Object Instance Identification in Dynamic Environments

Takuma Yagi, Md. Tasnimul Hasan, and Yoichi Sato
Industrial Institute of Science, The University of Tokyo
{tyagi,tasnim,ysato}@iis.u-tokyo.ac.jp

Abstract

We study the problem of identifying object instances in a dynamic environment where people interact with the objects. In such an environment, objects' appearance changes dynamically by interaction with other entities, occlusion by hands, background change, etc. This leads to a larger intra-instance variation of appearance than in static environments. To discover the challenges in this setting, we newly built a benchmark of more than 1,500 instances built on the EPIC-KITCHENS dataset which includes natural activities and conducted an extensive analysis of it. Experimental results suggest that (i) robustness against instance-specific appearance change (ii) integration of low-level (e.g., color, texture) and high-level (e.g., object category) features (iii) foreground feature selection on overlapping objects are required for further improvement.

1. Introduction

The ability to identify whether an object appearing in an image is the same instance as previously seen is one of the essential abilities to understand physical environments over time. For example, when a robot is instructed by the user to "bring my phone", the robot needs to identify *the phone* that has observed in the past from its appearance. This object instance identification task offers several applications such as product search, robot control, and assistive technology.

Previously, recognition of object instances has been primarily studied in a static environment where the appearance of the object does not change across time. However, in the real-world, objects are used for their intended use, and their appearance change through interactions. For example, a mug may be filled with liquids or washed by hands during its use, causing significant foreground and background clutters (Figure 1, top row). Although identifying instances in such a *dynamic environment* is demanded, challenges in this situation has not been seriously discussed yet.

We analyze the problem of object instance identification under dynamic environment where people interact with objects. To this end, we contribute a new annotation built on
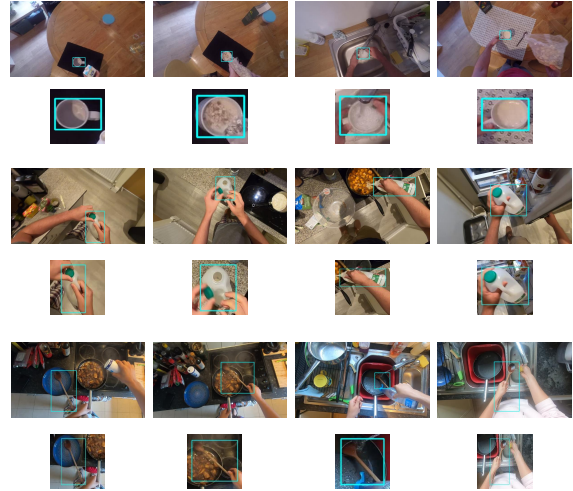

Figure 1. Examples of EK-Instance dataset. Each row denotes images from same instance.

the large-scale EPIC-KITCHENS dataset [3] that contains natural cooking activities. We choose first-person videos as a subject because it contains realistic hand-tool use and habitual activities across times. The new EK-Instance dataset (Figure 1) consists of more than 1,500 object instances, and contains various clutter in the foreground (*e.g.*, pouring drinks) and background (*e.g.*, overlap with other objects).

To discover the challenges in this setting, we introduce baseline methods using the off-the-shelf metric learning and clustering, and conduct (1) performance evaluation in EK-Instance dataset (2) cross-dataset transfer for analysis. The results suggest that (i) robustness against instance-specific appearance change (ii) integration of low-level and high-level features (iii) foreground feature selection are required for further improvement.

2. Related Work

Instance recognition/detection in a monotonous background [8] and re-identification in a static environment [2] have been studied, aiming for applications such as product recognition and robot control. Numbers of instance-

(a) Identification by appearance

(b) Single object per bounding box

(c) Unshared view
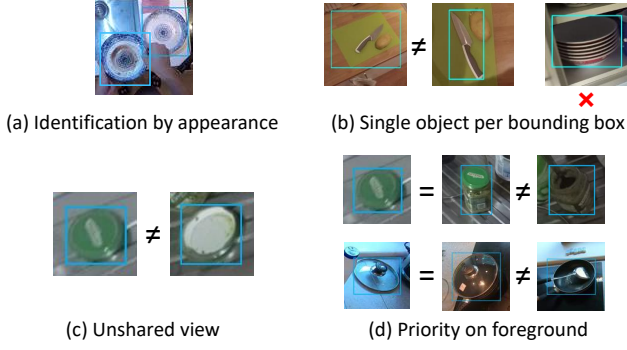
(d) Priority on foreground

Figure 2. Criteria for determining an instance.

level datasets have been proposed in parallel. Their source is roughly divided into (i) recording under controlled environments and (ii) Internet images. In the first group, at most, a few hundred instances were collected on a turntable or by hand-held. Robustness against pose change and lighting conditions is their primary interests [18]. Except for a few datasets that intentionally collected varied backgrounds [12, 17], most datasets are captured in a controlled environment where objects are captured without interaction with monotonous background. The second group collects data from E-Commerce sites for product recognition. Compared to the former, it is easier to scale up the number of instances to tens of thousands of instances (*e.g.*, [13]). However, for their original purpose, objects were shot with a simple background without real-world clutters. Different from the above, we collect a large number of object instances that appear in natural environments with interaction.

3. EK-Instance Dataset

3.1. Definition of Instance

We first clarify what *an instance* refers to. Compared to the previous works that captured objects apart from clutters, real-world environments include corner cases caused by interaction and it is not trivial how to determine the same instance. We propose four rules for consistent annotation.

(a) Identification by appearance: We define an instance as *an object which has a unique appearance that is distinguishable from other objects*. Visually identical objects (*e.g.*, set of tableware) are counted as a single instance.

(b) Single object per bounding box: We use a bounding box that surrounds the object just enough on determining the target object. When part of an object is occluded, we just cover the visible region. If a bounding box surrounds multiple objects (*e.g.*, stacked dishes), it is considered invalid and excluded from the evaluation (see Figure 2 (b)).

(c) Exception on unshared view: Ideally, a model should identify an object as the same even if it is viewed from different directions. However, as shown in Figure 2

(c), in some cases, it will be difficult to do so when there is no common part across views, such as the colors of the front and back of the lid being completely different. We exceptionally count as different instances in such a situation.

(d) Priority on foreground: Some objects can be separated into multiple parts. For example, a frying pan can be separated into a body and a lid. If we look at it from a top-down view, it will be impossible to distinguish between (i) a set of body and lid and (ii) the lid itself. While multiple interpretations are possible, we give priority to the part on the foreground and count as the same if objects share the foreground. In Figure 2 (d), we consider (i) and (ii) as the same while counting the body alone as a different instance.

3.2. Annotation Procedure

Based on the aforementioned criteria, annotations were made on the EPIC-KITCHENS dataset [3] which includes continuous cooking activities across dates. Because it is too labor-intensive to manually annotate instances across long videos, we adopt a semi-automatic procedure for annotation. Specifically, a category-agnostic object detector [9] and an object tracker are used to automatically extract and track the objects. Then we manually inspect whether the obtained trajectories contain a valid object and the same instance as other tracks. Finally, we sample detections by 1 fps and corrected the bounding boxes.

As a result, we collected 1,554 object instances, 38,479 tracks, and 92,376 frames from the 20 hours of video from 23 camera wearers in the EPIC-KITCHENS dataset. It consisted of containers (*e.g.*, bottle), cooking tools (*e.g.*, frying pan and paper roll), tableware (*e.g.*, knife), food (*e.g.*, apple), and electronics (*e.g.*, toaster). It also included objects which are difficult to define in a clear category (*e.g.*, fridge magnet, portafilter, and product package).

4. Experiments

Using the EK-Instance dataset, we evaluated the instance identification task in dynamic environments by building baseline models and comparing them across datasets.

4.1. Problem Statement

The problem is formulated as clustering of image tracks to a group of instances. Given a set of image tracks $\{X_1, \ldots, X_N\}$, we assign cluster labels $\{y_1, \ldots, y_N\}(y_i \in \{1, \ldots, K\})$, which is grouped by instance. In a novel environment, the true number of instances will not given in advance. Therefore, we formulated the task as clustering.

4.2. Image Track Encoder

We use ResNet-34 [7] pre-trained with the ImageNet dataset [4] as a backbone network to extract frame-level features. We pass the images to the backbone layers be-

Model	Clustering Method	AMI	ACC	F_P	F_B
ImageNet	K-means	0.755	0.590	0.547	0.628
	HAC	0.849	0.735	0.724	0.770
ArcFace	K-Means	0.794	0.631	0.606	0.670
	HAC	0.892	0.796	0.784	0.827
N-pair	K-Means	0.834	0.713	0.646	0.724
	HAC	**0.924**	**0.854**	**0.847**	**0.874**

Table 1. Performance on EK-Instance Dataset.

fore the final average pooling layer and obtain a 512-dimensional feature vector for each image. Next, they are averaged among images to aggregate image-level features into a track-level feature vector. Finally, we pass it to a single fully-connected layer to obtain a 256-dimensional track embedding. We used two loss functions for comparison:

Additive angular margin (ArcFace) [5]: Loss based on classification. Dicriminative features are learned by giving a penalty that the feature vectors between different classes are separated from each other by a certain amount or more.

N-pair [15]: Loss based on sample-wise comparison. Increase the similarity of pairs of feature vectors in the same class while decreasing the similarity of different classes.

In addition, we evaluated the **ImageNet** baseline which extracts 512-dimensional feature vectors from ImageNet pre-trained model as a strong baseline model.

4.3. Clustering and Evaluation

We performed clustering using the embeddings from the trained models. We compared Spherical K-means [6] and Hierarchical Agglomerative Clustering (HAC) [11]. Cosine similarity was used as a similarity measure in both algorithms. In Spherical K-means, the true number of clusters was given while in HAC the threshold was adjusted by the validation set. We used Adjusted Mutual Information (AMI) [16], unsupervised accuracy rate (ACC), Paired (F_P) and BCubed F-score (F_B) [1] as metrics.

4.4. Evaluation on EK-Instance Dataset

Table 1 shows the performance in the EK-Instance dataset. While ImageNet without fine-tuning showed moderate performance, the fine-tuned model in the training set showed better performance. In particular, the combination of N-pair loss and HAC showed the best performance. The reason why HAC exceeded K-means was that the K-means method tended to generate relatively even clusters, while the actual frequency distribution was biased.

Figure 3 shows an example of clusters extracted by the combination of N-pair and HAC. The model was not only robust to changes in the pose of the object and occlusion by hand (upper) but was also able to absorb foreground clutter such as putting a plastic bag on the trash can (lower).

However, some failure cases were seen as shown in Fig-

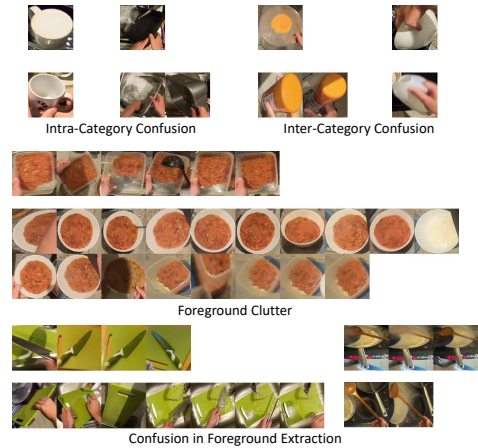Figure 3. Example of obtained clusters.

Figure 4. Failed HAC merging attepnts. First and second row of each item show tracks that are wrongly merged during HAC.

ure 4. First, there was intra/inter-category confusion with similar colors and textures. In particular, there were cases where completely different types of objects were confused by superficial features. As a unique failure in dynamic environments, confusion due to severe foreground clutter was seen when the same meal was served in another container as in the second line. Moreover, when multiple objects overlap, the model misunderstood the target object specified by a bounding box (*e.g.*, a cutting board and a knife on it).

4.5. Cross-Dataset Transfer

To investigate the properties of the learned feature representations in each dataset, we evaluated the performance when a model trained in one dataset was transferred to another dataset. We used Stanford Online Products (SOP) [13] as a dataset using internet images, INSTRE [17] and CORe50 [12] as a dataset containing instances on various backgrounds. CORe50 was used only for evaluation.

As shown in Table 2, the models trained by SOP and INSTRE did not generalize to EK-Instance (worse than ImageNet), while the model trained by EK-Instance showed improvement in SOP and CORe50, showing that it can respond to both foreground and background clutter.

Dataset	EK-Instance			Online Products			INSTRE			CORe50		
	AMI	F_P	F_B	AMI	F_P	F_B	AMI	F_P	F_B	AMI	F_P	F_B
ImageNet	0.849	0.724	0.770	0.332	0.023	0.434	0.944	0.836	0.890	0.497	0.311	0.446
EK-Instance	0.924	0.847	0.874	0.358	0.360	0.448	0.930	0.837	0.884	0.578	0.415	0.532
Online Products	0.800	0.620	0.694	0.649	0.321	0.633	0.870	0.756	0.773	0.090	0.063	0.216
INSTRE	0.795	0.634	0.703	0.304	0.047	0.445	0.984	0.963	0.972	0.667	0.440	0.599

Table 2. Results of cross-dataset transfer. Left column shows source dataset and top row shows target dataset. Numbers in bold and blue show best and second best results across source datasets, respectively.

5. Discussion

We draw the following suggestions: **(i) Robustness against instance-specific appearance change is needed:** The newly constructed EK-Instance dataset showed that the appearance of an object changes across interactions in a dynamic environment. Towards practical object instance identification, it is necessary to acquire not only viewpoint invariance but also robustness to changes in appearance due to interaction. **(ii) Both low-level/high-level features are important:** There was confusion between instances with similar low-level features such as color and texture but from completely different categories. We suspect that the category-level information was lost during training. A more robust approach is needed that uses both low-level and high-level features as clues for identification. **(iii) Selective foreground feature extraction is demanded:** In this work, bounding boxes were used to indicate the target object. However, when instances intersect with each other, confusion in determining the target object was observed. An explicit method to selectively filter features of the target object based on contextual information will be demanded.

6. Conclusion

In this study, we tackled the problem of object instance identification in a dynamic environment with human-object interaction. We proposed the EK-Instance dataset to analyze its difficulty. Specifically, we formulated the problem as clustering of image tracks and conducted an extensive analysis by within/cross-dataset evaluation. The analyses have revealed the existing biases in the trained model, and have discovered challenges for further development such as robustness against appearance change, feature fusion, and selective foreground feature extraction.

References

[1] Enrique Amigó, Julio Gonzalo, Javier Artiles, and Felisa Verdejo. A comparison of extrinsic clustering evaluation metrics based on formal constraints. *Information Retrieval*, 12(4):461–486, 2009. 3, 6

[2] Vaibhav Bansal, Gian Luca Foresti, and Niki Martinel. Where did i see it? object instance re-identification with attention. In *Proc. ICCVW*, pages 298–306, 2021. 1

[3] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, , Antonino Furnari, Jian Ma, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Rescaling egocentric vision. *IJCV*, 2021. 1, 2, 5

[4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proc. CVPR*, pages 248–255, 2009. 2

[5] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proc. CVPR*, pages 4690–4699, 2019. 3, 5

[6] Inderjit S Dhillon and Dharmendra S Modha. Concept decompositions for large sparse text data using clustering. *Machine Learning*, 42(1):143–175, 2001. 3

[7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. CVPR*, pages 770–778, 2016. 2

[8] David Held, Sebastian Thrun, and Silvio Savarese. Robust single-view instance recognition. In *Proc. ICRA*, pages 2152–2159, 2016. 1

[9] Dahun Kim, Tsung-Yi Lin, Anelia Angelova, In So Kweon, and Weicheng Kuo. Learning open-world object proposals without learning to classify. *RA-L*, 2022. 2

[10] Harold W Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955. 6

[11] Godfrey N Lance and William Thomas Williams. A general theory of classificatory sorting strategies: 1. hierarchical systems. *The Computer Journal*, 9(4):373–380, 1967. 3

[12] Vincenzo Lomonaco and Davide Maltoni. Core50: a new dataset and benchmark for continuous object recognition. In *Proc. CoRL*, pages 17–26, 2017. 2, 3, 6

[13] Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. Deep metric learning via lifted structured feature embedding. In *Proc. CVPR*, pages 4004–4012, 2016. 2, 3, 6

[14] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 618–626, 2017. 6

[15] Kihyuk Sohn. Improved deep metric learning with multiclass n-pair loss objective. In *Proc. NIPS*, pages 1857–1865, 2016. 3, 5

[16] Nguyen Xuan Vinh, Julien Epps, and James Bailey. Information theoretic measures for clusterings comparison: Variants,

properties, normalization and correction for chance. *JMLR*, 11:2837–2854, 2010. 3, 5

[17] Shuang Wang and Shuqiang Jiang. Instre: a new benchmark for instance-level object retrieval and recognition. *TOMM*, 11(3):1–21, 2015. 2, 3, 6

[18] Xiaohan Wang, Tengyu Ma, James Ainooson, Seunghwan Cha, Xiaotian Wang, Azhar Molla, and Maithilee Kunda. The toybox dataset of egocentric visual object transformations. *Proc. CVPRW*, 2018. 2

[19] Andrew Zhai and Hao-Yu Wu. Classification is a strong baseline for deep metric learning. In *BMVC*, 2019. 5

A. Details on EK-Instance Dataset

A.1. Statistics

As mentioned in the main text, we collected 1,554 instances, 38,479 tracks, and 92,376 frames from 23 participants in total. Following the split used in the EPIC-KITCHENS dataset [3], we split the dataset into training, validation, and test set. Concretely, we used instances from P33, P34, and P36 for validation and P09, P11, P18, and P32 for testing. Table 3 shows the statistics of each split. Figure 5 and 6 show the images of the instances included in validation and test set, respectively. By collecting more than 1,500 instances, we were able to train a metric learning-based neural network model. In addition, the subject-based split allowed us to evaluate the instance identification performance in novel environments.

The frequency of the instances is highly biased as it reflects the occurrence rate of objects in real-world videos. Figure 7 and 8 shows the frequency of the number of images per track and the number of tracks per instance, respectively. It can be seen that most instances appear a small number of times in a very short time, whereas a very small number of instances appear a long time and a large number of times.

A.2. Additional Examples

Figure 9 shows additional examples on the instances included in the EK-Instance dataset. Not only changes in the posture of the object and the lighting environment, but also various changes such as disturbances caused by hands and other objects, changes in appearance due to operation (*e.g.*, using the contents, peeling off the packaging), changes in the surrounding background, etc. are within the same instance. It is characteristic that it occurs in.

B. Details on Baseline Methods

B.1. Normalized Softmax (N-Softmax) Loss [19]

Normalized softmax loss is one of the simplest loss function for metric learning, extending the cross-entropy loss for classification task. It is almost the same as the cross-entropy

loss except the embedded vector and the weight vector of the final layer are L2 normalized.

$$L_{\text{nsoftmax}} = -\frac{1}{N} \sum_{i=1}^{N} \log \frac{e^{W_{y_i}^{\mathsf{T}} \mathbf{z}_i'/\tau}}{\sum_{k=1}^{K} e^{W_k^{\mathsf{T}} \mathbf{z}_i'/\tau}},$$

where $\mathbf{z}_i' = \frac{\mathbf{z}_i}{\|\mathbf{z}_i\|}$ is a vector by applying L2-normalization to $\mathbf{z}_i$ and τ denotes a temperature parameter that adjusts the scale of the numerator and denominator. We set $\tau = 0.2$.

B.2. Additive Angular Margin (ArcFace) Loss [5]

ArcFace loss is described as follows:

$$L_{\text{arcface}} = -\frac{1}{N} \sum_{i=1}^{N} \log \frac{e^{\cos(\theta_{y_i}+m)/\tau}}{e^{\cos(\theta_{y_i}+m)/\tau} + \sum_{k=1, k\neq y_i}^{K} e^{\cos\theta_j/\tau}},$$

where θ_k denotes the angle between L2-normalized weight W_k and a feature vector $\mathbf{z}_i'$ which belongs to class y_i. m denotes a margin term to increase the inter-class distance. This loss function tries to maintain a feature vector to have a margin of m against all the feature vectors that belongs to a different class. As a result, feature vectors of a same class will be distributed in a narrower space. We set $m = 0.5$ and $s = 1/30$.

B.3. N-pair Loss [15]

Different from the classification-based losses above, the n-pair loss tries to increase the similarity between feature vectors from the same class while decrease the similarity between feature vectors from different classes. Especially, the n-pair loss enables efficient training by taking all the possible pairs within the batch, without using classification weight W.

$$L_{\text{npair}} = -\frac{1}{N} \sum_{i=1}^{N} \log \frac{\sum_{j=1, y_j=y_i}^{N} e^{\mathbf{z}_i^{\mathsf{T}} \mathbf{z}_j/\tau}}{\sum_{j=1}^{N} e^{\mathbf{z}_i^{\mathsf{T}} \mathbf{z}_j/\tau}}.$$

We set $\tau = 0.07$.

C. Details on Evaluation

C.1. Data Split

We performed clustering within the images from the same participant and reported the average score among them.

C.2. Evaluation Metrics

We describe the four evaluation metrics in detail:

Adjusted Mutual Information (AMI) [16]: A metric based on mutual information. Given a set of ground truth clusters $C = \{c_i\}_{i=1}^{N}$ and predicted clusters $K = \{k_i\}_{i=1}^{N}$, we calculate AMI as follows:

$$AMI = \frac{I(C,K) - \mathbb{E}[I(C,K)]}{\frac{1}{2}[H(C) + H(K)] - \mathbb{E}[I(C,K)]},$$

Split	#participants	#instances	#tracks	#frames
Training	16	1051	27234	64300
Validation	3	193	4647	11851
Test	4	310	6598	16225
Total	23	1554	38479	92376

Table 3. Statistics of each split.

I and H denote mutual information and entropy, respectively. AMI will be 1.0 if the ground truth clusters and the predicted clusters are identical.

Unsupervised Accuracy (ACC): Given a set of ground truth clusters and a set of predicted clusters, the maximum accuracy can be achieved when a one-to-one correspondence between the two sets of clusters is made. Hungarian algorithm [10] is used to calculate the optimal assignment m.

$$ACC = \max_m \frac{\sum_{i=1}^{N} \mathbf{1}\{c_i = m(k_i)\}}{N}.$$

Paired F-score (F_P): Given N samples, takes all the $\frac{1}{2}N(N-1)$ combinations and counts them as correct if the relationship between the two samples is the same in predicted clusters and ground truth clusters.

BCubed F-measure (F_B) [1]: This metric calculates cluster-wise precision and recalls while F_P takes all the pairs regardless of the predicted clusters. Specifically, the *correctness* between samples i and j will be defined as follows:

$$Correctness(i,j) = \begin{cases} 1 & \text{if } L(i) = L(j) \text{ and } C(i) = C(j) \\ 0 & \text{otherwise,} \end{cases}$$

where $L(i)$ and $C(i)$ refers to the i-th predicted clusters and ground truth clusters. Finally, BCubed Precision and BCubed Recall will be calculated as follows:

$$BCubed \ Precision = \frac{1}{N} \sum_{i}^{N} \sum_{j \in C(i)} \frac{Correctness\ (i,j)}{|C(i)|},$$

$$BCubed \ Recall = \frac{1}{N} \sum_{i}^{N} \sum_{j \in L(i)} \frac{Correctness\ (i,j)}{|L(i)|}.$$

D. Details on Datasets

Stanford Online Products [13]: This dataset is composed of 120,053 images from 22,634 online product categories collected from eBay. Following the original split, we used 11,318 categories for training and 11,316 categories for testing, while the hyperparameter was tuned by further partitioning the training set for validation.

INSTRE [17]: This dataset consists of distinctive objects such as architecture, planar objects, and toys. The Internet images and manually collected images are mixed and divided into three subsets based on filming conditions. Notably, manually collected images were intentionally shot in 25 different backgrounds. For the evaluation, we used 100 instances from the INSTRE-S1 subset which contains 11,011 images, and 50 instances each are used for training and testing.

CORe50 [12]: This dataset contains 50 hand-held objects belonging to 10 categories such as scissors. Although the number of instances is as small as 50, we used as a target data set because it was intentionally photographed in various backgrounds with hand occlusion, which shares the difficulties in our EK-Instance dataset. Since this data set consists of videos for a total of 550 trials, images were extracted at 1 fps to compose a image track. Since the bounding box of the object attached to the data set was cut out looser than the target object, we crop 28 pixels at each of the four corners of the original 128-pixel square bounding box to form an input.

E. Visualization of Learned Features

To understand which region the model is focusing on, we applied Grad-CAM [14] and visualized the activation heatmap of each predicted instance (class). For simplicity, we used the model trained with N-softmax loss. Figure 10 shows the results of discriminative regions taken from the last layer of the ResNet-34 backbone.

First, we found that the model focused on the peripheral region of specific objects such as dishes and mugs (1-3rd row). Such unchanged regions may be discovered during the training since such objects typically have foreground clutters on their centers. Similarly, the model focused on regions other than hands that are not discriminative for instance identification (4th row). On the other hand, unreasonably sharp regions are generated in some instances even though the entire region contains unique patterns (5th and 6th row). Furthermore, as mentioned in the failure examples, the model focused on backgrounds in some instances, leaving room for improvement.
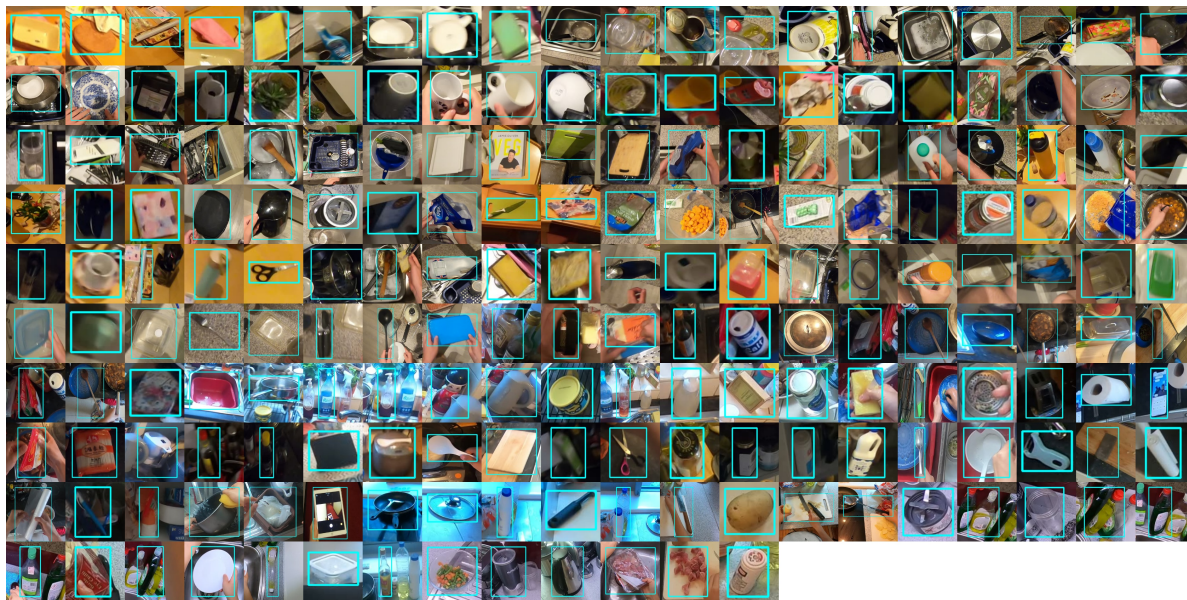
Figure 5. Images used in validation set (193 instances).

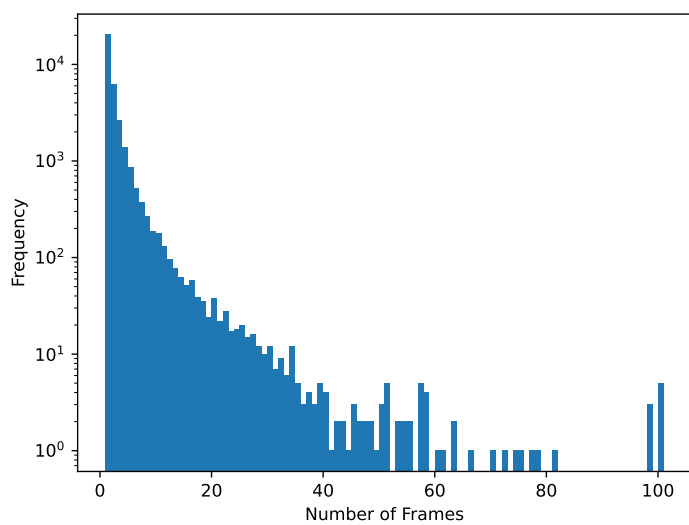Figure 6. Images used in test set (310 instances).
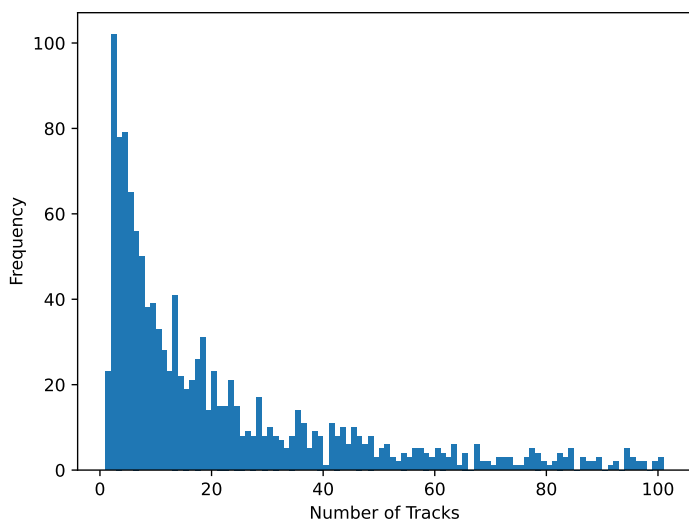
Figure 7. Distribution of number of images per track.

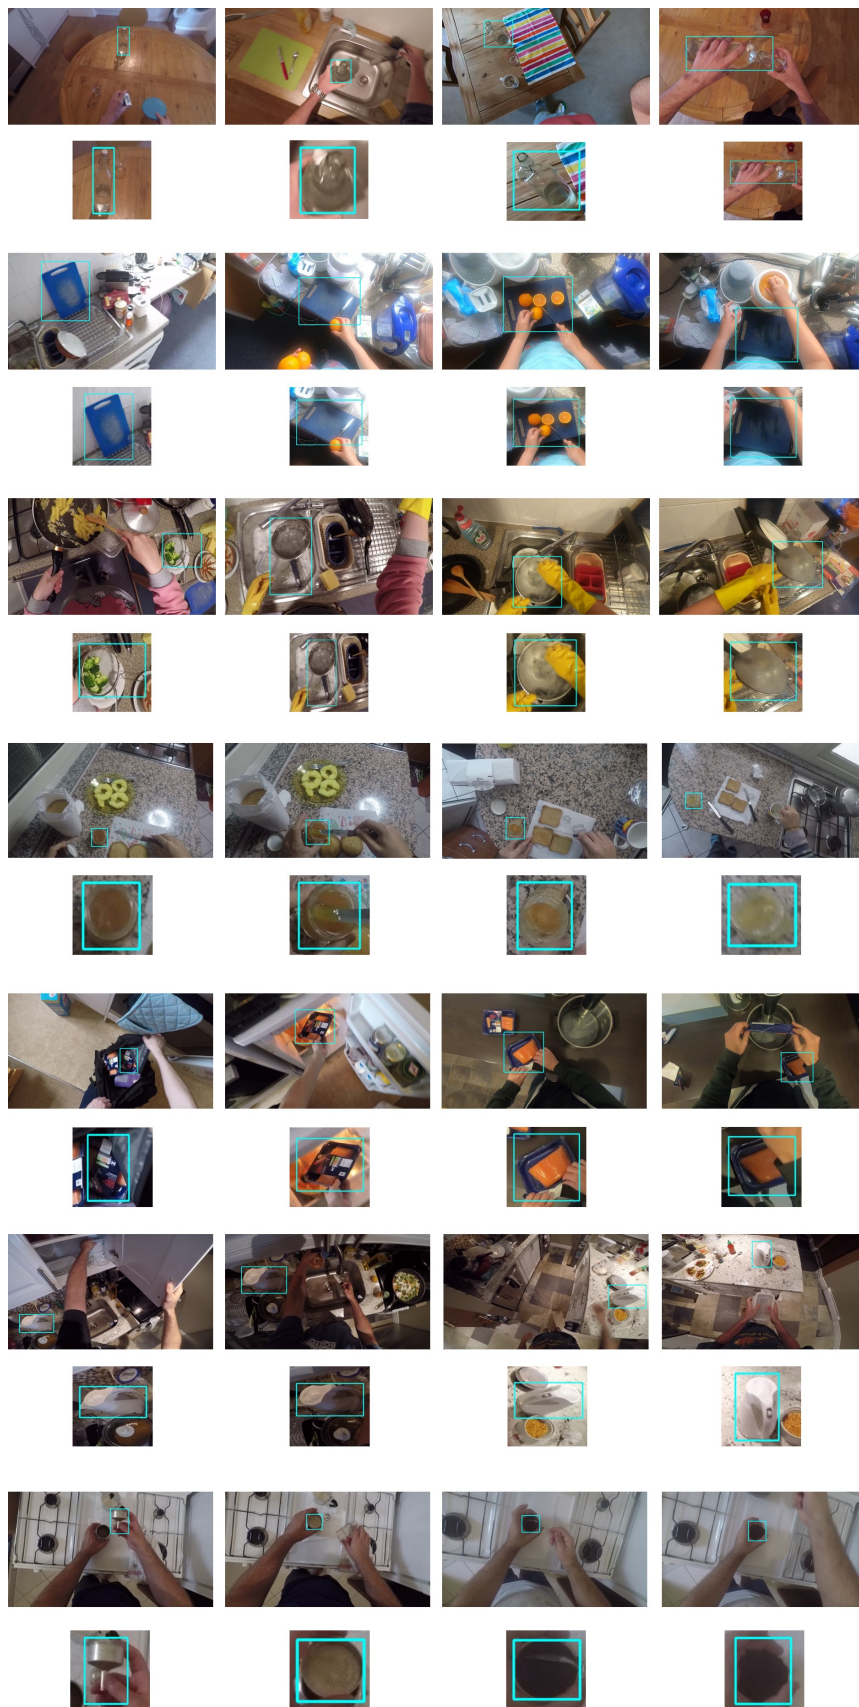Figure 8. Distribution of number of tracks per instance.

Figure 9. Additional examples of EK-Instance dataset. Each row refers to same instance.

Figure 10. Grad-CAM visualization on training set. First and second column are input image and its support for ground-truth category. Other columns are top-5 "hard" classes which showed high instance classification probability and their corresponding supports, respectively.