

# Fine-grained Affordance Annotation for Egocentric Hand-Object Interaction Videos

Zecheng Yu, Yifei Huang, Ryosuke Furuta, Takuma Yagi, Yusuke Goutsu, Yoichi Sato  
Industrial Institute of Science, The University of Tokyo

{zch-yu, hyf, furuta, tyagi, goutsu, ysato}@iis.u-tokyo.ac.jp

## Abstract

Object affordance is an important concept in hand-object interaction, providing information on action possibilities based on human motor capacity and objects' physical property thus benefiting tasks such as action anticipation and robot imitation learning. However, the definition of affordance in existing datasets often: 1) mix up affordance with object functionality; 2) confuse affordance with goal-related action; and 3) ignore human motor capacity. This paper proposes an efficient annotation scheme to address these issues by combining goal-irrelevant motor actions and grasp types as affordance labels and introducing the concept of mechanical action to represent the action possibilities between two objects. We provide new annotations by applying this scheme to the EPIC-KITCHENS dataset and test our annotation with tasks such as affordance recognition, hand-object interaction hotspots prediction, and cross-domain evaluation of affordance. The results show that models trained with our annotation can distinguish affordance from other concepts, predict fine-grained interaction possibilities on objects, and generalize through different domains.

## 1. Introduction

Affordance was first defined by James Gibson [9] as the possibilities of action that objects or environments offer. It is a non-declarative knowledge we have learned for automatically activating afforded responses on an object decided by both our motor capacity, *i.e.*, the motor actions suitable for human hands, and the object's physical properties such as shape. Recognizing affordance can benefit tasks like action anticipation and robot action planning by providing information on possible interactions with objects in the scene [15].

Many existing works [13, 17, 16, 20, 24, 18] in computer vision investigated affordance. They use verbs as affordance labels to describe the possible actions associ-




Figure 1. **Definitions used in our annotation scheme:** a) affordance labels are the combination of goal-irrelevant motor actions and grasp types that exists in hand-object interface; and b) mechanical action labels are action labels that can describe the interaction between two objects, exists in tool-object interface (figures adapted from [21]).

ated with objects. However, verbs like “cut”, “take”, and “turn off” do not correspond the definition of affordance. More specifically: a) “cut” is a possible action enabled by a knife, thus directly using “cut” as an affordance label fails to distinguish human natural motor capacities from the capacities extended by objects’ functionalities; b) using “take” as an affordance label overlooks changes in affordance when “take” is performed with different grasp types, which cannot provide fine-grained affordance annotations; and c) “turn-off” is a goal-related action, but not a goal-irrelevant affordance. The affordance utilized in “turn-off tap” should also apply in other interactions such as “press button”. Instead of the confusion of affordance and other concepts, verbs also can’t represent the diversity of affordance. For example, we cannot tell the differences between “pick-up bowl” and “pick-up knife” by the verb “pick-up”.

In light of these issues, we need a precise affordance definition to distinguish affordance from other concepts, with the capability of representing human motor capacity. Inspired by the findings of neuroscience [21], we address the shortcomings mentioned above by proposing an affordance annotation that considers hand-object interaction and tool-

object interaction separately. Our annotation scheme: a) defines affordance as a combination of goal-irrelevant motor actions and hand grasp labels. This can represent the possible motor actions enabled by human motor capacity and object’s physical property. This label structure can also provide a fine-grained affordance category to represent the diverse affordances in hand-object interactions; b) defines mechanical actions as the possible interactions between two objects, as shown in Figure 1.

Since annotating this information for a large-scale video dataset can be laborious, we propose an annotation method that leverages the consistency of affordance to simplify the annotation: the affordance will be the same when the same participant performs the same action on the same object. We apply this annotation scheme on the EPIC-KITCHENS dataset [5]. The annotations are available at <https://github.com/zch-yu/epic-affordance-annotation>.

We test our annotation on three tasks to evaluate its rationality. First we compare the recognition results of affordance, mechanical action, and tool / non-tool use action, the results show that models trained with our annotation clearly distinguish affordance from other concepts. Then we generate hand-object interaction hotspots by using affordance as weakly-supervision followed [18]’s method, our results show that the model can predict a more fine-grained, accurate interaction hotspots compared with using action labels. Finally we evaluate the generalization capacity of affordance, the performance of affordance recognition on unseen domains is much better than action recognition.

The main contributions of our work are as follows:

1. We point out the major shortcomings of existing affordance datasets, *i.e.*, affordance is wrongly confused with object functionalities and goal-related actions. In addition, verbs cannot completely describe affordance because of neglecting the grasp types.
2. We propose a fine-grained and efficient affordance annotation scheme for hand-object interaction videos to address the issues above, and provide annotations of affordance and other related concepts for a large-scale egocentric action video dataset: EPIC-KITCHENS.
3. We test our annotation on tasks such as affordance recognition, hand-object interaction hotspots prediction, and cross-domain evaluation of affordance. The results show that models trained with our annotation can distinguish affordance from other concepts, predict fine-grained interaction possibilities on objects, and generalize through different domains.

## 2. Related Works

### 2.1. Affordance Datasets

Earlier affordance datasets [20, 17] annotated possible actions and the exact regions where actions could occur for object images. Koppula *et al.* [13] provide affordance label annotation for human-object interaction video clips. Thermos *et al.* [24] and Fang *et al.* [6] annotate human-object interaction hotspot maps as object affordance for video clips associated with their action labels. Furthermore, Nagarajan *et al.* [18] use the action labels of the EPIC-KITCHENS dataset, which is egocentric, as weak supervision to learn to generate human-object interaction hotspot maps. As we can see in Table 2.1, these datasets all use verb as affordance label. They neither provide a clear definition of affordance nor consider humans’ motor capacity. And they also fail to represent the diversity of affordance. Therefore, we propose a fine-grained affordance annotation scheme considering both humans’ motor capacity and the object’s physical property for hand-object interaction videos.

### 2.2. Affordance Understanding

Affordance understanding methods can be divided into four categories: Affordance Recognition, Affordance Semantic Segmentation, Interaction Hotspots Prediction, and Affordance as Context. Given a set of images/videos, the task of affordance recognition [1] aims to estimate affordance labels from them. Affordance semantic segmentation [16, 20] aims at segmenting the input image / video frame into a set of regions that are labeled with affordance labels. Interaction hotspots prediction [18, 6] tries to predict the possible interaction hotspots of objects. Moreover, some works [14, 15, 19] also use affordance as a context for other tasks such as action anticipation.

All of these methods are influenced by simply using verbs as affordance labels. Firstly, verbs confuse affordance with other concepts such as object functionality. For example, our attention differs when observing the “cut” action and the “take” action [12, 11]: the former on the interacting object and the latter on the hand. Mixing them up may confuse models for affordance recognition. And verbs cannot represent the diverse affordances utilized in hand-object interaction. We may perform the same action with different affordances depending on the objects. For example, we may directly push or handle the doorknob to close a door. However, previous works overlook these details by simply using the action “close” as an affordance label. This leads to the failure to distinguish different affordance regions in affordance semantic segmentation and interaction hotspots prediction tasks.

	Format	Categories	Image/Video	Interaction Region	View	Affordance Labels
IIT-AFF [20]	RGB Image	9	8,835	✓	-	contain, cut, display, engine, grasp, hit, pound, support, w-grasp
RGB-D Part Affordance Dataset [17]	RGB-D Image	7	105	✓	-	grasp, cut, scoop, contain, pound, support, warp-grasp
CAD-120 [13]	RGB Video	6	130	×	third-person view	openable, cuttable, pourable, containable, supportable, holdable
SOR3D [24]	RGB-D Video	13	9,683	✓	third-person view	grasp, lift, push, rotate, open, hammer, cut, pour, squeeze, unlock, paint, write, type
OPRA [6]	RGB Video	7	11,505	✓	third-person view	hold, touch, rotate, push, pull, pick up, put down
EPIC-KITCHENS Affordance [18]	RGB Image	20	1,871	✓	egocentric	subset of EPIC-KITCHENS action set

Table 1. Comparative overview of existing affordance datasets

### 3. Proposed Affordance Annotation

Our goal is to develop a fine-grained and efficient annotation scheme for affordance and other related concepts for hand-object interaction videos.

#### 3.1. Definitions

Our proposed affordance annotation scheme is inspired by the three-action system model (3AS) [21], which clearly defines the concepts needed in hand-object interactions. The three-action system model includes affordance, mechanical action, and contextual relationship. We mainly focus on the first two concepts since they are closely related to our goal. Affordance is thus defined as

- **Hand-centered:** affordance only presence in the hand-object interface.
- **Animal-relative:** affordance is not only determined by the properties of objects but is also related to human motor capacity.
- **Goal-irrelevant object property:** the same affordance can be utilized for different purposes.

To fill the absence of affordance in object-object interactions, the 3AS [21] introduce mechanical actions as tool-centered, mechanical action possibilities between objects. According to the above, we separately consider the affordance and mechanical action for an instance of hand-object interaction as shown in Figure 1. We define affordance label as a combination of goal-irrelevant motor actions and grasp types. For mechanical actions, we use verbs that describe interactions between objects (such as *cut*, *stir*) as mechanical action labels.

#### 3.2. Annotation Scheme

To annotate an action video dataset, we first need to divide the original action labels of the dataset into tool-use actions and non-tool-use actions since mechanical actions are only present in tool-use actions. We then annotate mechanical actions for tool-use actions and affordances for both tool-use and non-tool-use actions. Directly annotating these labels can be laborious, thus we proposed an annotation scheme which utilizes existing annotations to reduce the difficulty of labeling.

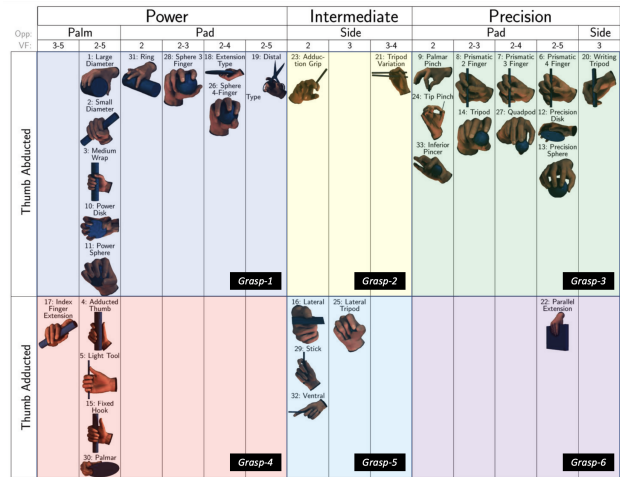


Figure 2. The 6-class grasp type taxonomy simplified from [8].




Figure 3. An affordance label is composed of a goal-irrelevant motor action label and a grasp type label.

**Tool-use / non-tool-use action annotation:** Tool-use / non-tool-use action annotation for action video datasets can be done by dividing original action labels of the dataset into three categories: tool-use action, non-tool-use action, and both, according to the meaning of each action label. For example, “take” is a non-tool-use action, while “cut” is a tool-use action. Some action labels could represent both the tool-use action and the non-tool-use action simultaneously, such as “wash”. We ignore these labels during annotation because of their ambiguity.

**Mechanical action annotation:** We only need to annotate mechanical actions for tool-use actions. According to the definition, we can use the verbs in original action labels as mechanical action labels. For example: in “stir food”, “stir” can represent the mechanical action between the slice

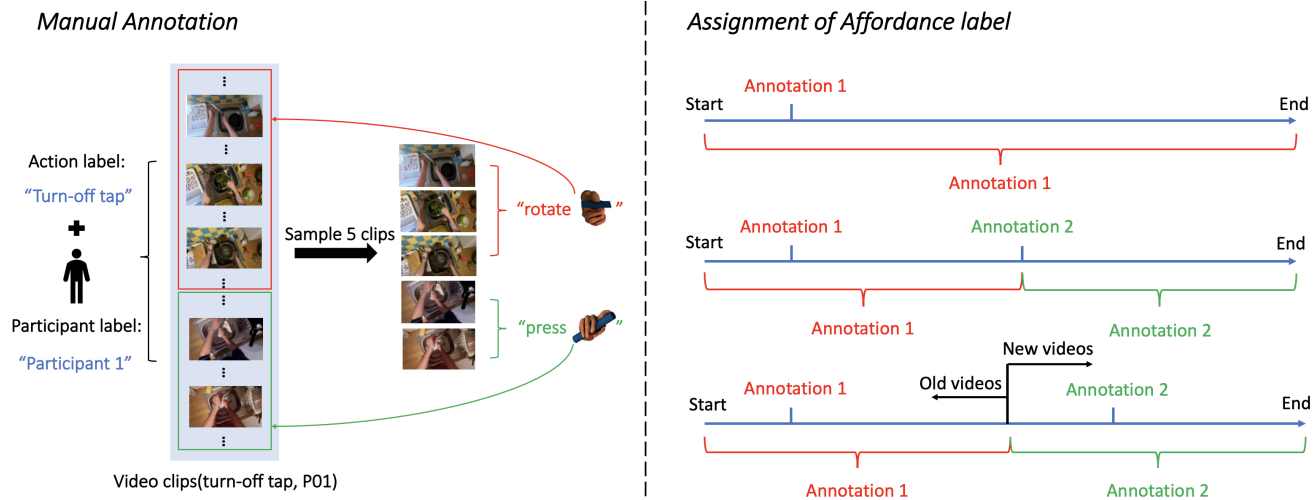


Figure 4. **The efficient annotation scheme.** **Left:** We first sample 5 video clips from the clips annotated with the same action-participant label, then manually annotate affordance label for them. **Right:** after manual annotation, the affordance labels are assigned to all video clips with the same action-participant label as follows: (a) If there is only one affordance label, then it is assigned to all videos; (b) If there are two affordance labels, we assign annotation 2 to video clips later than the video clip which is annotated with annotation 2, annotation 1 to video clips former than it; (c) If there is already a boundary which divides the video clips into groups, videos in each group are annotated with the annotation located in their group.

and the food. We can automatically annotate mechanical actions for all tool-use action video clips based on this rule, allowing significant reduction of annotation cost.

**Affordance annotation:** For affordance annotation, as shown in Figure 3, we annotate a goal-irrelevant motor action and a grasp type for each video clip. Given an unlabeled video clip, we first define a goal-irrelevant motor action according to the object property used in it. In the example of this figure, we use “pull” to represent the “pullable” property of the cupboard’s handle. Next, we chose a grasp type from a 6-class grasp type taxonomy. This taxonomy is simplified from a well-known 33-class grasp type taxonomy [8] based on the power of the grasp type and the posture of the thumb, as shown in Figure 2. Finally, we combine the goal-irrelevant motor action label with the grasp-type label as the affordance label. This form of affordance labels can model both the object’s physical property and the human motor capacity. And the combination of motor action and grasp type provides a fine-grained structure to represent diverse affordances.

To reduce the manual work of affordance annotation, we propose an efficient annotation method based on the assumption that the same person would interact with the same object in a fixed manner. As shown Figure 4(left), there are multiple video clips demonstrating the same participant performs the same action in the original dataset. We first sample five clips from video clips with the same action-participant (verb, noun, participant) annotation, then manually annotate affordances for them with the CVAT [22].

Then we assign these affordance labels to video clips with the same action-participant annotation, as shown in Figure 4. In some cases, videos with the same action-participant label have multiple affordance labels because of the scene variation (*e.g.*, the participant performs the same action in different rooms). To address this issue, our affordance annotation assignment scheme is as follows: (a) One affordance annotation: if there is no scene change among these video clips, we apply the only affordance annotation to all video clips with the same action-participant label. (b) Two scenes without a predefined boundary (*e.g.*, the boundary of EPIC-KITCHENS55 [4] and EPIC-KITCHENS100): we use affordance occurred later as a boundary, video clips earlier than it are annotated with annotation 1, those later than it are annotated with annotation 2. (c) Two scenes with a boundary: There exists a boundary which divides the video clips into two groups based on their relative position to the boundary. Videos of each group are annotated with the annotation inside their group. Note that there is a trade-off between annotation speed and accuracy. The more video clips we manually annotate the more accurate the automatic annotation results will be. The efficiency and accuracy of this method are shown in Section 4.1.

## 4. Experiments

We first annotate the EPIC-KITCHENS dataset with our proposed method and then train models on them. Our experiments evaluate the rationality of the annotation on four aspects: first, we test the recognition performance of tool use




Figure 5. Distribution of affordance classes

action / non-tool use action to demonstrate the importance of separating these action domains. Secondly, we compare the recognition models of affordance and mechanical action, evaluating the models’ capacity of distinguishing affordance from other concepts. Thirdly, we show their capability of representing the diverse affordances by following Nagarajan *et al.* [18]’s method. Finally, we compare the generalization ability of affordance and action.

#### 4.1. Dataset

We chose the EPIC-KITCHENS dataset among large scale video datasets [5, 3, 10], which contains egocentric hand-object interaction video clips annotated with action (verb, noun) labels, recorded by 21 participants in 45 kitchen scenes. We first annotate tool use action / non-tool use action for the dataset. There are 60 non-tool use actions and 33 tool use actions among the 97 verbs of EPIC-KITCHENS, which results in 51.5k non-tool use action video clips and 8.5k tool use action video clips, as shown in Table 2. Then the verb annotation of the tool use action video clips are used as their mechanical action labels, as shown in Figure 6. Next, using our annotation method, we sampled 91 most used action labels from the original annotation and manually annotated 1356 affordance labels for different action-participant pairs. We obtain 18, 613 video clips annotated with 60 affordance labels, as shown in Figure 5. For quality check, we randomly sample 1,113 instances from the annotated video clips and manually checked their affordance label, getting an accuracy of 87.95%, and then correct the affordance annotation with them.




Figure 6. Distribution of mechanical classes

Category	Video Clips	Action labels
Non-tool use	51.5k	take, put, open...(60)
Tool use	8.5k	cut, pour, mix, dry...(33)

Table 2. Tool-use / non-tool-use action labels of the EPIC-KITCHENS

#### 4.2. Identification of Affordance / Mechanical Action / Tool Use Action

To evaluate whether our annotation scheme can address the issue of confusing affordance with object functionality and goal-related actions. We train three recognition models for tool / non-tool use action, mechanical action, and affordance separately, and compare their visualization results to show the important regions each model focuses on.

**Tool Use / Non-Tool Use Action Recognition:** We train two SlowFast [7] models to recognize tool use action / non-tool use action from a given video clip, one with random

Dataset	Tool use actions	Non tool-use actions
Random annotation	0.4720	0.5282
Our annotation	0.8580	0.7867

Table 3. Tool-use / non-tool use action classification results.

	Top1 Acc	Top5 Acc
Affordance	0.5708	0.8771
Mechanical action	0.5190	0.8643

Table 4. Affordance / mechanical action recognition results

tool / non-tool use action annotation and another with our annotation. The results are shown in Table 3, which demonstrates the rationality of our automatic annotation scheme for tool / non-tool use actions.

**Mechanical Action Recognition:** We train a SlowFast model for mechanical action recognition with our annotation. For the 33-class mechanical action recognition task, as shown in Table 4, we get a recognition accuracy of 51.90%.

**Affordance Recognition:** We train a SlowFast model for affordance recognition with our annotation. For the 60-class affordance recognition task, as shown in Table 4, we get a recognition accuracy of 57.08%.

According to the three-action model, affordances are only present in the hand-object interface, and mechanical actions are only present between the object-object interface. Thus, interaction regions can help us distinguish affordance from object functionality. Here, we compare the visualization results of these models to see whether our proposed annotation scheme can separate affordance with other concepts. The visualization results generated by GradCam [23] are shown in Figure 7. From the first row, we can see that the affordance recognition model focuses more on the hand-object interaction. The second row shows that the mechanical action recognition model cares more about object interaction. When the second row is compared to the third row, it is evident that the mechanical action recognition model focuses on tool-object interactions (2nd row), and the tool use/non-tool use action classification model focuses on the existence of tools (3rd row). These results show that our affordance annotation provides a more precise affordance label than previous datasets, which is important for discriminating affordance with other concepts.

### 4.3. Hand-Object Interaction Hotspots

In this section, we evaluate our affordance annotation on the interaction hotspot prediction task, which generates an interaction hotspot map for a given object image. We follow the weakly-supervised method proposed in [18]. They first train an action recognition model with action labels and an




Figure 7. GradCam [23] visualization results of affordance recognition, mechanical action recognition, and tool-use/non-tool-use action recognition. The affordance recognition model focuses on the hand-object interaction, while the mechanical action recognition model cares more about object-object interaction. This demonstrate our proposed annotation labels clearly separate affordance with other concepts, and completely represent affordance without missing the human motor capacity.

anticipation network that can predict the “active state” of an inactive object. At the inference phase, the input image is fed to the anticipation network, followed by the action recognition model to obtain the action prediction. Then they derive gradient-weighted attention maps as the interaction hotspot map. In our experiment, we train the model with our affordance labels instead of action labels.

We first compare the ground-truth heatmaps to our affordance-supervised interaction hotspots and the action-supervised interaction hotspots. We train the affordance model with 4, 344 video clips including 43 different affordance labels, and the action model with 9, 236 video clips including 20 different action labels. Table 5 shows the results, we report the error as KL-Divergence [23], SIM, and AUC-J [2]. The affordance model outperforms the action model in these metrics, demonstrating that the model can better capture the interaction cues with the supervision of affordance labels. One reason is that the affordance label’s components: goal-irrelevant action and grasp type are more suitable for representing hand-object interactions. Furthermore, the granularity of affordance labels also helps the model avoid missing possible interactions.

In addition to the quantitative performance improvement, the predicted interaction heatmaps also benefit from the fine-grained affordance labels of our annotation. The interaction heatmaps generated by the action model and affordance model are shown in Figure 8, which highlight the interaction regions of different actions(affordances) on ob-

	KLD↓	SIM↑	AUC-J↑
Action label	2.661	0.382	0.758
Affordance label	1.305	0.399	0.776

Table 5. **Interaction hotspots prediction results generated using action model and affordance model.** The affordance model outperforms the action model on all metrics, demonstrating that our fine-grained affordance labels can help the model better capture the interaction cues.

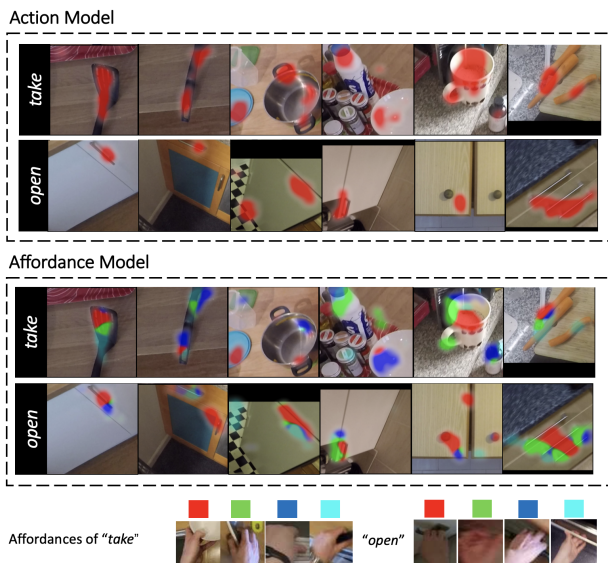


Figure 8. **Interaction heatmaps generated by action models and affordance models.** **Top rows:** Interaction heatmaps of *take*, *open*(red) generated using action model. **Bottom rows:** Interaction heatmaps of affordances(red, green, blue, cyan) used in *take*, *open* generated using affordance model. The affordance model trained with our annotation performs better on both capturing the right interaction region and representing the diversity of hand-object interaction.

jects. The top rows show that the action model can predict the possible interaction regions on different objects(*e.g.*, we can take the slice by grasping the red regions in the first image). Comparing these heatmaps with those in the bottom rows generated by the affordance model, we can see that the affordance heatmaps can better represent the diverse hand-object interactions(*e.g.*, the “take” action is replaced by four different affordances, representing different interaction ways of “taking”). This demonstrates our affordance annotation’s capacity to help the model learn a more granular representation of hand-object interaction.

#### 4.4. Generalization Ability of Affordance on Different Domains

In this section, we evaluate the generalization ability of the affordance recognition model. How well does the model



Figure 9. **Visualization results of the affordance recognition model(top) and the action recognition model(bottom).** Our affordance annotation guide the model to focus on the hand-object interaction which does not vary much with the domain. Thus, the learned affordance representation’s generalization ability is much better than the action representation.

trained on videos from one domain work on other domains? We followed the experiment setting of the unsupervised domain adaptation(UDA) challenge of the EPIC-KITCHENS dataset. First, we separate annotated video clips into source and target domains according to EPIC-KITCHENS’s UDA setting. We then train an affordance recognition model and an action recognition model in the source domain. Finally, the model which works the best on the source domain’s validation set is chosen to be evaluated on the target domain. Table 6 shows the comparison results of affordance recognition and action recognition. Here we use the same Slowfast model for both tasks. We can see that the performance drop of the affordance recognition model in the target domain is smaller than that of the action model, which demonstrates the generalization ability of the learned affordance representation.

Figure 9 shows the visualization results of the affordance recognition model(top), and the action recognition model(bottom). We can see that the affordance model focuses more on the hand-object interaction, but the action model cares more about the target object. The reason is that the affordance labels(*goal-irrelevant action*, *grasp type*) focus more on how we interact with the object, such as the appearance and motion of the hand, and the parts of the object to be interacted with. These do not vary much with different domains. However, the action labels(*verb*, *noun*) focus on the object and human-object interaction, which are difficult to be consistent through different domains. Objects with the same label may change a lot through domains(*e.g.*, knives in a different color, size.). Furthermore, human-object interactions are also changeable since different people may perform the same action differently. Thus, the affordance recognition model works better in the target domain.

Task	Target Top-1 Acc	Target Top-5 Acc	Source Top-1 Acc	Source Top-5 Acc
Action Recognition	12.67	30.99	16.16	32.91
Affordance Recognition	20.69	54.59	24.68	53.58

Table 6. **Cross-domain recognition results of action / affordance:**the affordance recognition models work better in the unseen target domain.

## 5. Conclusion

In this study, we proposed a fine-grained affordance annotation scheme for hand-object interaction videos, distinguishing affordance from other concepts like object functionality and reducing the manual annotation burden. We successfully applied our proposed annotation scheme to the EPIC-KITCHENS dataset and evaluated them on three tasks. The results of recognition tasks on affordance, mechanical action, and tool / non-tool use action show that our proposed annotation clearly distinguishes affordance from other concepts and completely represents affordance without missing human motor capacities. The interaction hotspots prediction results demonstrate that our fine-grained affordance labels can better represent the diverse hand-object interactions than verbs. Moreover, our affordance annotation also exhibits its generalization ability in different domains.

## Acknowledgments

This work is supported by JST AIP Acceleration Research Grant Number JPMJCR20U1 and JSPS KAKENHI Grant Numbers JP20H04205 and JP22K17905.



## References

- [1] Ryunosuke Azuma, Tetsuya Takiguchi, and Yasuo Arika. Estimation of Object Functions Using Visual Attention. page 4, 2018.
- [2] Zoya Bylinskii, Tilke Judd, Aude Oliva, Antonio Torralba, and Frédo Durand. What do different evaluation metrics tell us about saliency models? *IEEE transactions on pattern analysis and machine intelligence*, 41(3):740–757, 2018.
- [3] Joao Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *arXiv preprint arXiv:1907.06987*, 2019.
- [4] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The epic-kitchens dataset. In *European Conference on Computer Vision (ECCV)*, 2018.
- [5] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Evangelos Kazakos, Jian Ma, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Rescaling egocentric vision: collection, pipeline and challenges for epic-kitchens-100. *International Journal of Computer Vision*, 130(1):33–55, 2022.
- [6] Kuan Fang, Te-Lin Wu, Daniel Yang, Silvio Savarese, and Joseph J Lim. Demo2vec: Reasoning object affordances from online videos. In *CVPR*, pages 2139–2147, 2018.
- [7] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, pages 6202–6211, 2019.
- [8] Thomas Feix, Javier Romero, Heinz-Bodo Schmiedmayer, Aaron M Dollar, and Danica Kragic. The grasp taxonomy of human grasp types. *IEEE Transactions on human-machine systems*, 46(1):66–77, 2015.
- [9] James J Gibson. The concept of affordances. *Perceiving, acting, and knowing*, 1, 1977.
- [10] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The” something something” video database for learning and evaluating visual common sense. In *Proceedings of the IEEE international conference on computer vision*, pages 5842–5850, 2017.
- [11] Yifei Huang, Minjie Cai, Zhenqiang Li, Feng Lu, and Yoichi Sato. Mutual context network for jointly estimating egocentric gaze and action. *IEEE Transactions on Image Processing*, pages 7795–7806, 2020.
- [12] Yifei Huang, Minjie Cai, Zhenqiang Li, and Yoichi Sato. Predicting gaze in egocentric video by learning task-dependent attention transition. In *ECCV*, pages 754–769, 2018.
- [13] Hema Swetha Koppula, Rudhir Gupta, and Ashutosh Saxena. Learning human activities and object affordances from rgb-d videos. *The International Journal of Robotics Research*, 32(8):951–970, 2013.
- [14] Hema S Koppula and Ashutosh Saxena. Anticipating human activities using object affordances for reactive robotic response. *IEEE transactions on pattern analysis and machine intelligence*, 38(1):14–29, 2015.
- [15] Miao Liu, Siyu Tang, Yin Li, and James M Rehg. Forecasting human-object interaction: joint prediction of motor attention and actions in first person video. In *ECCV*, pages 704–721. Springer, 2020.
- [16] Timo Luddecke and Florentin Worgotter. Learning to segment affordances. In *ICCVW*, pages 769–776, 2017.
- [17] Austin Myers, Ching L Teo, Cornelia Fermüller, and Yiannis Aloimonos. Affordance detection of tool parts from geometric features. In *ICRA*, pages 1374–1381. IEEE, 2015.
- [18] Tushar Nagarajan, Christoph Feichtenhofer, and Kristen Grauman. Grounded human-object interaction hotspots from video. In *ICCV*, pages 8688–8697, 2019.
- [19] Tushar Nagarajan, Yanghao Li, Christoph Feichtenhofer, and Kristen Grauman. Ego-topo: Environment affordances from egocentric video. In *CVPR*, pages 163–172, 2020.
- [20] Anh Nguyen, Dimitrios Kanoulas, Darwin G Caldwell, and Nikos G Tsagarakis. Object-based affordances detection with convolutional neural networks and dense conditional random fields. In *IROS*, pages 5908–5915. IEEE, 2017.
- [21] François Osiurak, Yves Rossetti, and Arnaud Badets. What is an affordance? 40 years later. *Neuroscience & Biobehavioral Reviews*, 77:403–417, 2017.
- [22] Boris Sekachev, Nikita Manovich, Maxim Zhiltsov, Andrey Zhavoronkov, Dmitry Kalinin, Ben Hoff, TOsmanov, Dmitry Kruchinin, Artyom Zankevich, DmitriySidnev, Maksim Markelov, Johannes222, Mathis Chenuet, a andre, telenachos, Aleksandr Melnikov, Jijoong Kim, Liron Ilouz, Nikita Glazov, Priya4607, Rush Tehrani, Seungwon Jeong, Vladimir Skubriev, Sebastian Yonekura, vugia truong, zliang7, lizhming, and Tritin Truong. *opencv/cvat: v1.1.0*, Aug. 2020.
- [23] 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 *ICCV*, pages 618–626, 2017.
- [24] Spyridon Thermos, Georgios Th Papadopoulos, Petros Daras, and Gerasimos Potamianos. Deep affordance-grounded sensorimotor object recognition. In *CVPR*, pages 6167–6175, 2017.