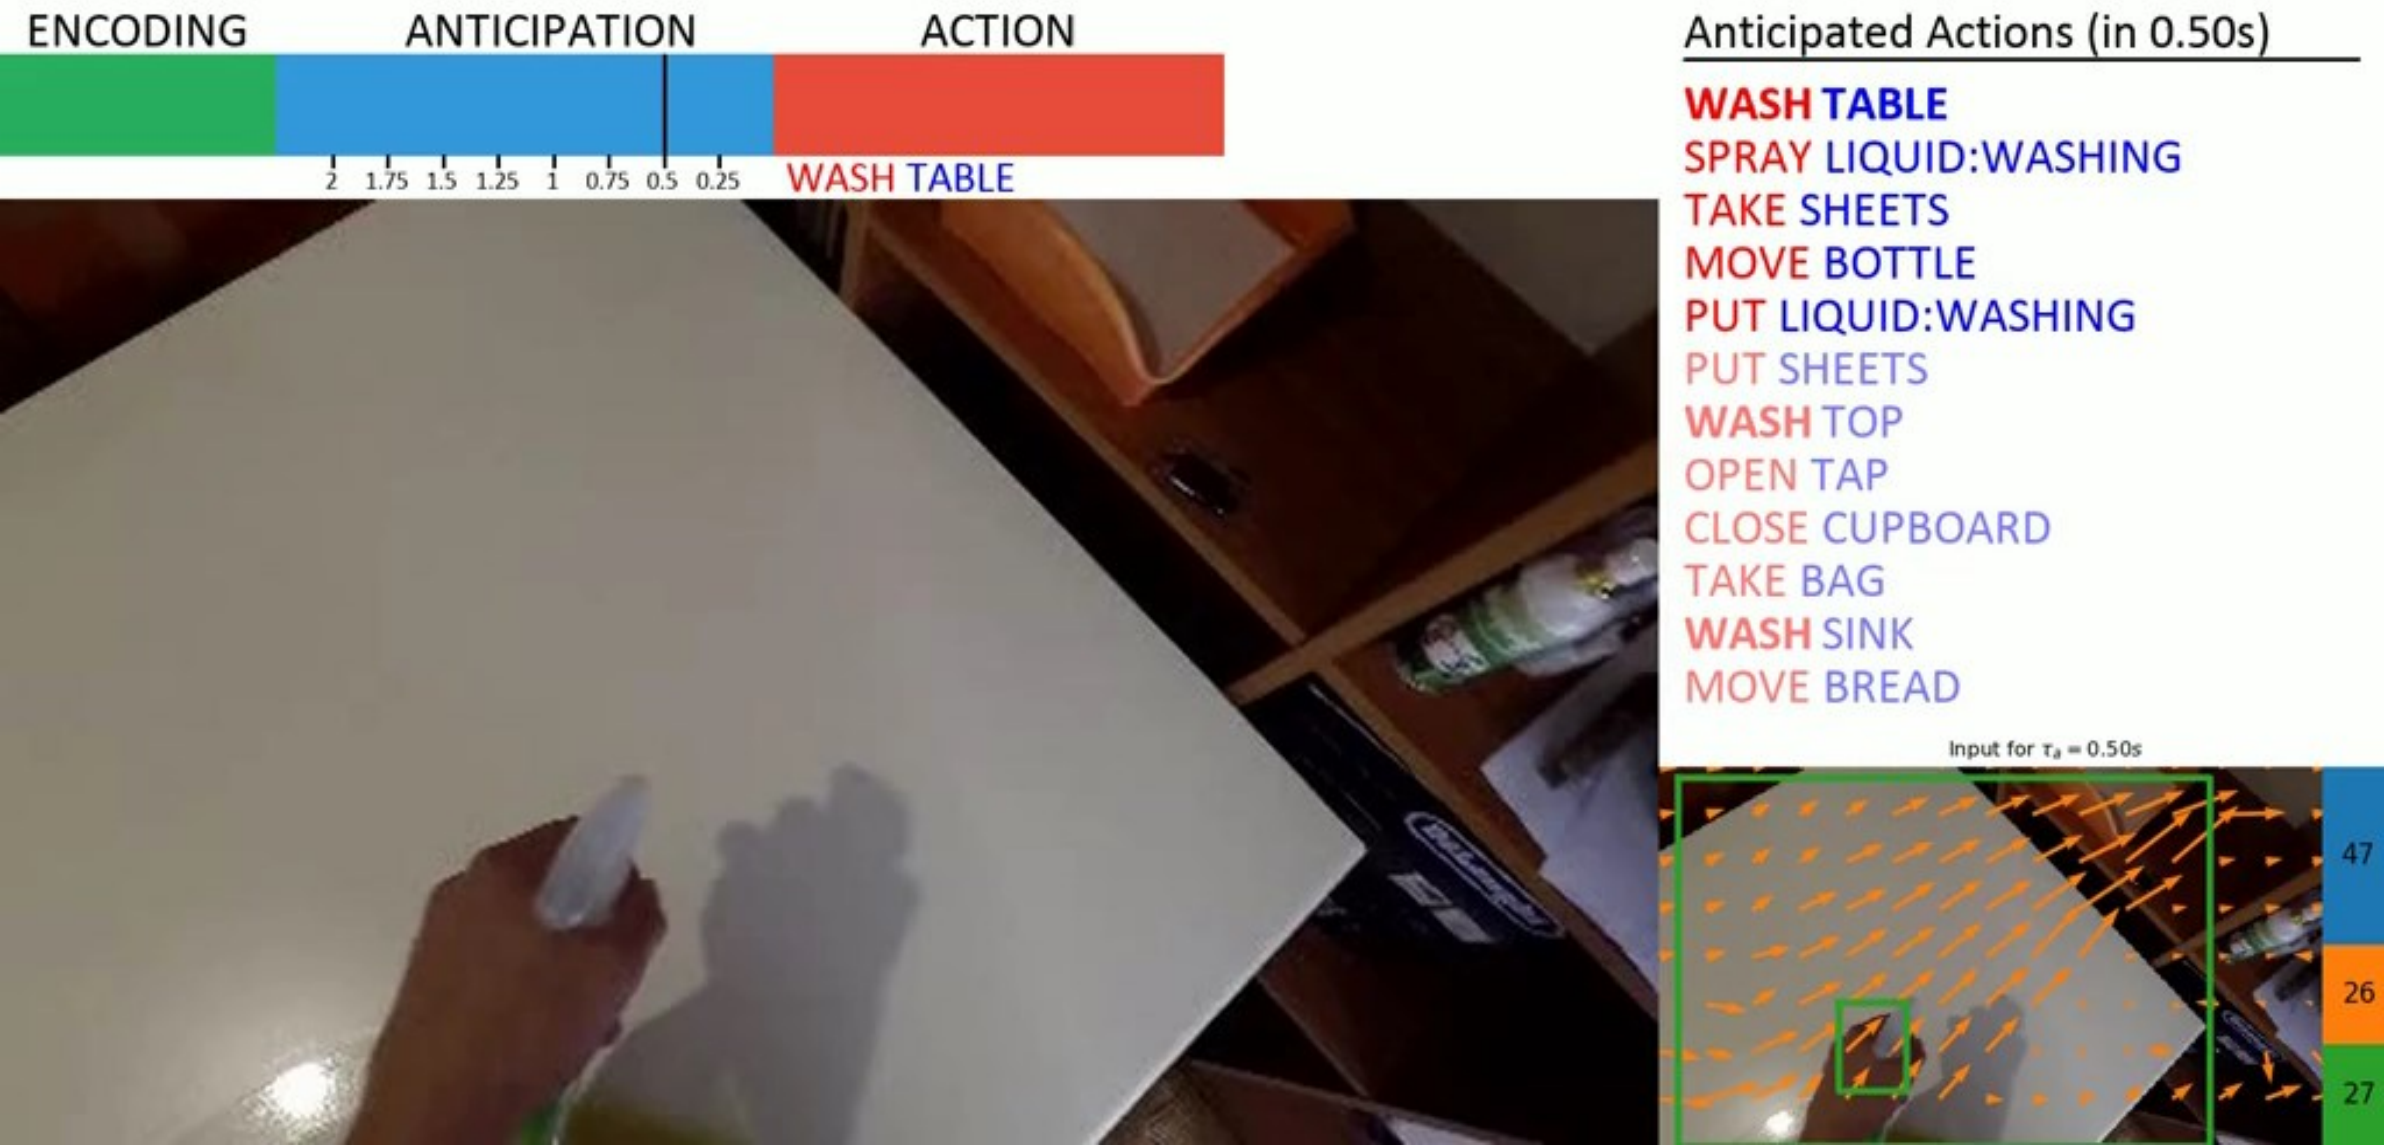

A. Furnari, G. M. Farinella, What Would You Expect? Anticipating Egocentric Actions with Rolling-Unrolling LSTMs and Modality Attention. ICCV 2019 (ORAL).

A. Furnari, G. M. Farinella. Rolling-Unrolling LSTMs for Action Anticipation from First-Person Video. TPAMI 2020. http://iplab.dmi.unict.it/rulstm

Anticipation – Next-Active-Objects

http://iplab.dmi.unict.it/NextActiveObjectPrediction/

Use egocentric object trajectories to distinguish passive from next-active-objects (i.e., those which will be used soon by the user).

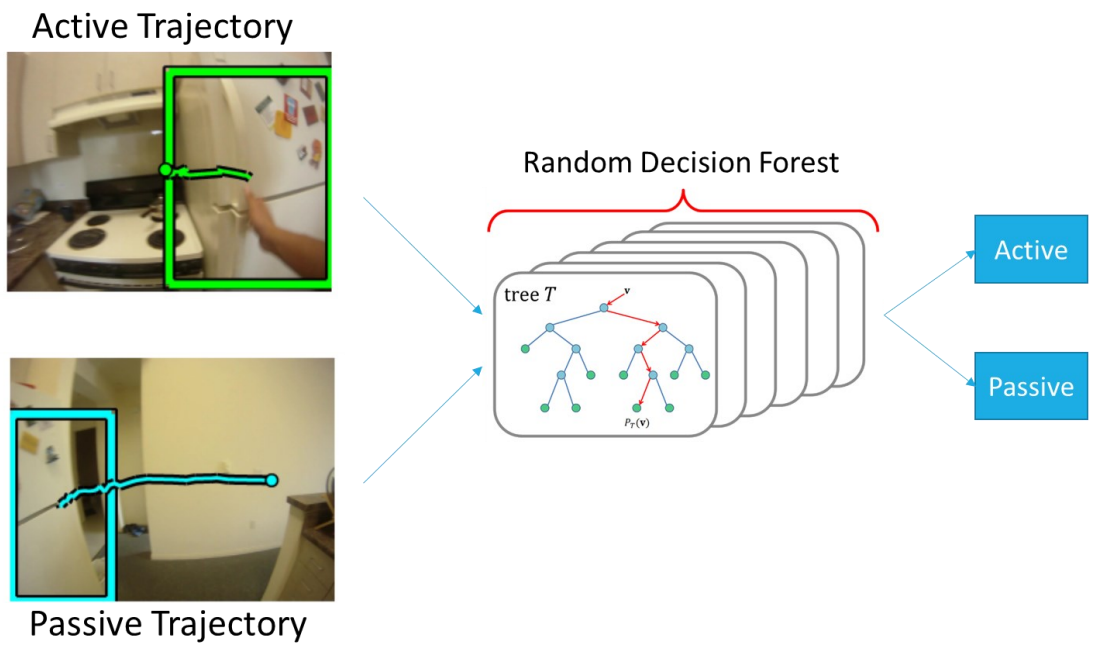

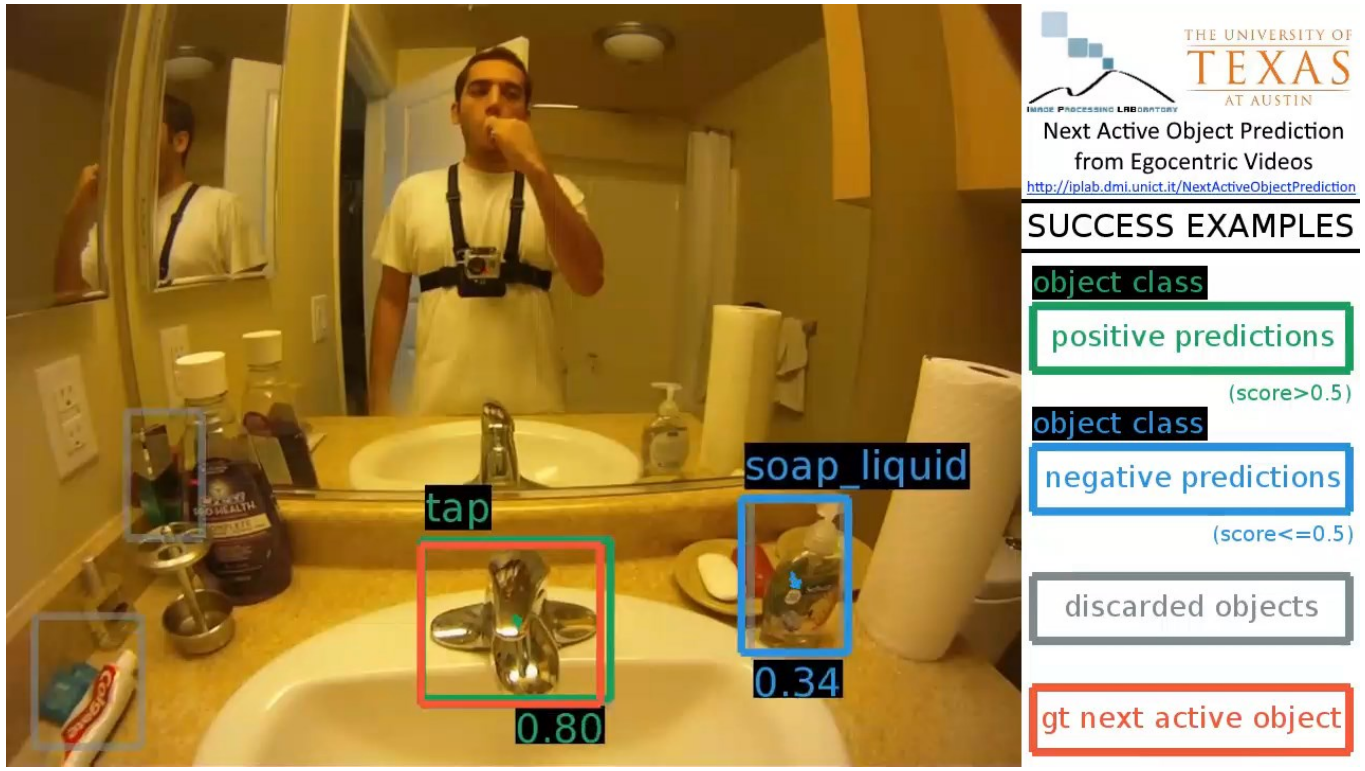

A. Furnari, S. Battiato, K. Grauman, G. M. Farinella, Next-Active-Object Prediction from Egocentric Videos, Journal of Visual Communication and Image Representation, 2017

Short-Term Object Interaction Anticipation

prediction

bbox = [1391,101,531,713]
noun = *wooden block*
verb = *take*
ttc = 0.75s
score = 0.83

Last observed frame (V_t)

Unobserved future frame ($V_{t+\delta}$)

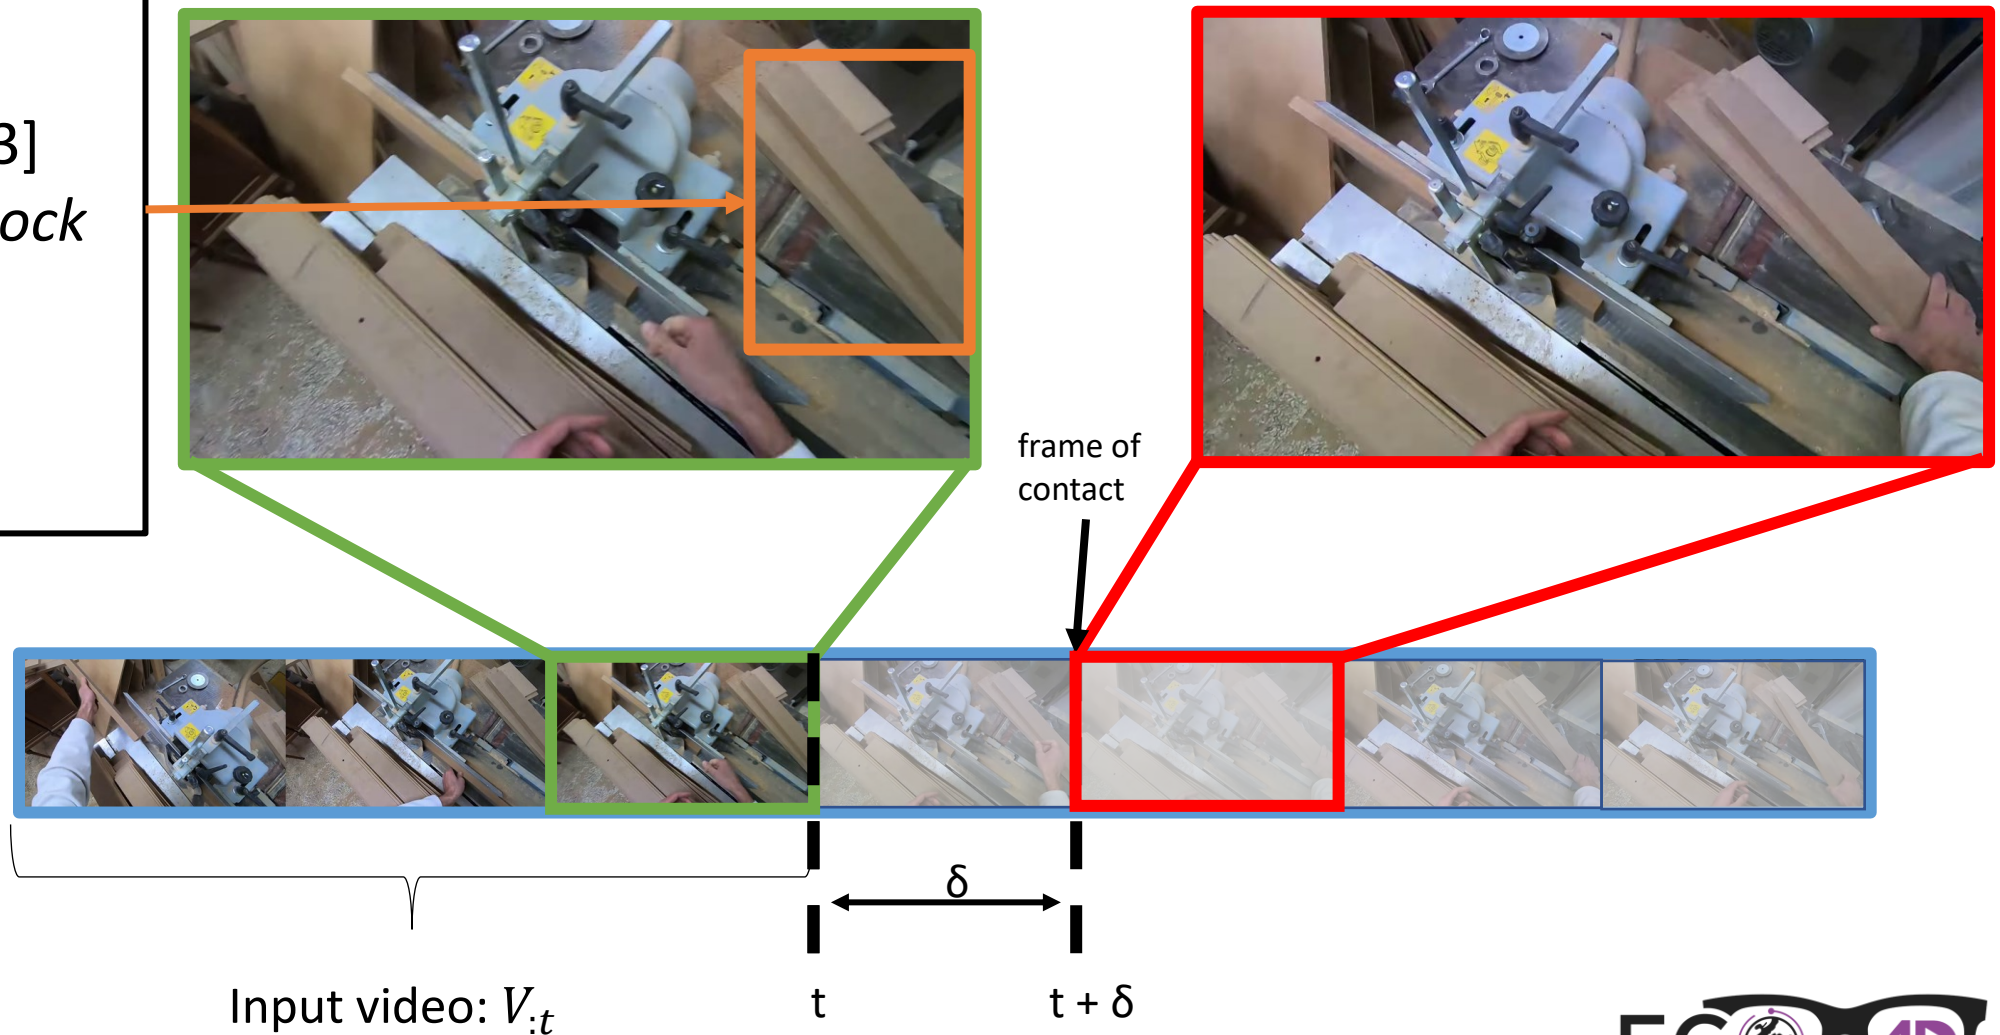

StillFast Network

An end-to-end approach for predicting next-active-objects based on an 2D-3D backbone taking as input a high resolution image and a video clip.

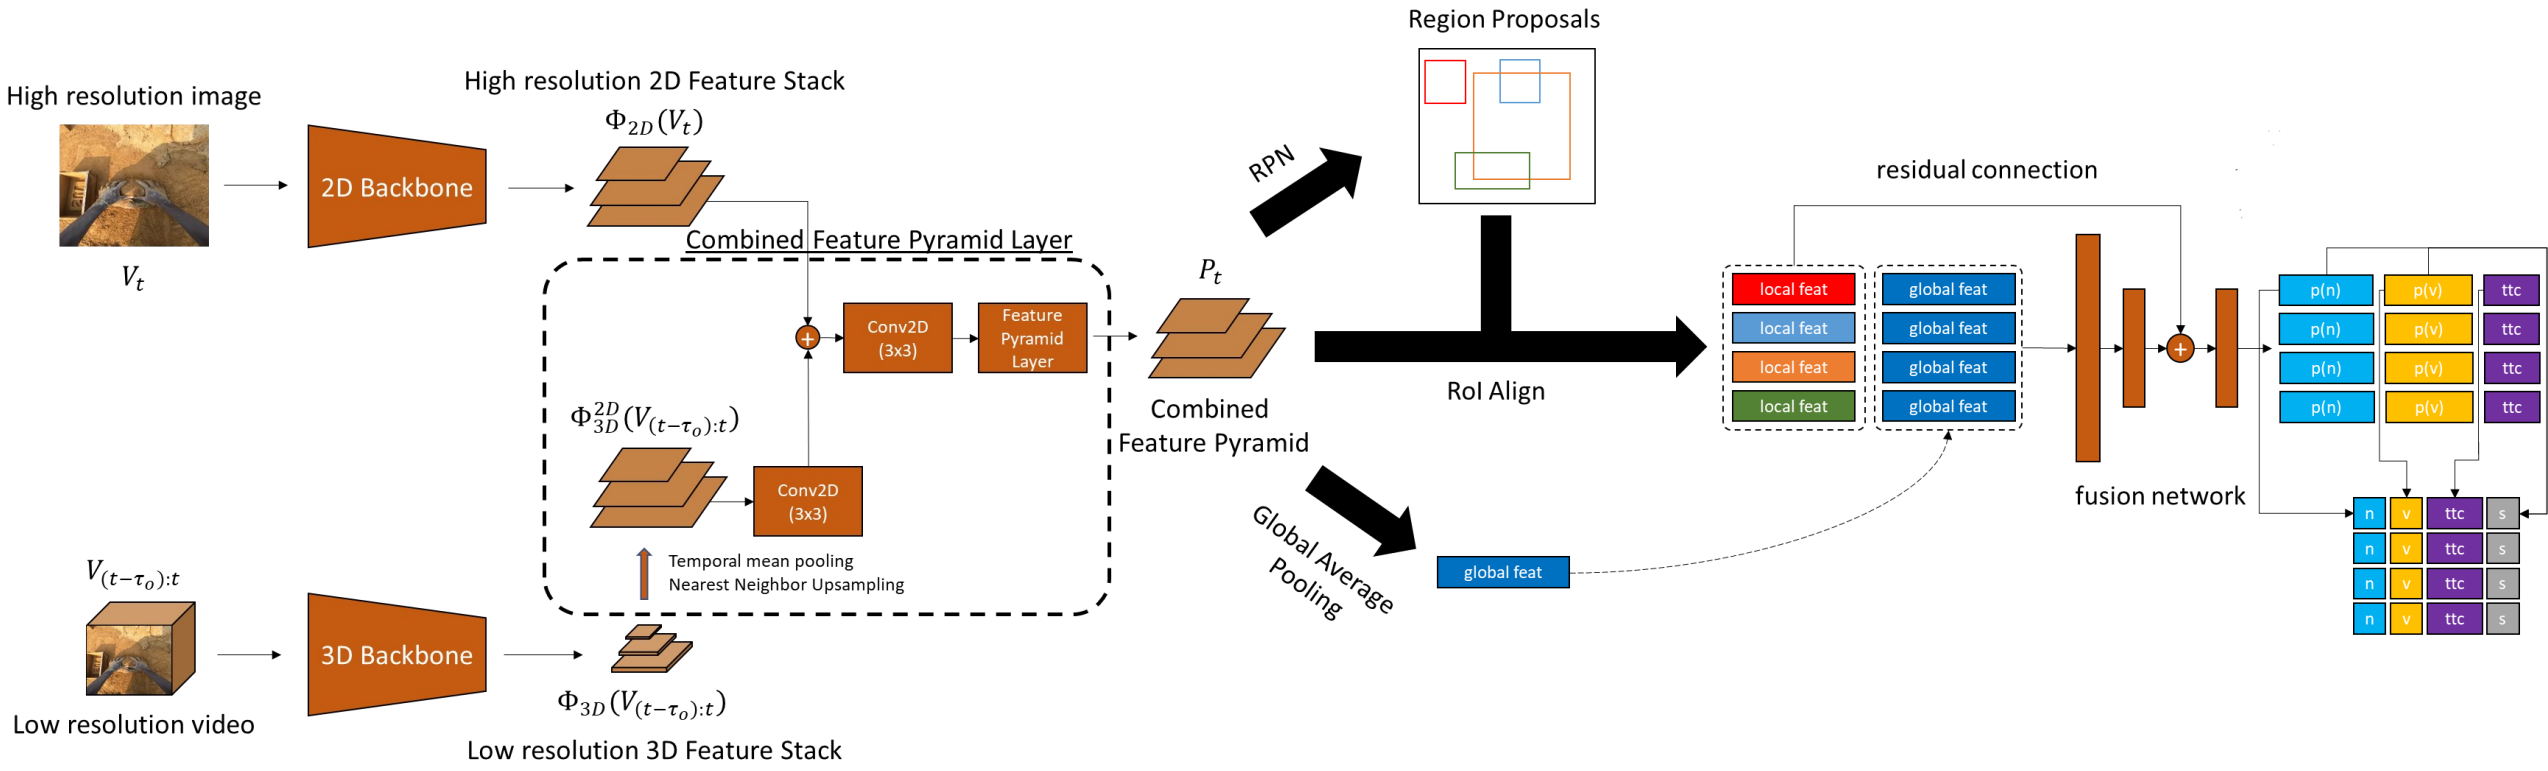

Francesco Ragusa, Giovanni Maria Farinella, Antonino Furnari (2023). StillFast: An End-to-End Approach for Short-Term Object Interaction Anticipation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops.

StillFast Qualitative Results

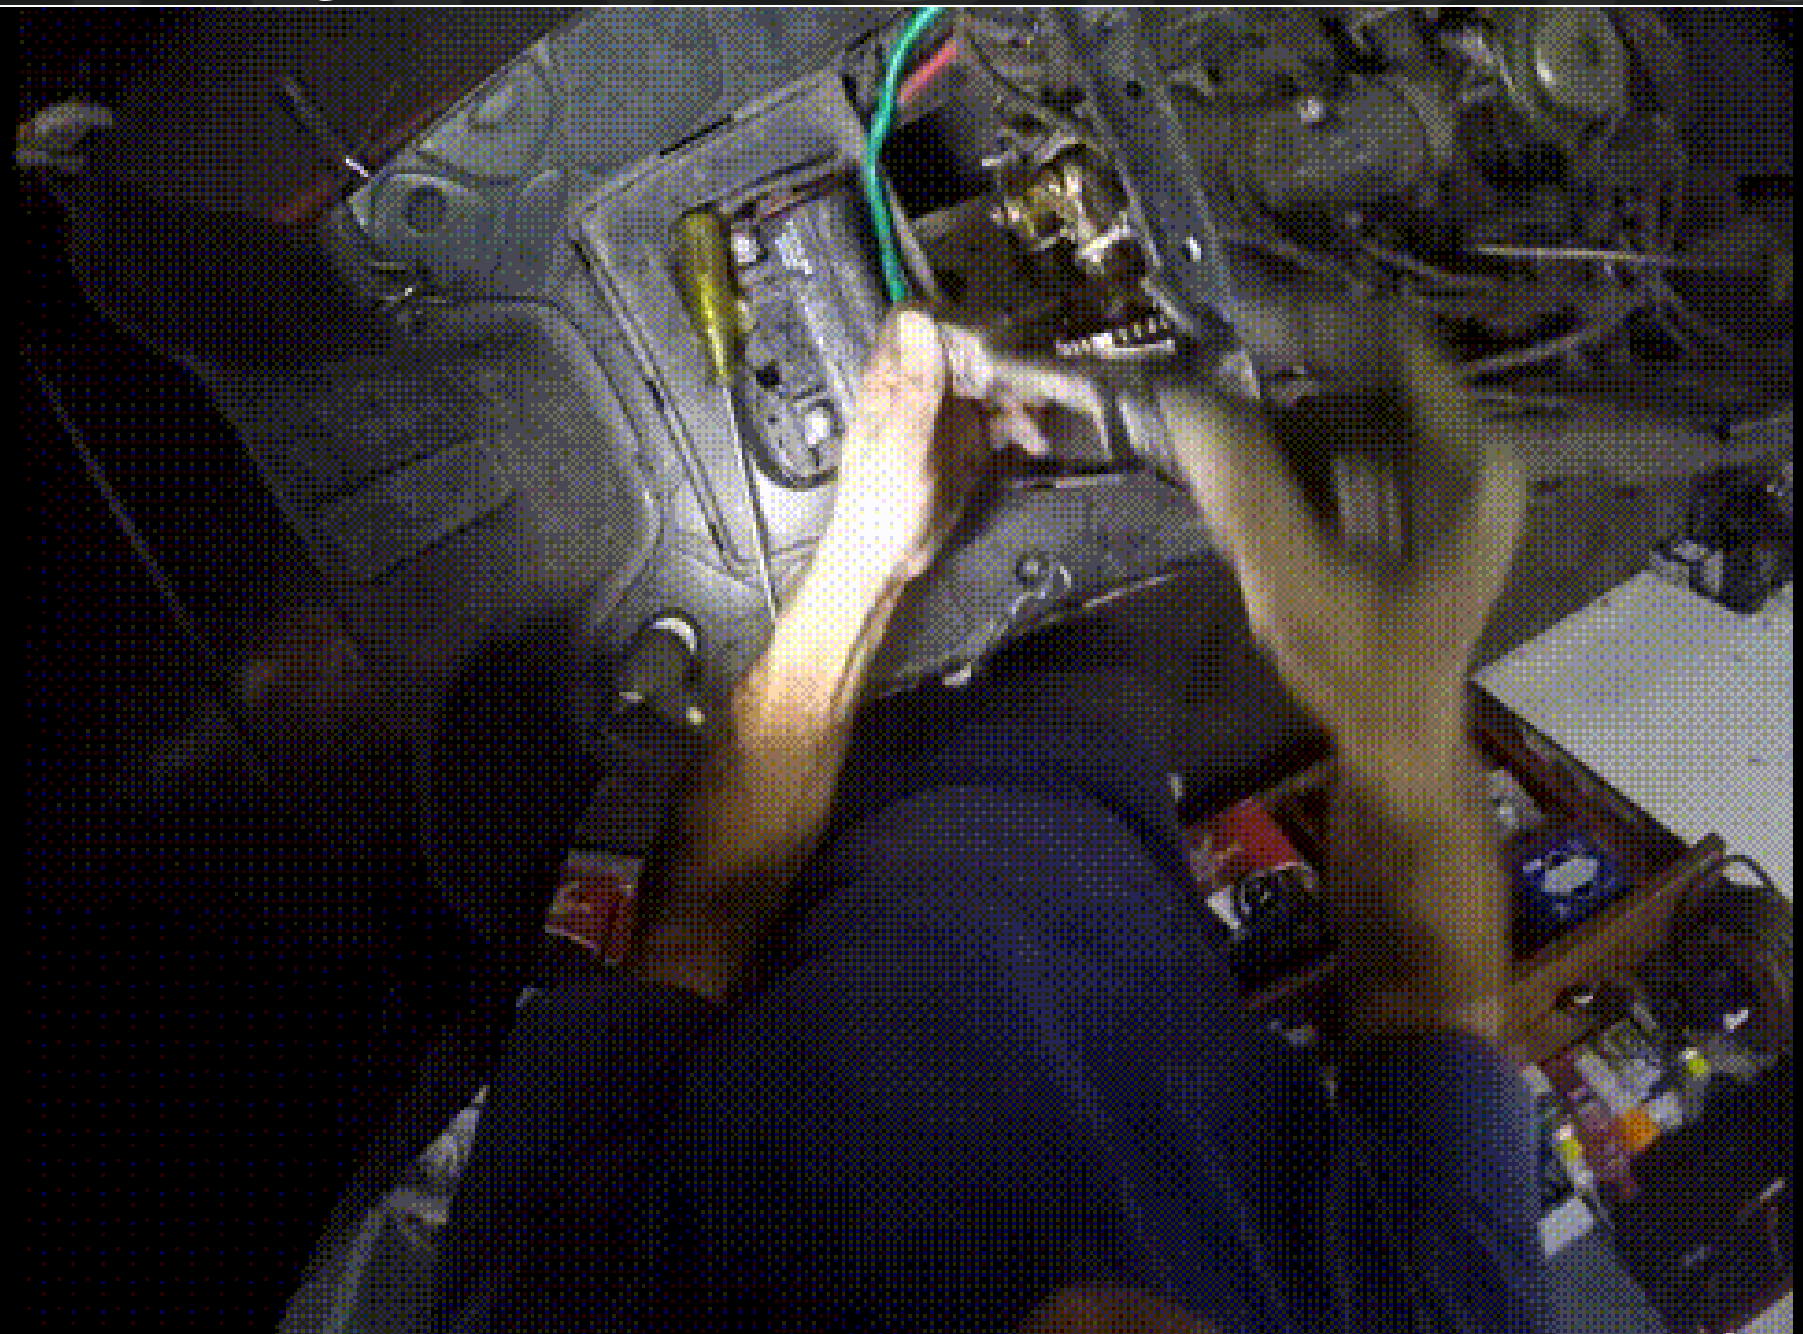

Leveraging Attention and Affordances for STA

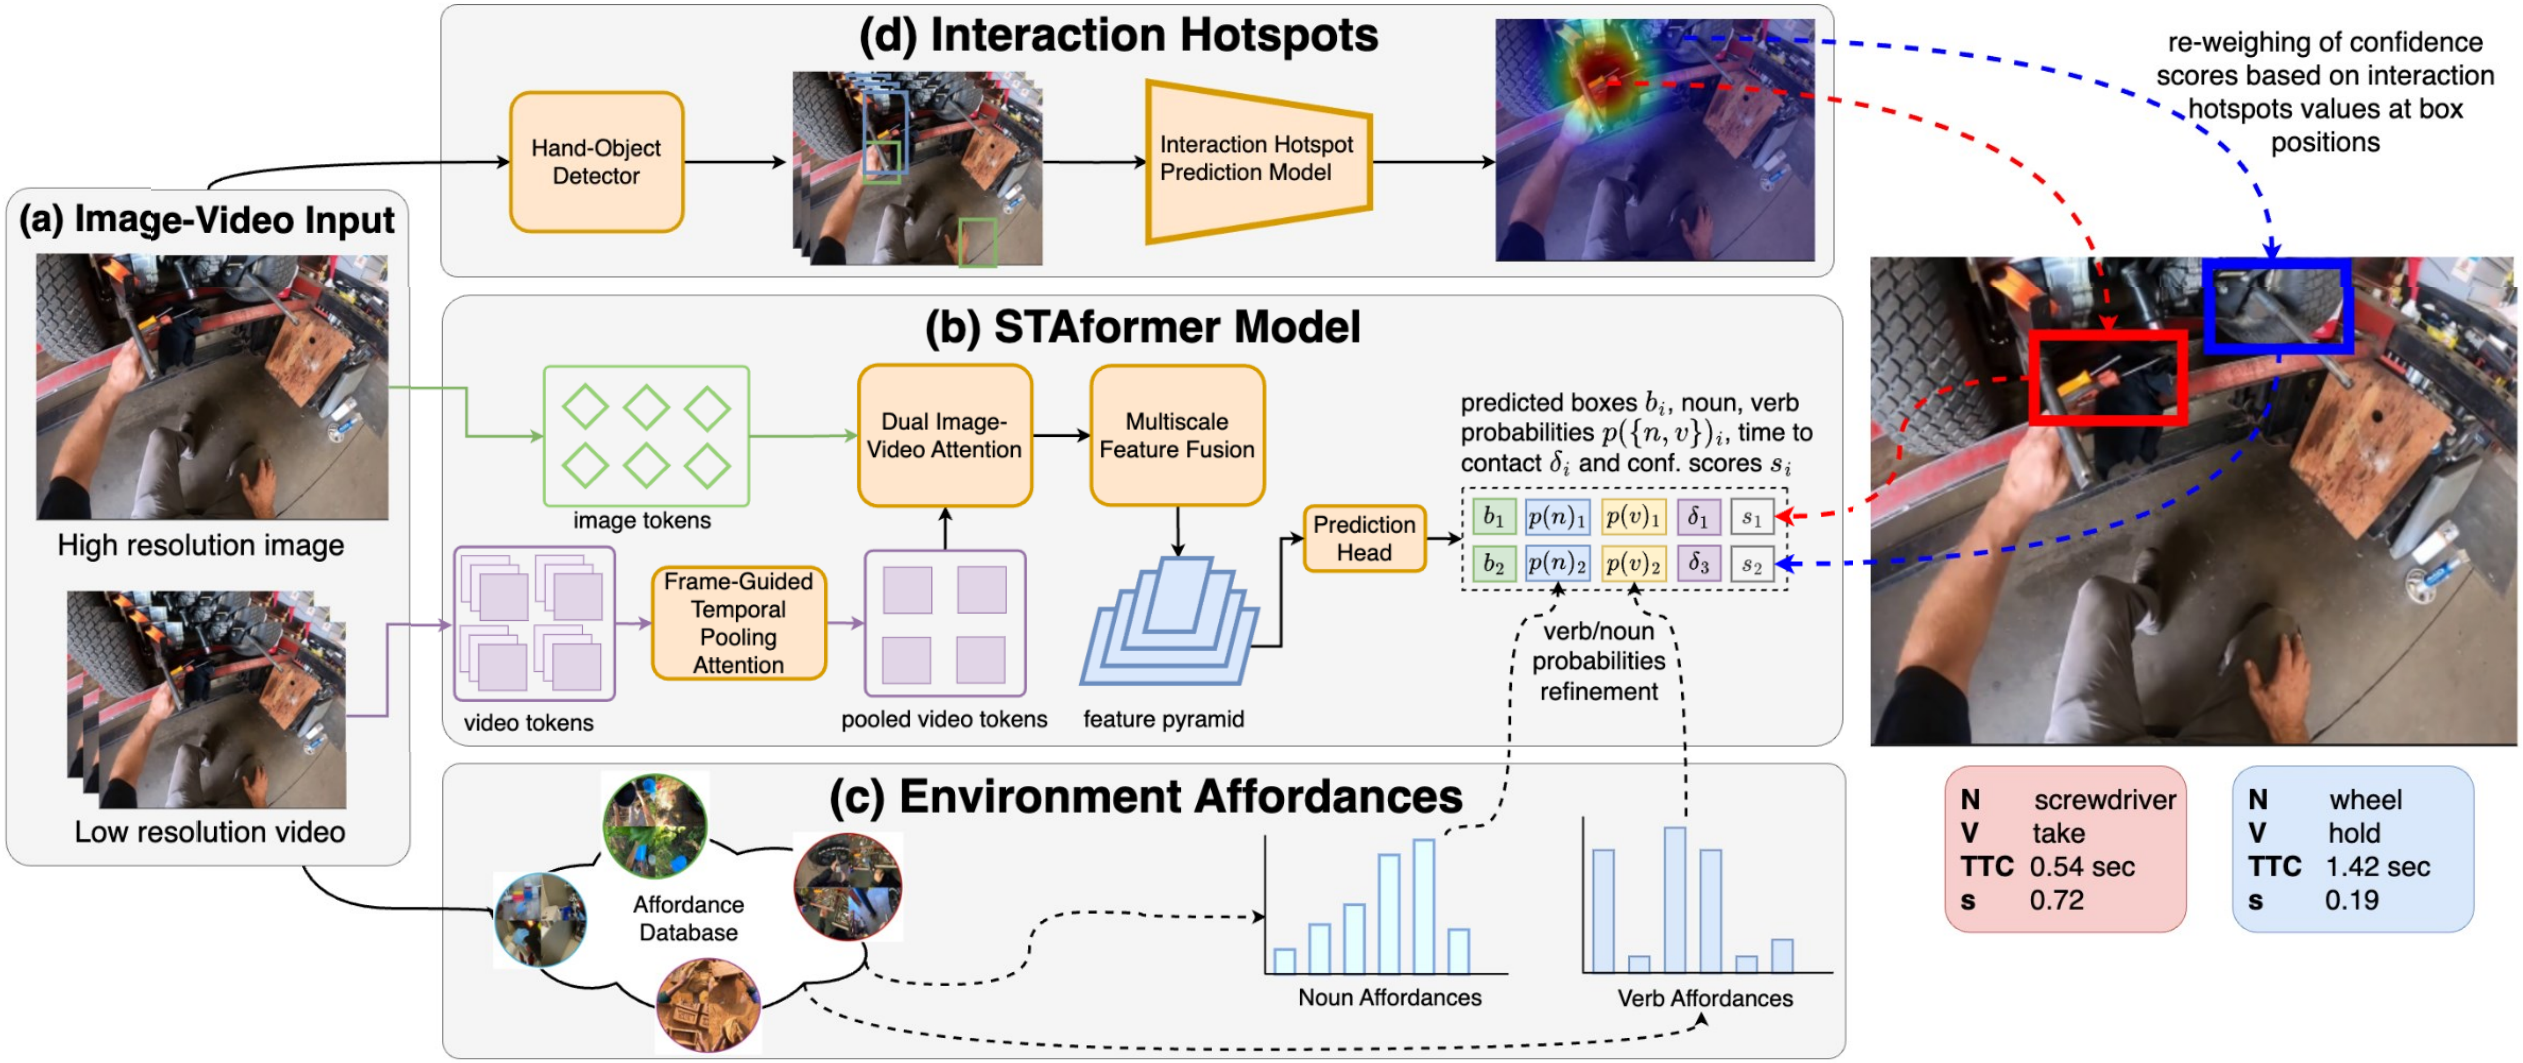

Lorenzo Mur-Labadia, Ruben Martinez-Cantin, Josechu Guerrero, Giovanni Maria Farinella, Antonino Furnari. AFF-ttention! Affordances and Attention models for Short-Term Object Interaction Anticipation . In European Conference on Computer Vision (ECCV) 2024

Can we bring egocentric vision to industry?

- The factory is a natural place for a wearable assistant;
- Closed-world assumption;
- Current research has considered different scenarios;
- No datasets in industrial-like scenarios;

The MECCANO Dataset

Data HERE -> https://iplab.dmi.unict.it/MECCANO/

We asked subjects to record egocentric videos while assembling a toy motorbike.

The assembly required to interact with several parts and two tools.

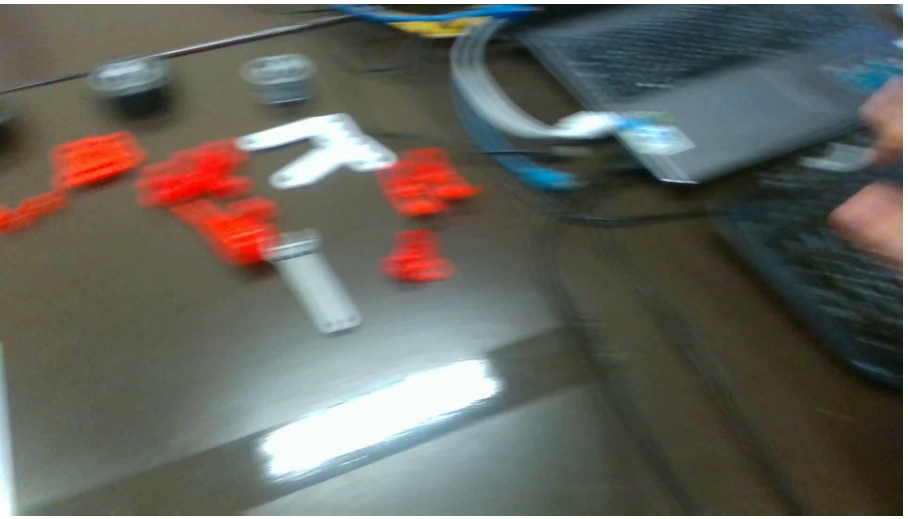

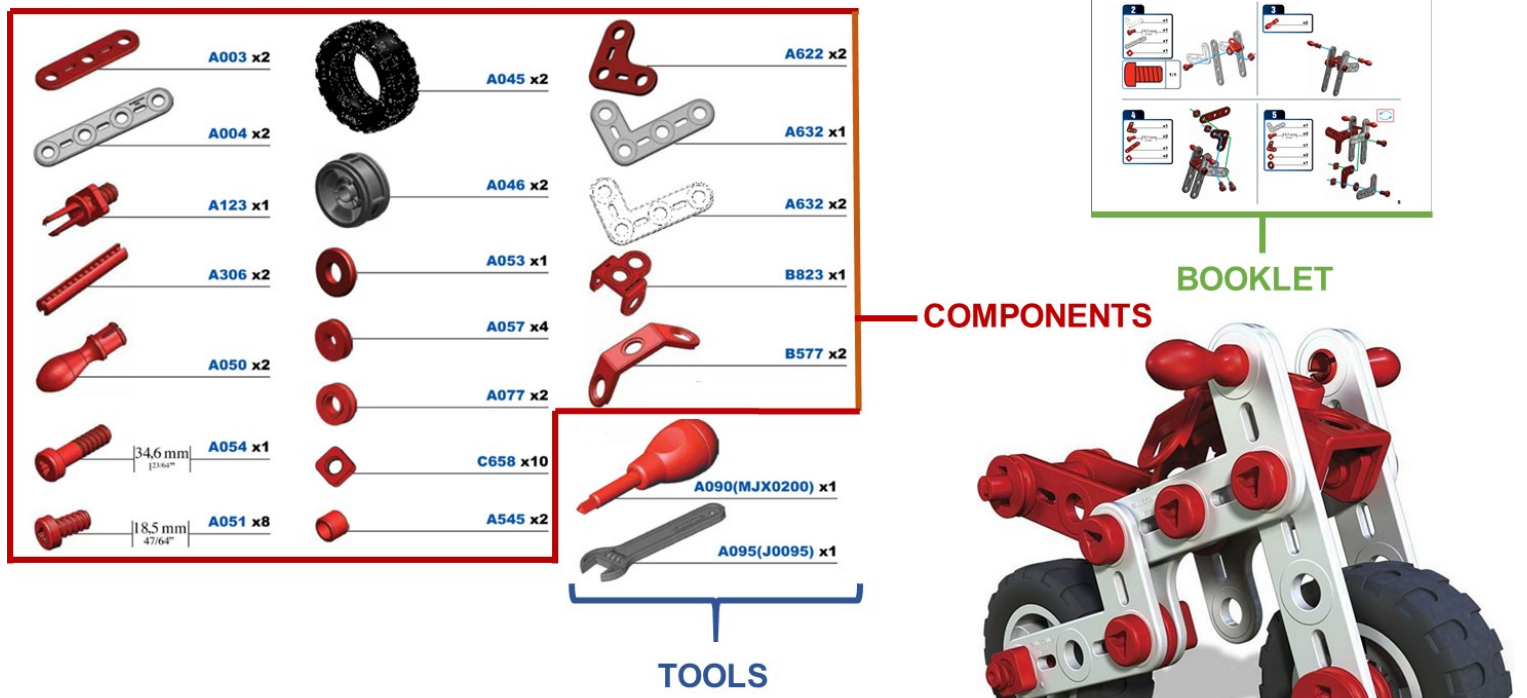

The scenario is industrial-like, with subjects undertaking interactions with tiny objects and tools in a sequential fashion to reach a goal.

Francesco Ragusa, Antonino Furnari, Salvatore Livatino, Giovanni Maria Farinella (2021). The MECCANO Dataset: Understanding Human-Object Interactions from Egocentric Videos in an Industrial-like Domain. In IEEE Winter Conference on Application of Computer Vision (WACV).

The ENIGMA-51 Dataset

We designed two procedures consisting of instructions that involve humans interacting with the objects present in the laboratory to achieve the goal of repairing two electrical boards

Low-Voltage

Hight-Voltage

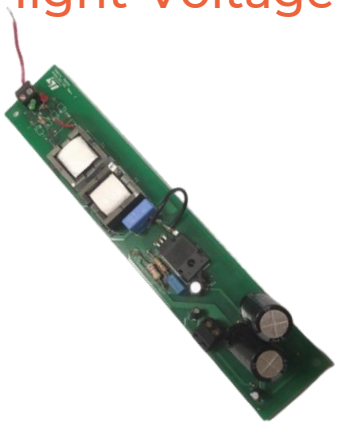

ENIGMA-51: Towards a Fine-Grained Understanding of Human Behavior in Industrial Scenarios. F. Ragusa R. Leonardi, M. Mazzamuto, C. Bonanno, R. Scavo, A. Furnari, G. M. Farinella. WACV (2024).

Industrial Applications

NEXT VISION

Spin-off of the University of Catania

https://www.nextvisionlab.it/

IMAGE PROCESSING LABORATORY

- Intelligent Navigation
- Image-based Localization
- Augmented Reality
- Multi-platform

Founders of Next Vision are authors of patents related to the developed technologies

https://drive.google.com/file/d/1lle4yF6b1kLp9P3ywqKOi77koTvn5OuE/view?usp=share_link

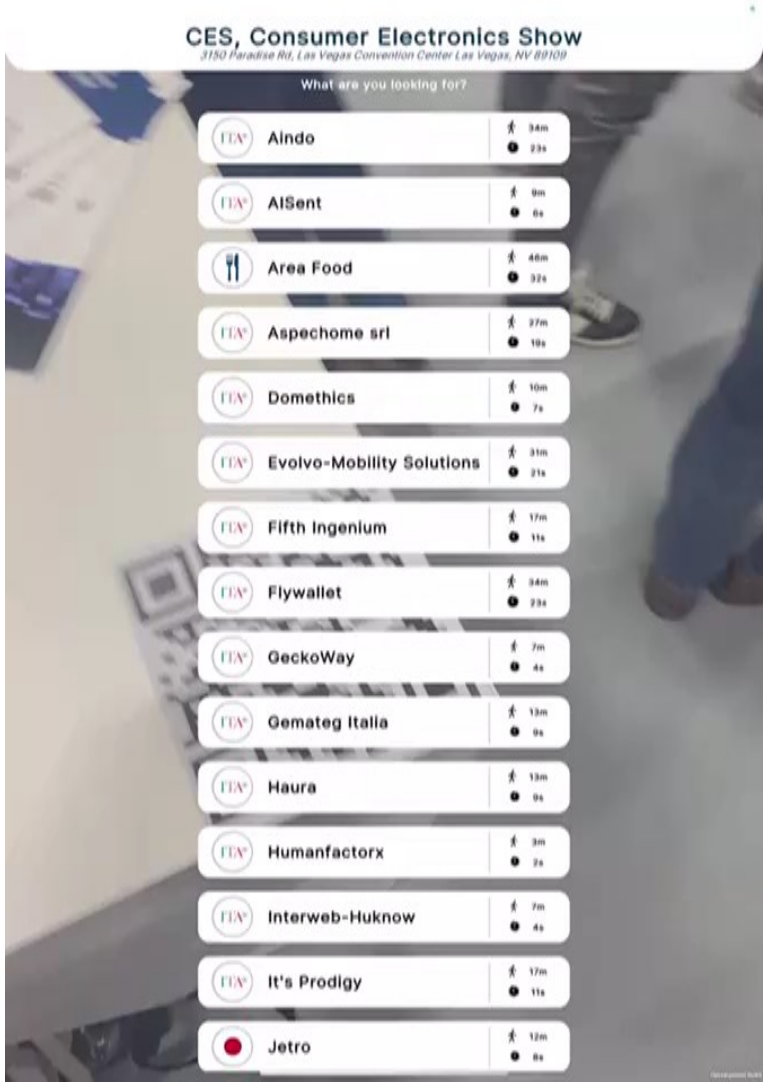

https://drive.google.com/file/d/1FAkLceBz_wCkDCsAJFq-nYBwFPZVciQV/view?usp=drive_link

NOKIA

Università di Catania

UNIONE EUROPEA REPUBBLICA ITALIANA REGIONE SICILIANA FESR SICILIA 2014-2020

CES

NAIROBI for Rehabilitation

Michele Mazzamuto, Francesco Ragusa, Antonino Furnari, Irene D'Ambra, Antonia Guarriera, Armando Sorbello, Giovanni Maria Farinella (2024). A Mixed Reality Application to Help Impaired People Rehabilitate Outside Clinical Environments. In IEEE International Conference on Metrology for eXtended Reality, Artificial Intelligence and Neural Engineering (MetroXRAINE).

NAIROBI for Rehabilitation

Michele Mazzamuto, Francesco Ragusa, Antonino Furnari, Irene D'Ambra, Antonia Guarriera, Armando Sorbello, Giovanni Maria Farinella (2024). A Mixed Reality Application to Help Impaired People Rehabilitate Outside Clinical Environments. In IEEE International Conference on Metrology for eXtended Reality, Artificial Intelligence and Neural Engineering (MetroXRAINE).

NAIROBI for Rehabilitation

Michele Mazzamuto, Francesco Ragusa, Antonino Furnari, Irene D'Ambra, Antonia Guarriera, Armando Sorbello, Giovanni Maria Farinella (2024). A Mixed Reality Application to Help Impaired People Rehabilitate Outside Clinical Environments. In IEEE International Conference on Metrology for eXtended Reality, Artificial Intelligence and Neural Engineering (MetroXRAINE).

- **NAOMI** is an AI Assistant able to support humans to monitor interactions, predict/anticipate next interactions, verify correctness in a sequence of interactions.

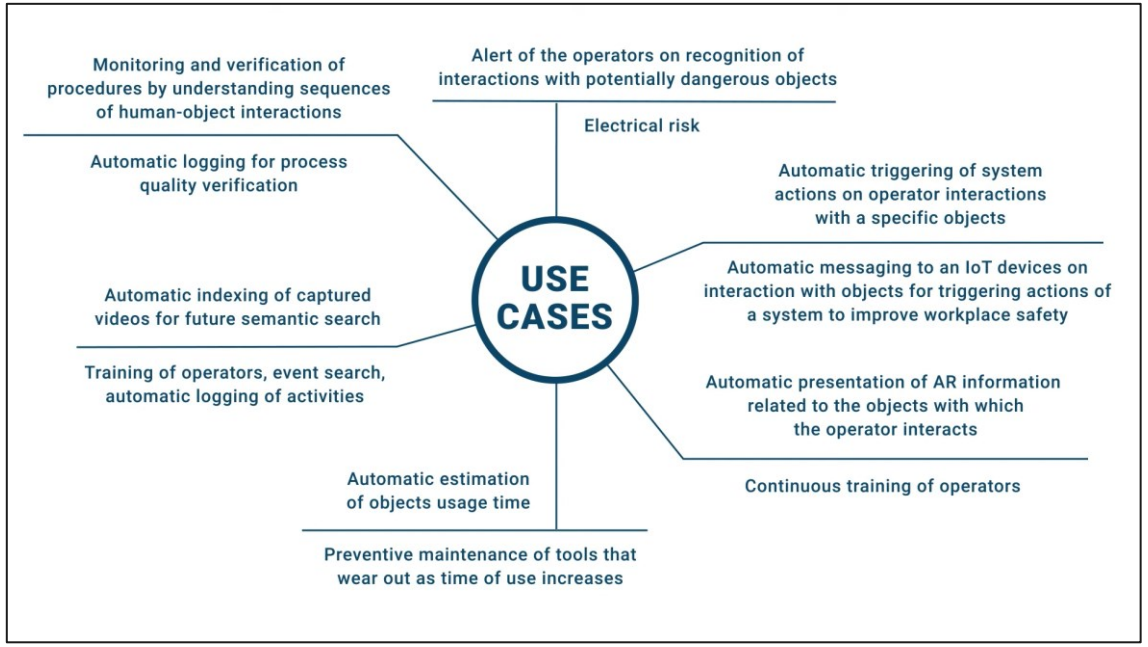

Use cases

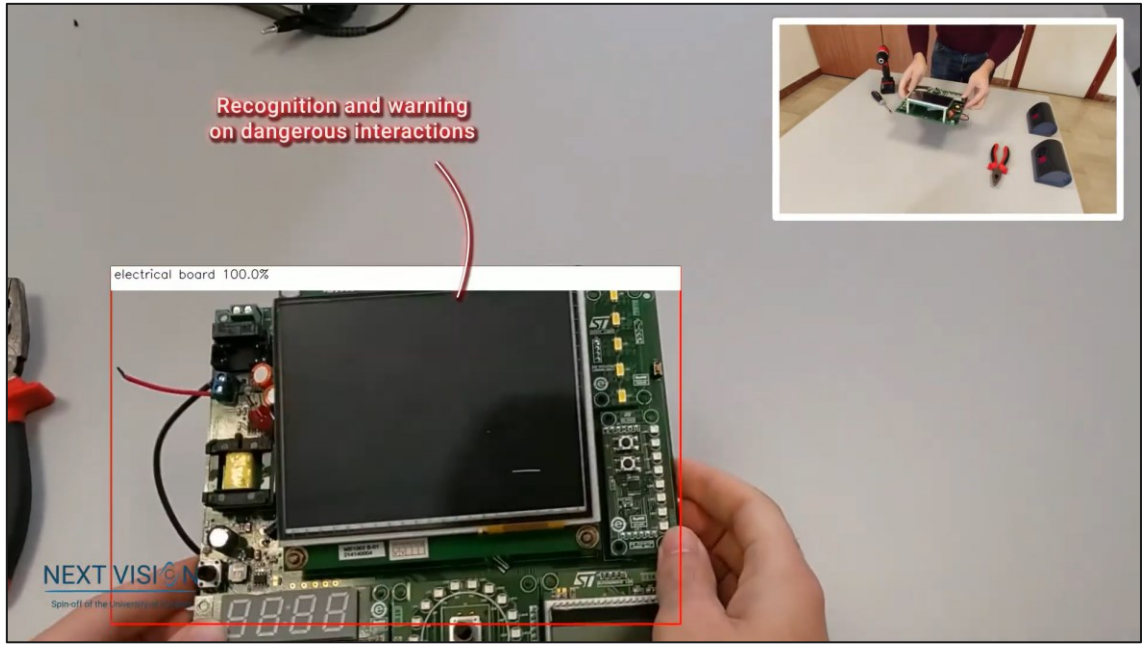

The video shows an example of object interaction monitoring.
The operator is notified on an interaction with a dangerous object.

https://drive.google.com/file/d/1oOvhVbyyR7AZ35I-V90Zy7RyRTR7lkD4/view?usp=drive_link

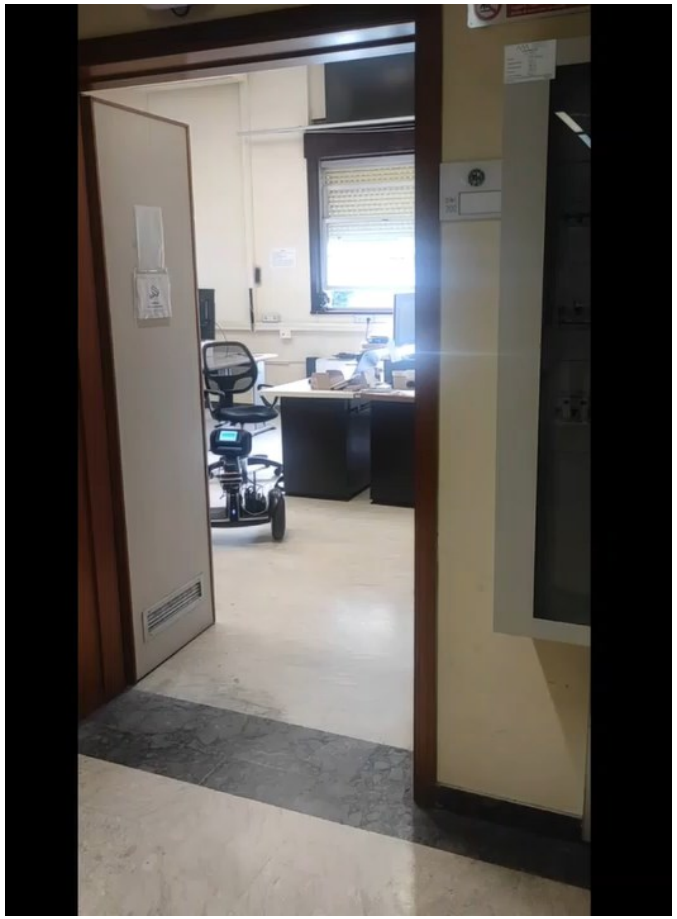

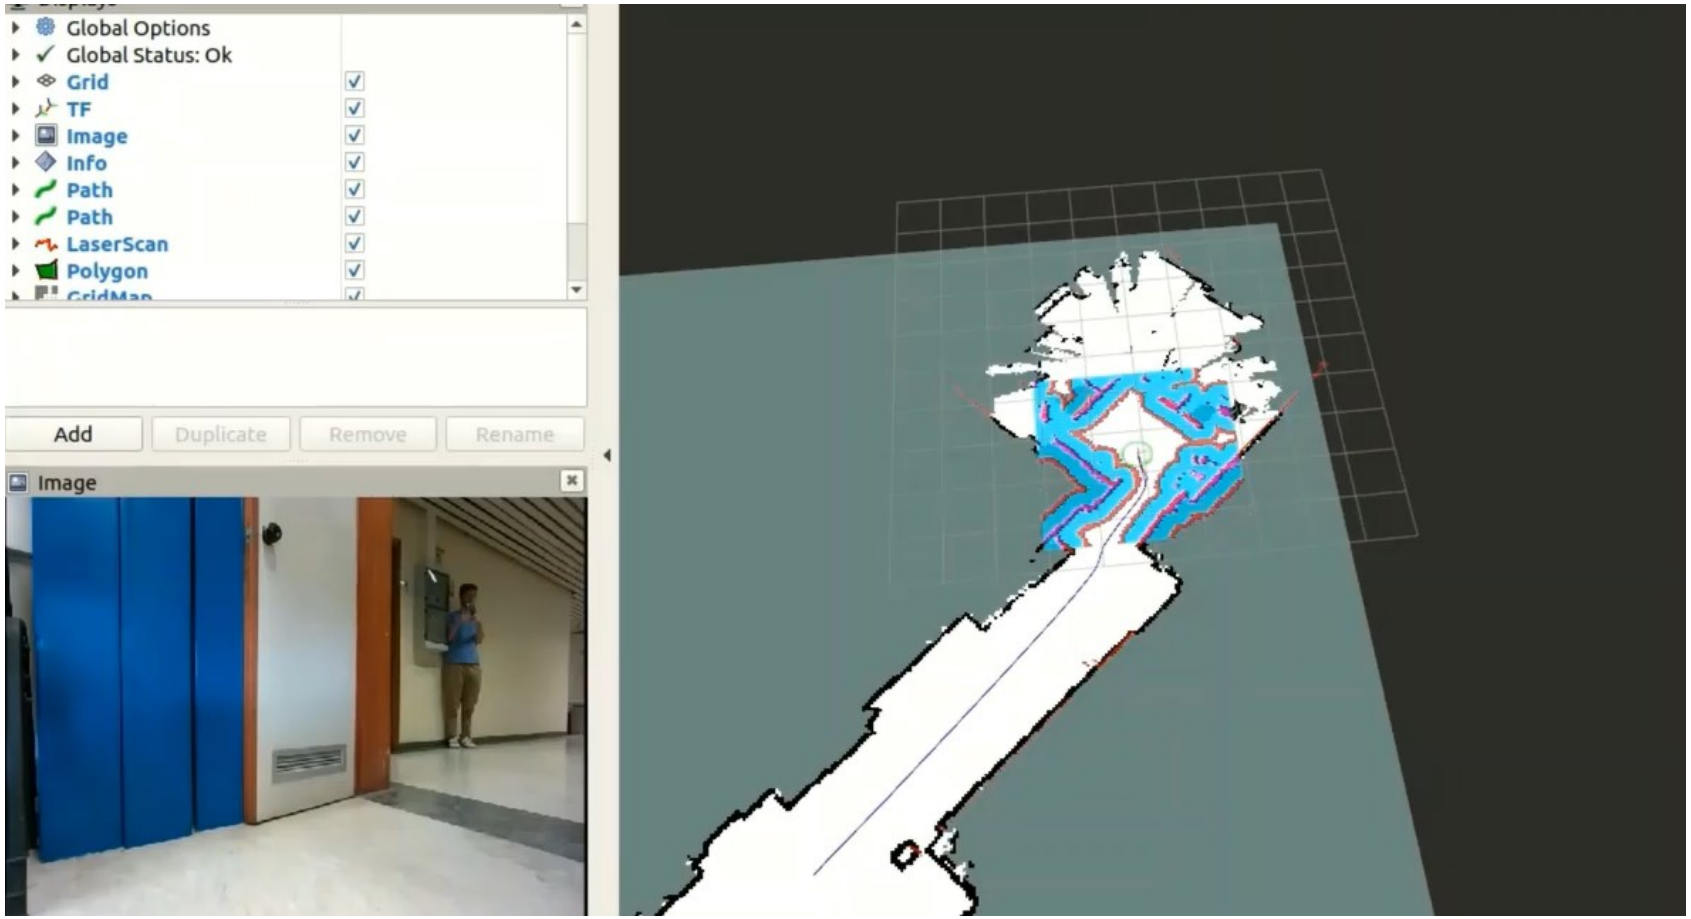

https://drive.google.com/file/d/17XrD-syy7pUm5MO4WYm7ZZxgRQsbSl2M/view?usp=drive_link

Marco Rosano, Francesco Ragusa, Antonino Furnari, Giovanni Maria Farinella (2023). MOVING: a MOdular and flexible platform for embodied VIsual NaviGation . In International Conference on Image Analysis and Processing (ICIAP)

Marco Rosano, Francesco Ragusa, Antonino Furnari, Giovanni Maria Farinella (2023). MOVING: a MOdular and flexible platform for embodied VIsual NaviGation . In International Conference on Image Analysis and Processing (ICIAP)

Università di Catania

NEPTUNE Project

Funded by the European Union NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

RAISE

In this conference!
Poster session 3A
Tomorrow 10:45 – 11:45

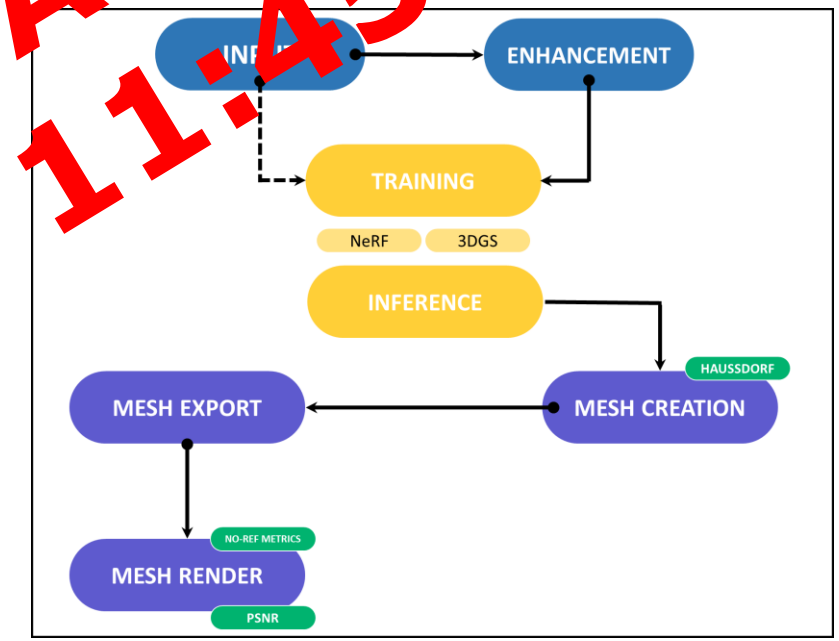

Neural Rendering

Mesh Rendering

S. M. Carota, A. Privitera, D. Di Mauro, A. Furnari, G. M. Farinella, F. Ragusa (2025). Benchmarking Neural Rendering Approaches for 3D Reconstruction of Underwater Environments. In International Conference on Computer Vision Theory and Applications (VISAPP)

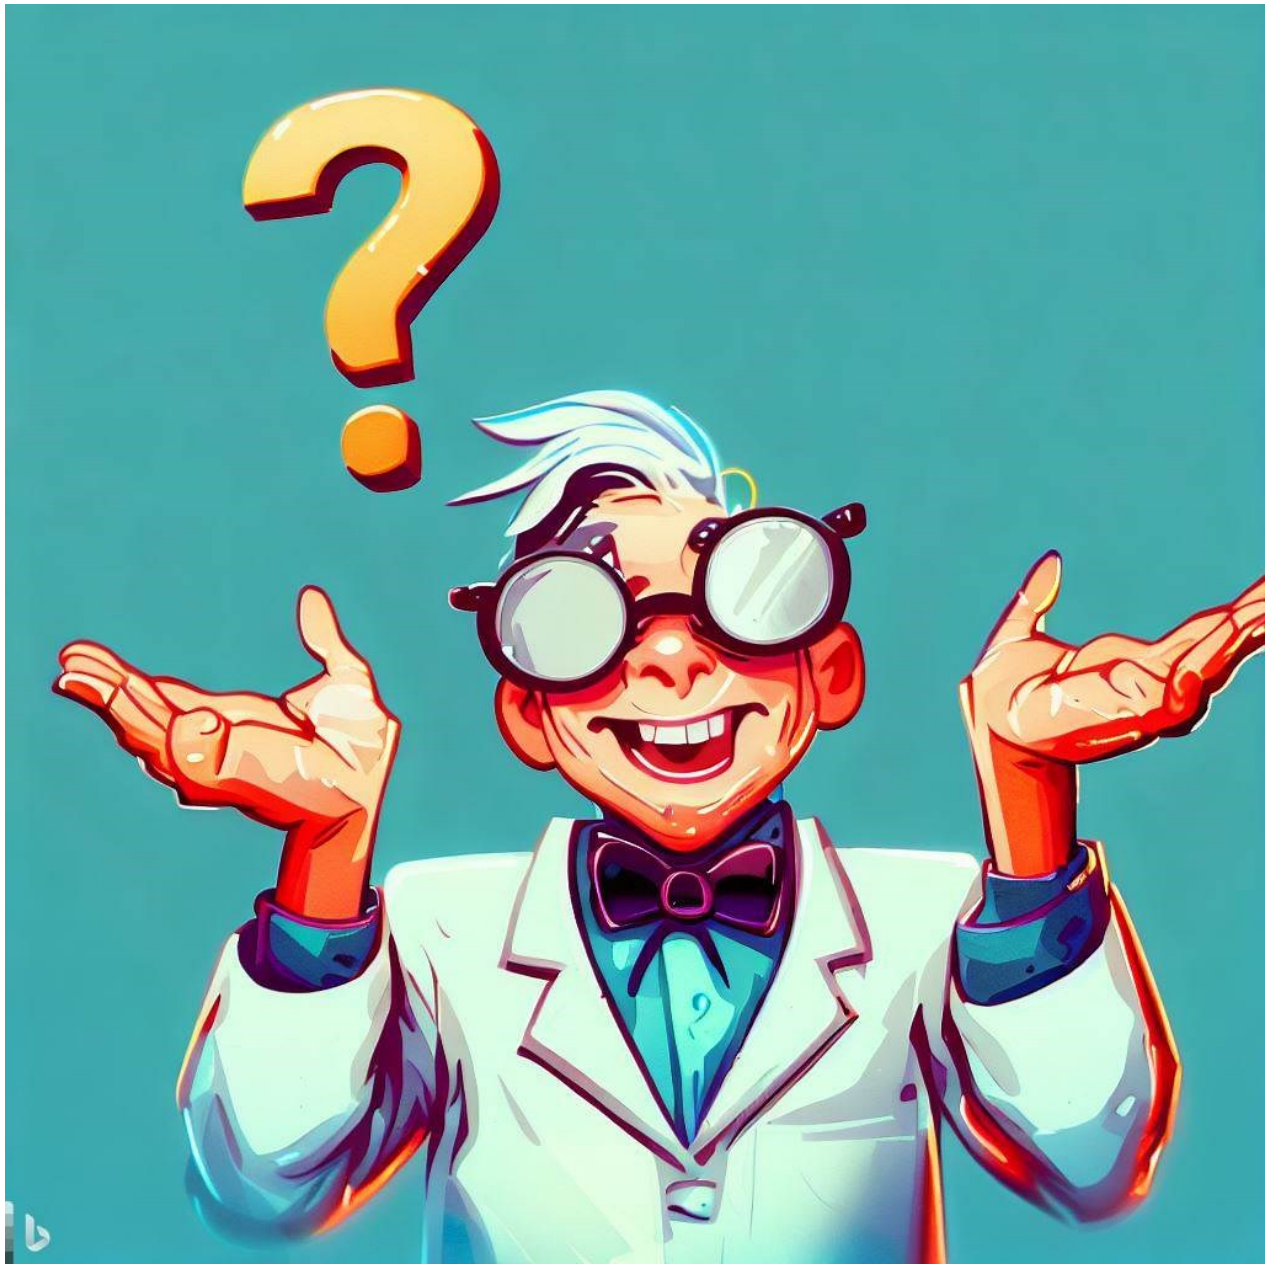

Doing Research in Egocentric Vision: Where to start?

Accessing Modern Datasets – License

Data nowadays carries a lot of privacy/social/economic implications, so modern datasets are usually licensed.

! pay attention to which uses are permitted!

EPIC KITCHENS

ABOUT STATS DOWNLOADS CHALLENGES TEAM

Disclaimer

EPIC-KITCHENS-55 and EPIC-KITCHENS-100 were collected as a tool for research in computer vision. The dataset may have unintended biases (including those of a societal, gender or racial nature).

Copyright

All datasets and benchmarks on this page are copyright by us and published under the Creative Commons Attribution-NonCommercial 4.0 International License. This means that you must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use. You may not use the material for commercial purposes.

For commercial licenses of EPIC-KITCHENS and any of its annotations, email us at uob-epic-kitchens@bristol.ac.uk

EGO4D License Agreement

Obtaining the dataset or any annotations requires you first review our license agreement and accept the terms. Go here (ego4ddataset.com) to review and execute this agreement, and you will be emailed a set of AWS access credentials when your license agreement is approved, which will take ~48hrs. In the meantime, you can check out data overview & sample notebooks here to get familiar with the dataset, and can download the CLI & dataloaders to get setup in advance.

Note that licenses have the option to execute our license agreements as either an individual or on behalf of your institution. You will likely sign the license as an individual. Typically, only institutional signatories at a director or executive level can agree to license terms on behalf of an entire organization.

Also note that once approved your access credentials will expire in 14 days - you're expected to download the data locally, not to consume it from AWS. You can easily renew your license once it expires though: license renewal FAQ

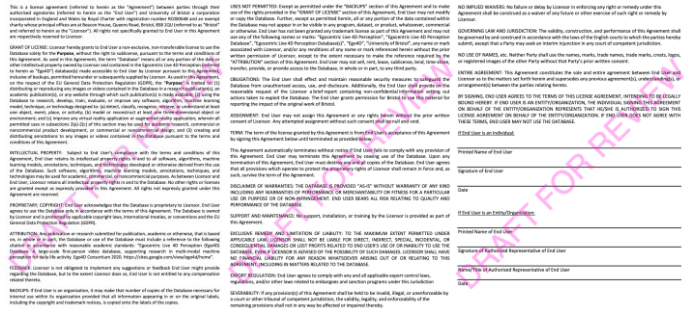

Ego4D Dataset

This information you enter below will be used to generate a data usage agreement. You will receive an email from HelloSign which will step you through the process of signing all the agreements. You can review the data usage agreement at —

http://ego4d.github.io/pdfs/Ego4D-Licenses-Draft.pdf

Note: Only official signatories can sign as organisation

(•) Individual ( ) Organization

First name | Last name

Email

Home Address

City | State / Province / County | Country

Submit

Accessing Modern Datasets – Command Line Interfaces

Modern datasets are HUGE!

- EPIC-KITCHENS ~ 796 GB
- EGO4D ~ 30+ TB

https://github.com/epic-kitchens/epic-kitchens-download-scripts

Download only certain data types

We provide videos, RGB/optical flow frames, GoPro's metadata (for the extension only) and object detection frames (for EPIC KITCHENS-55's videos only). You can also download the consent form templates.

If you want to download only one (or a subset) of the above, you can do so with the following self-explanatory arguments:

- `--videos`
- `--rgb-frames`
- `--flow-frames`
- `--object-detection-images`
- `--masks`
- `--metadata`
- `--consent-forms`

If you want to download only videos, then:

```
python epic_downloader.py --videos
```

Note that these arguments can be **combined** to download multiple things. For example:

```
python epic_downloader.py --rgb-frames --flow-frames
```

Will download both RGB and optical flow frames.

Specifying participants

You can use the argument `--participants` if you want to download data for only a subset of the participants. Participants can be specified with their numerical *or* string ID.

You can specify a single participant, e.g. `--participants 1` or `--participants P01` for participant `P01`, or a comma-separated list of them, e.g. `--participants 1,2,3` or `--participants P01,P02,P03` for participants `P01`, `P02` and `P03`

This argument can also be combined with the aforementioned arguments. For example:

```
python epic_downloader.py --videos --participants 1,2,3
```

Will download only videos from `P01`, `P02` and `P03`.

https://github.com/facebookresearch/Ego4d/tree/main/ego4d/cli

Data download

Canonical videos and annotations can be downloaded using the following command:

```
python -m ego4d.cli.cli --output_directory="~/ego4d_data" --datasets full_scale annotations --benchmarks FHO
```

v2.0 annotations can be downloaded with:

```
python -m ego4d.cli.cli --output_directory="~/ego4d_data" --datasets annotations --version v2
```

Detailed Flags

Flag Name	Description
`--dataset`	[Required] A list of identifiers to download: [annotations, full_scale, clips] Each dataset will be stored in folders in the output directory with the name of the dataset (e.g. output_dir/v2/full_scale/) and manifest.
`--output_directory`	[Required]A local path where the downloaded files and metadata will be stored
`--metadata`	[Optional] Download the primary `ego4d.json` metadata at the top level (Default: True)
`--benchmarks`	[Optional] A list of benchmarks to filter dataset downloads by - e.g. Narrations/EM/FHO/AV
`-y` `--yes`	[Optional] If this flag is set, then the CLI will not show a prompt asking the user to confirm the download. This is so that the tool can be used as part of shell scripts.
`--aws_profile_name`	[Optional] Defaults to "default". Specifies the AWS profile name from ~/.aws/credentials to use for the download
`--video_uids`	[Optional] List of video or clip UIDs to be downloaded. If not specified, all relevant UIDs will be downloaded.
`--video_uid_file`	[Optional] Path to a whitespace delimited file that contains a list of UIDs. Mutually exclusive with the `video_uids` flag.
`--universities`	[Optional] List of university IDs. If specified, only UIDs from the S3 buckets belonging to the listed universities will be downloaded.
`--version`	[Optional] A version identifier - e.g. "v1" or "v2" (default)
`--no-metadata`	[Optional] Bypass the `ego4d.json` metadata download
`--config`	[Optional] Local path to a config JSON file. If specified, the flags will be read from this file instead of the command line

Datasets

The following datasets are available (not exhaustive):

Dataset	Description
annotations	The full set of annotations for the majority of benchmarks.
full_scale	The full scale version of all videos. (Provide `benchmarks` or `video_uids` filters to reduce the 5TB download size.)
clips	Clips available for benchmark training tasks. (Provide `benchmarks` or `video_uids` filters to reduce the download size.)
video_540ss	The downscaled version of all videos - rescaled to 540px on the short side. (Provide `benchmarks` or `video_uids` filters to reduce the 5TB download size.)
annotations_540ss	The annotations corresponding to the downscaled `video_540ss` videos - primarily differing only in spatial annotations (e.g. bounding boxes).
3d	Annotations for the 3D VQ benchmark.
3d_scans	3D location scans for the 3D VQ benchmark.
3d_scan_keypoints	3D location scan keypoints for the 3D VQ benchmark.
imu	IMU data for the subset of videos available
slowfast8x8_r101_k400	Precomputed action features for the Slowfast 8x8 (R101) model
omnivore_video_swinl	Precomputed action features for the Omnivore Video model
omnivore_image_swinl	Precomputed action features for the Omnivore Image model
fut_loc	Images and annotations for the future locomotion benchmark.
av_models	Model checkpoints for the AV|Social benchmark.
lta_models	Model checkpoints for the Long Term Anticipation benchmark.
moments_models	Model checkpoints for the Moments benchmark.
nlq_models	Model checkpoints for the NLQ benchmark.
sta_models	Model checkpoints for the Short Term Anticipation benchmark.
vq2d_models	Model checkpoints for the 2D VQ benchmark.

EGO 4D

Command Line Interfaces Provided to Simplify Download

Lots of Tasks Out There – Challenges!

EPIC KITCHENS

ABOUT STATS DOWNLOADS CHALLENGES TEAM

EPIC-KITCHENS-100 2025 CHALLENGES

Challenge Details with links to ★NEW★ Codalab Leaderboards

leaderboards are now open for the challenge phase from Jan 2025.
In 2024, we have 7 open challenges. These are

- Semi-Supervised Video Object Segmentation Challenge
- EPIC-SOUNDS Audio-Based Interaction Recognition
- EPIC-SOUNDS Audio-Based Interaction Detection
- Action Recognition
- Action Detection
- UDA for Action Recognition
- Multi-Instance Retrieval

EPIC-Kitchens 2025 Challenges

Feb 1st 2025,	All leaderboards are open
May 19th 2025,	Server Submission Deadline at 00:00:00 UTC
May 23rd 2025,	Deadline for Submission of Technical Reports on CMT HERE
June 17 2024,	Results announced at 2nd EgoVis workshop in Nashville EgoVis@CVPR2025 workshop

Challenges Guidelines

The eight challenges below and their test sets and evaluation servers are available via CodaLab. The leaderboards will decide the winners for each individual challenge. For each challenge, the CodaLab server page details submission format and evaluation metrics.

This year, we offer four new challenges in: Semi-Supervised Video Object Segmentation using the VISOR annotations, Hand-object-segmentations using the VISOR annotations, single-object tracking and audio-based action recognition using the epic-sounds dataset.

To enter any of the nine competitions, you need to register an account for that challenge using a valid institute (university/company) email addres
To enable your account, fill this form with your team's details. A single registration per research team is allowed. We perform a manual check for

https://epic-kitchens.github.io/2025

Ego4D and EgoExo4D Challenge 2024

Ego4D and EgoExo4D Challenge 2024

Overview

At CVPR 2024, we will host **16** challenges including 2 new challenges (Goal Step and Ego Schema), representing each of Ego4D's five benchmarks. Included in the 16 challenges hosted at CVPR are two teaser Ego-Exo4D challenges (Ego-Pose Body and Ego-Pose Hands). Please find details below on the challenges:

Ego4D challenges

Episodic memory:

- Visual queries with 2D localization (VQ2D) and Visual Queries 3D localization (VQ3D): Given an egocentric video clip and an image crop depicting the query object, return the most recent occurrence of the object in the input video, in terms of contiguous bounding boxes (2D + temporal localization) or the 3D

https://ego4d-data.org/docs/challenge/

Challenges – Train/Val/Test scheme

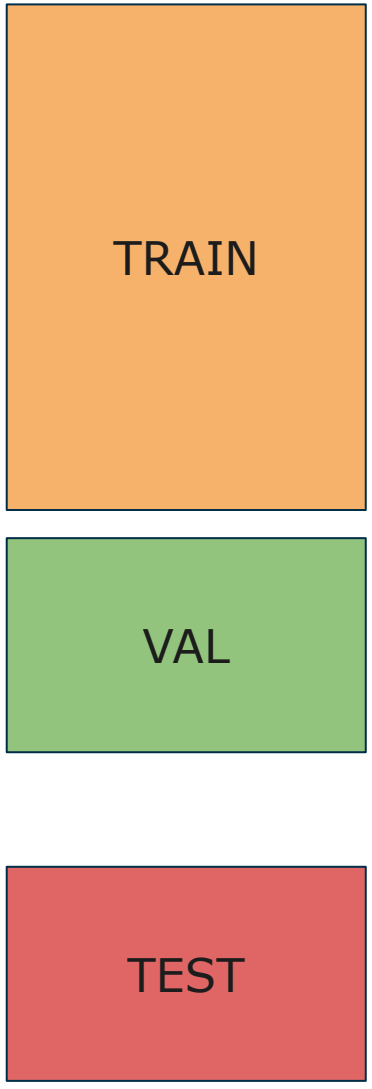

- Datasets are usually divided into train/val/test splits;
- All videos are publicly released;
- Train annotations are publicly released and meant for training models for the different challenges;
- Val annotations are publicly released and meant for model development and hyperparameter search;
- Test annotations are private and meant for assessing the performance of models avoiding bias in model design and optimization;
- Hence, the only way to obtain results on the test set is to send model predictions to an evaluation server.

Challenges – Evaluation Server

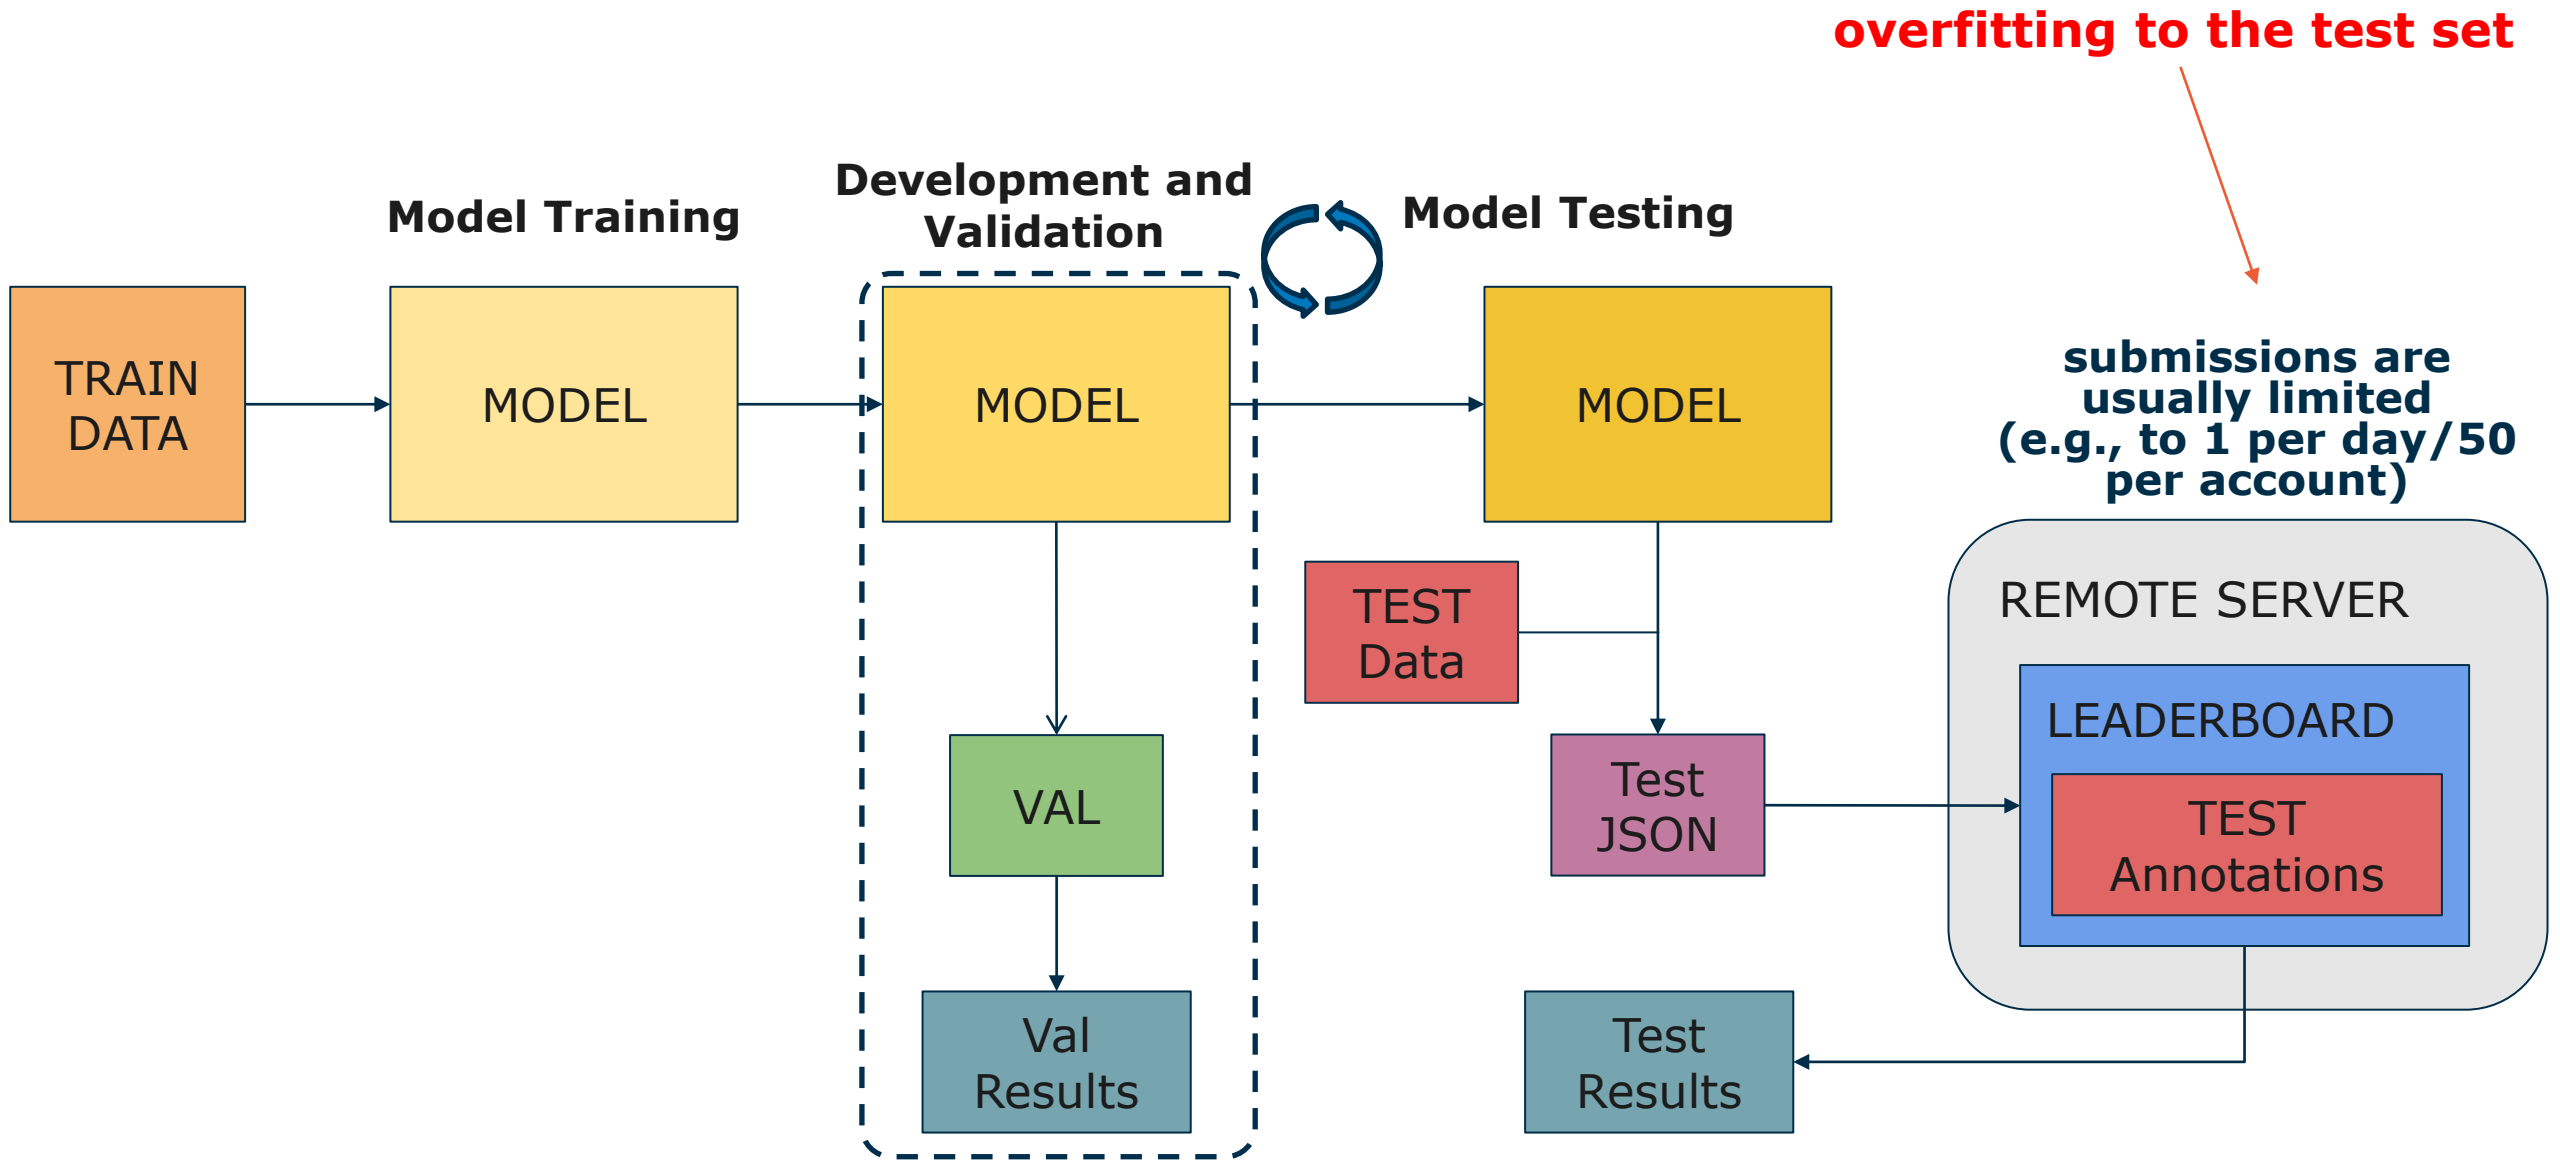

Challenges – Evaluation Server

EPIC-KITCHENS-100 Action Detection

Organized by antonino - Current server time: Feb. 23, 2025, 7:28 p.m. UTC

First phase	End
2024 Open Testing	Competition Ends
Oct. 1, 2024, 8 a.m. UTC	Jan. 31, 2025, midnight UTC

Test Set (Mean Average Precision - mAP)

#	User	Entries	Date of Last Entry	Team Name	SLS			mAP@0.1 (%)			mAP@0.2 (%)			mAP@0.3 (%)			mAP@0.4 (%)			mAP@0.5 (%)			Avg. mAP (%)		
					PT	TL	TD	Verb	Noun	Action	Verb	Noun	Action	Verb	Noun	Action	Verb	Noun	Action	Verb	Noun	Action	Verb	Noun	Action
1	shuming	12	05/02/24	KAUST-4Paradigm-MoonshotAI-Nvidia	2.0 (2)	3.0 (1)	4.0 (1)	34.11 (1)	40.66 (1)	36.09 (1)	32.75 (1)	38.63 (1)	34.70 (1)	30.48 (1)	36.32 (1)	32.67 (1)	28.03 (1)	32.55 (1)	29.91 (1)	24.73 (1)	27.98 (1)	26.50 (1)	30.02 (1)	35.23 (1)	31.97 (1)
2	xyx	23	05/30/24	dg_team(deepglint)	3.0 (1)	3.0 (1)	4.0 (1)	30.28 (5)	33.73 (3)	29.29 (2)	29.27 (5)	32.39 (3)	28.33 (2)	27.32 (4)	30.44 (3)	26.80 (2)	25.32 (3)	27.48 (2)	24.77 (2)	22.14 (3)	23.74 (2)	22.06 (2)	26.87 (4)	29.56 (2)	26.25 (2)
3	TIM_method	1	04/06/24	Oxford+Bristol	2.0 (2)	3.0 (1)	3.0 (2)	32.14 (2)	34.88 (2)	28.13 (3)	30.01 (3)	32.99 (2)	26.74 (3)	27.84 (3)	30.57 (2)	25.01 (3)	25.24 (4)	26.60 (3)	22.29 (3)	20.37 (5)	21.78 (3)	18.86 (3)	27.12 (3)	29.36 (3)	24.21 (3)
4	mzs	5	05/27/23	mzs	2.0 (2)	3.0 (1)	4.0 (1)	31.01 (3)	30.32 (5)	25.54 (4)	30.04 (2)	28.76 (5)	24.54 (4)	28.01 (2)	27.20 (4)	23.16 (4)	25.44 (2)	24.28 (4)	21.04 (4)	22.32 (2)	20.74 (4)	18.35 (4)	27.36 (2)	26.26 (4)	22.52 (4)
5	lijun	18	06/01/22		2.0 (2)	3.0 (1)	4.0 (1)	30.67 (4)	30.96 (4)	24.57 (5)	29.40 (4)	29.36 (4)	23.50 (5)	26.81 (5)	26.78 (5)	21.94 (5)	24.34 (5)	23.27 (5)	19.65 (5)	20.51 (4)	18.80 (5)	16.74 (6)	26.35 (5)	25.83 (5)	21.28 (5)
6	tzzcl1	19	06/01/22	4Paradigm-UWMadison-NJU	2.0 (2)	3.0 (1)	4.0 (1)	26.97 (7)	28.61 (6)	23.90 (6)	25.91 (7)	27.14 (6)	22.98 (6)	24.21 (6)	24.92 (6)	21.37 (6)	21.77 (6)	22.14 (6)	19.57 (6)	18.47 (6)	18.69 (6)	16.94 (5)	23.47 (7)	24.30 (6)	20.95 (6)
7	Haniel	44	06/01/23	Bristol-MaVi	2.0 (2)	3.0 (1)	3.0 (2)	27.57 (6)	24.18 (8)	19.64 (7)	26.16 (6)	22.92 (8)	18.66 (7)	24.02 (7)	20.85 (8)	17.29 (7)	21.48 (7)	18.24 (8)	15.55 (7)	18.42 (7)	15.21 (8)	13.49 (7)	23.53 (6)	20.28 (8)	16.93 (7)
8	Alibaba-MMAI-Research	1	12/15/21	CVPR 2021 Challenges	2.0 (2)	3.0 (1)	3.0 (2)	22.77 (10)	26.44 (7)	18.76 (8)	22.01 (9)	24.55 (7)	17.73 (8)	19.63 (11)	22.30 (7)	16.26 (8)	17.81 (10)	19.82 (7)	14.91 (8)	14.65 (11)	16.25 (7)	12.87 (8)	19.37 (11)	21.87 (7)	16.11 (8)
9	ctai	7	05/30/23	ctai	2.0 (2)	3.0 (1)	3.0 (2)	23.37 (8)	21.21 (9)	17.09 (9)	22.57 (8)	20.22 (9)	16.49 (9)	21.67 (8)	19.00 (9)	15.71 (9)	19.51 (8)	17.01 (9)	14.24 (9)	16.87 (8)	13.94 (9)	12.38 (9)	20.80 (8)	18.28 (9)	15.18 (9)

https://codalab.lisn.upsaclay.fr/competitions/707

Ego4D Short Term Object Interaction Anticipation Challenge

★ 15

Organized by: Ego4D
Starts on: Oct 25, 2022 2:00:00 AM CET (GMT + 1:00)
Ends on: Jun 1, 2099 1:59:59 AM CET (GMT + 1:00)

Overview | Evaluation | Phases | Participate | Leaderboard | Discuss

Leaderboard

Overall Top-5 mAP

Phase: Test Phase, Split: Test Split

Order by metric

B - Baseline * - Private V - Verified

Visible Metrics

Rank	Participant team	Noun (↑)	Noun_Verb (↑)	Noun_TTC (↑)	Overall (↑)	Last submission at	Meta Attributes
1	123456ABCD (j)	31.08	16.18	12.41	7.21	8 months ago	View
2	Zarrio (IH_new)	33.50	17.26	11.77	6.75	8 months ago	View
3	ICL@SNU (YOLO + CLIP)	34.89	17.61	10.91	6.22	9 months ago	View
4	Language NAO (TransFusion, Ego4D v2)	30.09	13.58	10.39	5.41	1 year ago	View
5	PAVIS (GANO_v2)	25.67	13.60	9.02	5.16	2 years ago	View

https://eval.ai/web/challenges/challenge-page/1623/leaderboard/3910

Conclusion

- First Person Vision paves the way to a variety of user-centric applications;
- However, we are still missing solid building blocks related to fundamental problems of First Person Vision such as action recognition, object detection, action anticipation and human-object interaction detection;
- Consumer devices are starting to appear, but the near future of First Person Vision is in focused applications such as the ones in industrial scenarios.

We Are Hiring!

francesco.ragusa@unict.it – fragusa@nextvisionab.it

THANK YOU!

Egocentric Vision: Exploring User-Centric Perspectives

Francesco Ragusa

LIVE Group @ UNICT - https://iplab.dmi.unict.it/live/

Next Vision - http://www.nextvisionlab.it/

Department of Mathematics and Computer Science - University of Catania

francesco.ragusa@unict.it - https://francescoragusa.github.io/