

# Segmenting Object Affordances: Reproducibility and Sensitivity to Scale

Tommaso Apicella<sup>1,2</sup>, Alessio Xompero<sup>2</sup>, Paolo Gastaldo<sup>1</sup>, and  
Andrea Cavallaro<sup>3,4</sup>

<sup>1</sup> University of Genoa, Italy

<sup>2</sup> Queen Mary University of London, United Kingdom

<sup>3</sup> Idiap Research Institute, Switzerland

<sup>4</sup> École Polytechnique Fédérale de Lausanne, Switzerland

{t.apicella, a.xompero}@qmul.ac.uk

paolo.gastaldo@unige.it

a.cavallaro@idiap.ch

**Abstract.** Visual affordance segmentation identifies image regions of an object an agent can interact with. Existing methods re-use and adapt learning-based architectures for semantic segmentation to the affordance segmentation task and evaluate on small-size datasets. However, experimental setups are often not reproducible, thus leading to unfair and inconsistent comparisons. In this work, we benchmark these methods under a reproducible setup on two single objects scenarios, tabletop without occlusions and hand-held containers, to facilitate future comparisons. We include a version of a recent architecture, Mask2Former, re-trained for affordance segmentation and show that this model is the best-performing on most testing sets of both scenarios. Our analysis shows that models are not robust to scale variations when object resolutions differ from those in the training set.

**Keywords:** Affordances · Semantic segmentation · Benchmarking

## 1 Introduction

Affordances are the potential actions that objects of interest present in the environment offer to an agent (e.g., a person or a robot) [16]. This definition implies that an object can suggest two main interactions to accomplish an action: one between the agent and the object (agent-object) and another between the object and the environment (object-object) [41]. For example, to fill a cup (action), a person needs to first grasp the cup (agent-object) and then pour the filling into the cup (object-object). Visual affordance segmentation identifies the affordance regions of an object directly from an image [21, 37]. In the filling example, the external surface of the cup is a *graspable* region, whereas the internal part can *contain* the filling. Affordance segmentation is important for assistive robotic and human-robot collaboration [37, 40, 42, 44, 45, 58–60], action recognition [30] and anticipation [51], and augmented and virtual reality [20].

However, variations in the physical properties and appearance of the objects of interest as well as occlusions make the task challenging. Objects of the same type, such as drinking cups or knives, can vary in shape, color, and material [37, 40], or different

object types can be designed for different actions, but their perception in an image can result in visually similar regions. For example, the surface of a trowel to *scoop* with and the surface of a turner to *support* with. Occlusions can occur when other objects are present in front of the object of interest and with respect to the camera observing the object, or when the object is hand-held by a person (hand-occlusion) [1, 56]. These occlusions result in the object of interest only partially observed in the image and increase the difficulty of accurately identifying the visible affordance regions. Moreover, the recognition and segmentation of these regions is further affected by the perspective of the object in more challenging poses when hand-held by a person [1, 56].

To tackle these challenges, most of the existing methods independently assign an affordance class to each pixel and group regions with pixels belonging to the same class [26, 39, 40]. However, when the object is not close to the camera or not easily distinguishable from the background, these methods can misclassify areas outside of the object region. Other methods [17, 60, 65, 67] use attention mechanisms to select relevant pixel positions or relevant channels within low-resolution feature maps extracted from the image, while avoiding the high computational cost at higher resolutions. However, the coarse spatial information of low-resolution feature maps can affect the affordance segmentation, degrading the details of object regions. In the attention mechanism, latent vectors can replace feature maps as query, enabling higher resolutions through a more efficient computation [27]. Latent vectors, also referred to as object queries, are compact representations independent from the input and can be learned during training time. Latent vectors can be combined with the feature maps to predict specific object classes location [3] or segmentation masks [8, 9, 48]. Despite these advantages, no existing work modelled latent vectors for visual affordance segmentation. Moreover, a fair comparison with previous works is not always possible due to unavailable implementations or weights of the model, and performance is compared quantitatively in spite of changes in the training setup [14, 17, 39]. Previous works evaluated models using the benchmarks and comparing results, without considering how model could perform in a real scenario, where the camera might be closer or farther from the object of interest. In case of robotic experiments, the camera pose was fixed in a way that could resemble the distance of the object observed during training [6, 14, 39, 40, 59, 60].

In this paper, we provide an up-to-date benchmark by re-implementing and re-training existing methods under the same framework and setup, along with a recent architecture for semantic segmentation. Through our benchmark, the community can further investigate the affordance segmentation problem in a fair and reproducible way<sup>1</sup>. We show that models trained under the same setup have lower performance than their performance in the respective papers in the scenario of a single object on a tabletop. We assess generalisation performance of the models by varying the object scale and using the object occupancy i.e., the ratio of pixels belonging to the object mask and the total number of pixels, as proxy for the scale. To complement the variety of architectures that tackle the affordance segmentation problem, we retrain Mask2Former [8], one of the most recent methods using latent vectors for segmentation. Our experiments show that Mask2Former outperforms other methods on the dataset with single object on a tabletop, and on most datasets with hand-occluded objects.

<sup>1</sup> Code and trained models are available at <https://apicis.github.io/aff-seg>

**Table 1:** Characteristics and comparison of visual affordance segmentation models. Note that we report the best-performing backbone for each model, and we do not consider additional parts of the pipelines, such as a separate object detector.

Model	Source architecture	Backbone	FPN	IF	Attention				Aff. Object				CRF	FS	
					Sp	Ch	Sa	Mc	C	E	C	S			L
AffordanceNet [14]	Mask R-CNN [23]	VGG-16 [47]	o	o	o	o	o	o	o	o	o	o	o	o	o
B-Mask R-CNN [36]	Mask R-CNN [34, 35]	RNX-101 [55]	•	o	o	o	o	o	o	o	o	o	o	o	o
4C-RPN-5C [34]	AffordanceNet [14]	SE-RNX-101 [24]	o	o	o	o	o	o	o	o	o	o	o	o	o
A-Mask R-CNN [4]	AffordanceNet [14]	RN-50 [22]	•	o	o	o	o	o	o	o	o	o	o	o	o
BPN [59]	AffordanceNet [14]	RN-50 [22]	•	o	•	•	o	o	•	•	•	o	o	o	o
SEANet [60]	DFE [25]	RN-50 [22]	o	•	•	•	o	o	•	•	•	o	o	o	o
BB-CNN [40]	DeepLab [7]	VGG-16 [47]	o	o	o	o	o	o	o	o	o	o	o	•	•
DeepLab [46]	DeepLab [7]	RN-101 [22]	o	o	o	o	o	o	o	o	o	o	o	•	o
ADOSMNet [5]	PSPNet [66]	RN-101 [22]	o	o	o	o	o	o	o	o	o	o	o	o	o
CNN [39]	SegNet [2]	VGG-16 [47]	o	o	o	o	o	o	o	o	o	o	o	o	o
RN50-F [26]	Fast-FCN [54]	RN-50 [22]	o	o	o	o	o	o	o	o	o	o	o	o	o
ACANet [1]	UNet [43]	RN-18 [22]	o	o	o	o	o	o	o	o	o	o	o	o	o
DRNAtt [17]	DANet [15]	DRN [61]	o	o	•	•	o	o	o	o	o	o	o	o	o
GSE [65]	HRNet [49, 64]	RNS-101 [63]	o	•	•	•	o	o	o	o	o	o	o	o	o
RANet [67]	EncNet [62]	RN-50 [22]	o	o	o	•	o	o	•	•	•	o	o	o	o
STRAP [12]	SINN [38]	RN-50 [22]	o	•	o	o	•	o	•	o	o	o	o	•	o
M2F-Aff	Mask2Former [8, 9]	RN-50 [22]	•	o	o	o	•	•	•	o	o	o	o	o	o

KEYS – •: considered, o: not considered, Sp: spatial attention, Ch: channel attention, Mc: masked cross-attention, Sa: self-attention, Aff.: affordance, C: classification, E: edge segmentation, S: segmentation, L: localisation, RN: ResNet, RNX: ResNeXt, RNS: ResNeSt, SE-RNX: squeeze and excite ResNeXt, DRN: Dilated Residual Network, CRF: conditioned random fields, IF: intermediate feature maps fusion, FS: full supervision.

## 2 Related work

Methods for visual affordance segmentation use semantic and instance segmentation to predict affordance regions on the objects [2, 7, 15, 23, 43, 54, 66]. Table 1 summarises the characteristics of existing models for affordance segmentation discussed in this section. The training and evaluation of models is performed with different datasets (see Table 2), based on the chosen object instances and the affordance classes. Benchmarks are testing sets of datasets, held out during training.

### 2.1 Methods for affordance segmentation

Following instance segmentation, AffordanceNet [14] and A-Mask R-CNN [4] modify Mask R-CNN [23] to predict the affordance masks instead of the object masks for each object localised in an image. 4C-RPN-5C [34] and BPN [59] modify AffordanceNet to align the region of interest with the feature maps at different resolutions and predict the Intersection over Union of bounding boxes and boundaries of affordance regions. However, the additional object detection component can fail to detect objects present in the image or can cause false alarms that result in affordance regions predicted within image parts outside of actual objects. Moreover, the additional affordance edge segmentation task of BPN fails to predict precise affordance contours when edges are blurred in images or not clearly defined (e.g., occlusions or transparent objects). To avoid the dependence from the prediction of the object location, other methods learn to independently assign each pixel of the image to an affordance class during a supervised training (per-pixel affordance segmentation) [1, 17, 26, 39, 60, 65, 67]. Methods such as CNN [39],

RN50-F [26], and ACANet [1] can predict false positives outside the object region when objects are occluded or boundaries are not clearly defined [1].

To focus only on relevant information, recent methods belonging to both approaches use attention mechanisms to weigh feature maps extracted from the image [17, 59, 60, 65, 67]. GSE [65] learns the channels weight with supervision of the affordances. RANet [67] uses the additional supervision of object classes. DRNAtt [17] and SEANet [59] learn similarities between channels or positions in the feature maps without direct supervision, updating the weights of the layers that compute the attention map during training. However, both DRNAtt and GSE process feature maps at low-resolutions for computational reasons, causing the feature map to degrade important details for the affordance segmentation, e.g., edges, when the object is not close to the camera. In RANet, occlusions can cause mistakes in the attention weights and hence a mismatch between the predicted object classes and the segmented affordances.

Unlike previous methods, classification and segmentation of affordances can be decoupled, performing the class assignment at the level of each segmentation mask [12]. For example, STRAP is a multi-branch architecture that learns the affordance classification from the classes present in an image in one branch and the segmentation masks via weakly supervision from a point annotation of a region in another branch [50]. STRAP uses Conditional Random Fields to process the pixel position and color, leading to inaccurate segmentation when the object color is not clearly distinguished from the background [39, 46]. Moreover, the self-attention mechanism used in the transformer encoder processes low-resolution feature maps extracted by the backbone and may lead to lose details about the object in the image if the object scale is small. Mask2Former is a recent hybrid architecture that combines an encoder-decoder convolutional neural network with a transformer decoder to decouple the classification of classes by the segmentation, tackling different types of segmentation, e.g., semantic, instance, and panoptic segmentation [8]. Mask2Former introduced a masking operation in the cross-attention mechanism that combines the latent vectors with the features extracted from the image, ignoring the pixel positions outside the object region. This type of processing, not considered by previous methods, can improve the learning in tasks such affordance segmentation, in which the majority of pixels belongs to the background. In the following, we refer to the version of Mask2Former that segments affordances as M2F-AFF.

## 2.2 Datasets and benchmarks

Table 2 shows the main characteristics of existing datasets for affordance segmentation. Methods are usually compared and benchmarked on the testing split of datasets with images of real scenarios. In general, existing benchmarks do not have a complete overlap of affordance classes, hence benchmarking the performance of methods trained using different datasets is impossible.

Some of the training datasets are synthetic i.e., created using computer graphics, due to the need of a large number of images [1, 10, 11]. These datasets are annotated by labelling the regions of CAD models, rendering the CAD model in a virtual scene, and then projecting back the new annotation in the camera frame. This type of segmentation annotation has a perfect spatial overlapping with the object rendered by the virtual camera. UMD-Synth [11] has objects in different poses placed on a tabletop while varying

**Table 2:** Characteristics of existing datasets for visual affordance segmentation.

Dataset	Images	Categories		Real data	Transp.	TPV	Hand-Occ.
		Objects	Affordances				
AFF-Synth [10]	30,245	21	7	◦	◦	•	◦
UMD-Synth [11]	37,200	17	7	◦	◦	•	◦
Multi-View [29]	47,210	37	15	•	◦	•	◦
HANDAL [18]	308,000	17	1	•	◦	•	◐
TRANS-AFF [28]	1,346	3	3	•	•	•	◦
UMD [37]	28,843	17	7	•	◦	•	◦
IIT-AFF [40]	8,835	10	9	•	◐	◐	◐
CAD120-AFF [46]	3,090	11	6	•	◦	•	◐
FPHA-AFF [26]	4,300	14	8	•	◐	◦	•
CHOC-AFF [1]	138,240	3	3	◐	◐	•	◐

KEYS – Transp.: transparency, TPV: third person view, Hand-Occ.: hand-occlusion, •: considered, ◦: not considered, ◐: mixed (partially considered)

colors and texture. AFF-Synth [10] has images generated using domain randomization techniques to overcome the gap between simulated and real data. Unlike UMD-Synth, each image of AFF-Synth includes multiple objects and some of them are used as distractors during training as they do not have affordance annotation. The main issue with synthetic datasets is the sim-to-real gap in performance: models trained on completely synthetic datasets tend to have poor generalisation performance to real data.

In real datasets images are captured by a real camera and manually annotated to label the pixels of the object regions with an affordance class, following the convention on classes names and definition established by the dataset creators [10, 11, 28, 29, 37, 40, 46]. In general, annotators label object regions with the action an agent can perform *with* or *on* that region (e.g., the handle of an hammer is *graspable* while the hard top allows to *pound* other objects). Manual annotation is a time consuming task, requires a certain amount of effort, and is prone to errors and imperfections. For example, some objects in the scene might not be annotated, especially in the case of clutter or small object scale. Moreover, the mask annotation might be incomplete or outside the object boundaries, depending on the precision of annotators in distinguishing between regions of the object and the background. During training, these mistakes can influence the models learning, and during evaluation, the ranking and values of the performance measure. A model that has imperfect outputs, but more similar to the annotation error, will obtain a higher score than a model predicting masks that are visually better than the annotated mask.

UMD [37] is composed by images of objects placed on a blue rotating table in the same environment. Similarly, Multi-View [29] contains images of objects placed on a white rotating table keeping the same lighting. Multi-view has a higher number of objects and affordances than UMD, however the training set consists only of object crops, hence models cannot be trained using an entire scene. HANDAL [18] varies objects placement in outdoor and indoor environments, such as a street or a living room, which also contain clutter. To annotate images, HANDAL uses a hybrid approach between manual and automatic annotation based on recent off-the-shelf methods. BundleSDF predicts the 6D pose of the objects in each frame and reconstruct the CAD model [53], then the annotator labels the handle of the CAD models as *graspable*. The final annotation mask is obtained by projecting the annotated CAD model in the camera frame through the estimated object and camera pose. The main limitation of HANDAL is the

**Table 3:** Training setup of affordance segmentation methods for the UMD dataset [37]. Our setup is applied to all models to ensure a fair comparison.

Training setup	Image resolution	Data augmentation				Image resize procedure	
		Flipping	Scaling	Rotating	Color jittering	Training set	Testing set
[39]	320 × 240	◦	◦	◦	◦	centre-crop	sliding window
[14]	1000 × 600	◦	◦	◦	◦	unknown	unknown
[65]	400 × 400	•	•	◦	◦	crop	unknown
[67]	224 × 224	◦	◦	◦	◦	centre-crop	unknown
[59]	1000 × 600	•	•	•	•	unknown	unknown
[17]	320 × 240	◦	◦	◦	◦	centre-crop	unknown
Ours	640 × 480	•	•	◦	•	centre-crop	no cropping

KEYS – •: considered, ◦: not considered

focus on objects having a handle, limiting the application of trained models to other types of object. TRANS-AFF [28] is focused on few categories of transparent objects placed on a tabletop, increasing the difficulty to predict affordances, as objects are also in different poses. IIT-AFF [40] is composed by images of objects placed in a cluttered scene to better reflect a scenario with clutter and occlusions. One part of the dataset is collected from other datasets (e.g., ImageNet [13]) by varying the instances and the setting of the images. In CAD120-AFF [46], images are sampled from videos of humans performing activities in a realistic setting, e.g., kitchen or office, with more than one object in the scene and also with hand occlusions. FPHA-AFF [26] has images of hand-held objects acquired from an egocentric point of view. However, the affordance annotation of this dataset is currently not publicly available and egocentric images contains arms from the bottom of the image, resulting in objects highly occluded by the hands. CHOC-AFF [1] is the only available dataset with a high number of images (more than 129,000) with hand-occluded objects, varying hand and object pose.

Table 3 highlights the main limitation of previous works when training models on UMD. The different setups used by the methods make their direct comparisons unfair [14, 17, 65]. In this work, we tackle affordance segmentation of a single object in two scenarios, on a tabletop and hand-occluded, and revisit the UMD and CHOC-AFF benchmarks. We select these two datasets due to the trade-off between real images of the entire scene, number of images, scenario, number of object and affordance categories.

### 3 Unoccluded object on a tabletop

#### 3.1 Experimental setup

We consider UMD [37], a standard dataset used to evaluate the visual affordance segmentation of a single object without occlusions. UMD has 28,843 images of tools and containers, each placed on a blue rotating table, and acquired from a fixed view with the object in the center [37]. Object instances are not evenly distributed among the 17 object categories, e.g., there are 10 instances of spoon and 2 instances of pot. UMD is annotated with 7 affordance classes: *grasp*, *cut*, *scoop*, *contain*, *pound*, *support*, *wrap-grasp*. UMD provides a pre-defined split of the dataset into training (14, 823 images) and testing sets (14, 020 images), holding out approximately half of object instances per

category. The training set is subsequently split into training and validation sets, sampling 30% of the images for the latter.

We compare M2F-AFF with three existing methods: AffordanceNet [14], CNN [39] and DRNAtt [17]. CNN is based on an encoder-decoder architecture to segment affordances. AffordanceNet is a two-stage method that detects the object and segments affordances. DRNAtt is an attention-based architecture and is the best-performing method in tabletop settings [17]. Due to the unavailable implementations, we re-implemented DRNAtt and CNN, and trained all the methods using the same experimental setup to ensure fair conditions.

Previous works evaluated methods performance using  $F_{\beta}^w$  [33] as performance measure.  $F_{\beta}^w$  associates a different weight to the prediction errors based on the Euclidean distance to the annotated mask. Despite the positive aspect of attributing different importance to prediction mistakes,  $F_{\beta}^w$  was proposed to evaluate foreground maps and therefore ignores the classes that are not present in the annotation. To also count prediction mistakes of classes that are not in the annotation, and when comparing models predictions  $\{\hat{M}_i\}_{i=1}^N$  with annotations  $\{M_i\}_{i=1}^N$ , where  $N$  is the number of images in a dataset, we use per-class Jaccard index ( $J$ ) to measure how many pixels predicted as class  $s$  are correct, among all pixels (predicted and annotated):

$$J = \frac{\sum_{i=1}^N \sum_x^{W \times H} TP}{\sum_{i=1}^N \sum_x^{W \times H} TP + FP + FN}, \quad (1)$$

where we consider a pixel  $x$  having value  $s$  in the prediction  $\hat{M}$  and in the annotation  $M$  as true positive ( $TP$ ), a pixel having value  $s$  in  $\hat{M}$ , but not in  $M$  as false positive ( $FP$ ), a pixel having value  $s$  in  $M$ , but not in  $\hat{M}$  as false negative ( $FN$ ).

### 3.2 Strategy and parameters setting

Table 3 compares the training setups of the previous works with ours. Although almost each paper used a different setup for training the method, some previous works made straight comparisons of performance values with the results in related works [14, 17, 65]. CNN [39], GSE [65], and RANet [67] use an input resolution lower than the UMD images ( $640 \times 480$ ), cropping the images around the annotation to fit the resolution. This procedure implies the presence of a perfect object detector to locate the object in the scene. Unlike these setups, ours uses the full image resolution for every model, avoiding the perfect detector assumption for all the models. Methods that use the additional object detection task, such as AffordanceNet [14] and BPN [59], resize the image to a resolution with minimum height 600 pixels and maximum with 1,000 pixels. Thus, AffordanceNet and BPN process feature maps with higher resolutions than the ones obtained with  $640 \times 480$  images (UMD resolution). Only GSE [65] and BPN [59] use augmentation to limit the overfitting on the images having all similar background, i.e., the blue rotating table. In this work, we train CNN and DRNAtt using cross-entropy loss, Adam optimiser and batch size 4. For AffNet, we used a combination of cross-entropy and smooth-L1 losses for detection and cross-entropy for segmentation, mini-batch gradient descent with weight decay 0.001 and batch size 2 to fit our available GPU memory.

We initialise all the mentioned architectures with a learning rate of  $10^{-3}$ . We train M2F-AFF with AdamW optimiser [32] with batch size 4 and learning rate  $10^{-4}$ . We follow Mask2Former setup [8], using hungarian algorithm [31] to match the prediction from each latent vector with the corresponding annotation. We minimize the linear combination among the cross-entropy for classification  $L_{cls}$ , with the binary dice loss  $L_{dice}$  and the binary cross-entropy  $L_{ce}$  for segmentation. Losses are weighted using hyperparameters  $\lambda_{ce}$ ,  $\lambda_{dice}$ , and  $\lambda_{cls}$ .

For all models, we decrease the learning rate by a factor of 0.5, if there is no increase of the mean Intersection over Union in the validation set for 3 consecutive epochs. We set the maximum number of epochs to 100. We also use early stopping with a patience of 5 epochs to reduce over-fitting. We use flipping with a probability of 0.5, scaling by randomly sampling the scale factor in the interval [1, 1.5] and center-cropping to simulate a zoom-in effect, color jitter with brightness, contrast, and saturation sampled randomly in the interval [0.9, 1.1], and hue sampled randomly in the interval [-0.1, 0.1]. These augmentations allow to increase variability in the training set. We initialise all backbones with weights trained on ImageNet [13].


### 3.3 Implementation details

For DRNAtt and CNN, we used the backbones described in the corresponding paper due to design requirements. In DRNAtt, DRN22 backbone allows to maintain higher resolution compared to a ResNet backbone. In CNN, the features down-sampling is not performed by VGG16 blocks but by down-sampling layers and the indices of down-sampled position are saved to be re-used in the decoder, while a ResNet backbone decreases the resolution using stride 2 in convolutional filters. We used an available implementation of AffordanceNet [10], that differs from the original one by implementing the ResNet-50 backbone with the Feature Pyramid Network instead of using a VGG backbone. Moreover, to fit the available memory, each up-sampling layer in the segmentation decoder contains 128 channels as opposed to 512 in AffordanceNet. The first up-sampling layers has kernel size 4 instead of 8 and the stride is 2 instead of 4. For M2F-AFF, we chose the Mask2Former version that uses the ResNet-50 backbone, multi-scale deformable attention [68], 3 decoder layers repeating the sequence of a masked cross-attention layer with 8 heads, a self-attention layer with 8 heads, a feed-forward layer, and 100 queries. During training, we follow Mask2Former setup, setting the default loss weights  $\lambda_{ce} = 5$ ,  $\lambda_{dice} = 5$ ,  $\lambda_{cls} = 2$  and to compute losses values, we sample sets of  $K = 12,544$  points for prediction and ground truth masks [8]. During inference, we use a confidence threshold of 0.5 to separate background from an affordance region.

### 3.4 Comparative analysis and segmentation results

Fig. 1 compares the performance of affordance segmentation methods on the UMD testing set using the Jaccard Index. On average, M2F-AFF outperforms other methods. Previous works simply reported the  $F_{\beta}^w$  even though the experimental setup used was different. The ranking of methods is similar to the one in previous works [6, 17], but with a different range of performance values ( $F_{\beta}^w$  results in the Supplementary Material document). DRNAtt is the previous work that has the best performance compared





**Fig. 1:** Comparison of segmentation accuracy between methods on the testing set of UMD [37]. Methods are re-implemented and trained using the same experimental setup. KEYS – G: *grasp*, CU: *cut*, SC: *scoop*, CO: *contain*, P: *pound*, SU: *support*, WG: *wrap-grasp*, AVG: *average*, — CNN [39], — AffordanceNet [14], — DRNAtt [17], — M2F-AFF.

to CNN and AffordanceNet, especially for the class *support*. M2F-AFF has the highest improvement in Jaccard Index for the class *cut* (+13 percentage points) and *scoop* (+16 percentage points) compared to DRNAtt.


Fig. 2 compares the predictions of methods with the affordance segmentation annotations. For some images, the annotation of the affordance mask is either wider or smaller than the actual region in the image (see first two columns). In these cases, models visually segments the affordance regions more accurately than the annotation. For example, the predicted *pound* region covers the whole top of the hammer (second column). For other cases, such as the images of a saw and a mug, both annotation and M2F-AFF prediction are similar (see third and fourth columns). Furthermore, M2F-AFF predictions are visually accurate despite class confusions (last two columns), even compared to methods that segment the correct class (AffNet in the last column).

## 4 Hand-held containers

### 4.1 Dataset

We follow the experimental setup in previous evaluations [1] and we report the results from their validation. This setup includes a mixed-reality dataset and two testing sets of real images, all publicly available; Jaccard index as performance measures (see Eq. 1); and four methods, all trained on the mixed-reality dataset and evaluated on both the mixed-reality dataset and the two testing sets.

CORSMAL Hand-Occluded Containers Affordance (CHOC-AFF) [1,52] has 138,240 images of 48 synthetic containers hand-held by synthetic hands rendered on top of 30 real backgrounds, and with annotations of the object affordances (*graspable* and *contain*) and forearm masks (*arm*). The locations and poses of the hand-held objects vary across images, with 129,600 images of hand-held objects rendered above a table, and 8,640 images with objects rendered on a tabletop. CHOC-AFF has four splits: a training set of 103,680 images, combining 26 backgrounds and 36 objects instances; a validation set of 17,280 images, using all 30 backgrounds and holding out 6 objects instances; a testing



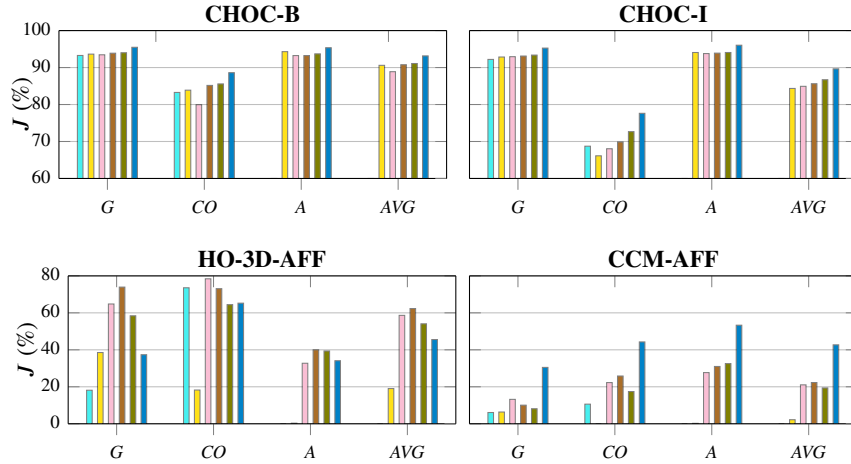
**Fig. 2:** Comparison between the ground truth and models prediction on the UMD testing set. Legend: ■ *graspable*, ■ *cut*, ■ *scoop*, ■ *contain*, ■ *pound*, ■ *support*, ■ *wrap-grasp*

set of 13,824 images (CHOC-B), used to assess the generalisation performance to backgrounds never seen during training; and another testing set of 17,280 images (CHOC-I), to assess the generalisation performance to object instances never seen during training.

HO-3D-AFF and CCM-AFF are the two testing sets of real images used to assess the generalisation performance of the models [1]. HO-3D-AFF sampled 150 images from HO-3D [19], a dataset with images of hands holding and manipulating objects, such as tools and containers (only mugs, boxes, and cans, in different poses, were selected for the sampled images). HO-3D-AFF complemented the existing annotations with affordance masks. CCM-AFF sampled 150 images from CORSMAL Containers Manipulation (CCM) [57], a dataset consisting of audio-visual recordings of people manipulating containers. Images were selected from a lateral view, while varying object instances and conditions, such as lighting, filling type and amount, and presence of the tablecloth. CCM-AFF annotated the affordance and forearm masks. Both testing sets use the same classes of CHOC-AFF: *graspable*, *contain*, and *arm*.

## 4.2 Methods under comparison

We compare the performance of our re-implemented versions of RN50-F [26], DRNAtt [17], RN18-U [43], ACANet [1] with M2F-AFF trained with CHOC-AFF affordance classes. RN50-F uses a ResNet-50 encoder with a pyramid scene parsing module [66] to segment only the object affordances *graspable* and *contain* [26]. DRNAtt uses position and channel attention mechanisms in parallel after the feature extraction



**Fig. 3:** Comparison of the affordance and arm segmentation results between the models on the two mixed-reality testing sets (top row) and on the two real testing sets (bottom row). Note the different y-axis limits. KEYS – G: *grasp*, CO: *contain*, A: *arm*, AVG: *average*, — RN50F [26], — DRNAtt [17], — RN18-U [1], — ACANet [1], — ACANet50 [1], — M2F-AFF [8].

stage [17]. The outputs of the attention modules are summed element-wise and the result is up-sampled through a learnable decoder. RN18-U and ACANet [1] are UNet-like models that gradually down-sample feature maps in the encoder and up-sample them in the decoder, preserving the information via skip connections [43]. Both models implement a ResNet encoder [22]. ACANet separately segments object and hand regions, using these masks to weigh the feature maps learnt in a third branch for the final affordance segmentation. Additionally, we trained a version of ACANet with ResNet-50 to have the same backbone of M2F-AFF and RN50-F.

M2F-AFF uses the same architecture described in Sec. 3.2 and 3.3, only the number of output classes changes to fit the hand-occluded settings. We used the same training setup of compared methods [1], initialising the backbone with ImageNet weights. We use early stopping with a patience of 10 epochs, setting the maximum number of epochs to 100. The learning rate decreases by a factor of 0.5 if there is no increase of the mean Intersection over Union in the validation set for 3 consecutive epochs. We use the cropping window technique described in the previous work [1], resize images by a factor sampled in the interval [1, 1.5] and center crop the resized image with a  $480 \times 480$  window. To simulate the other arm holding the object, we use horizontal flip with a probability of 0.5. M2F-AFF uses batch size 4, learning rate 0.0001, and an additional Gaussian noise augmentation with variance in range [10, 100] to increase variability in training data.

### 4.3 Segmentation results under hand-occlusions

Fig. 3 compares the results of M2F-AFF and previous methods. Models achieve high and similar performance on the testing sets of CHOC-AFF, whereas there are signifi-



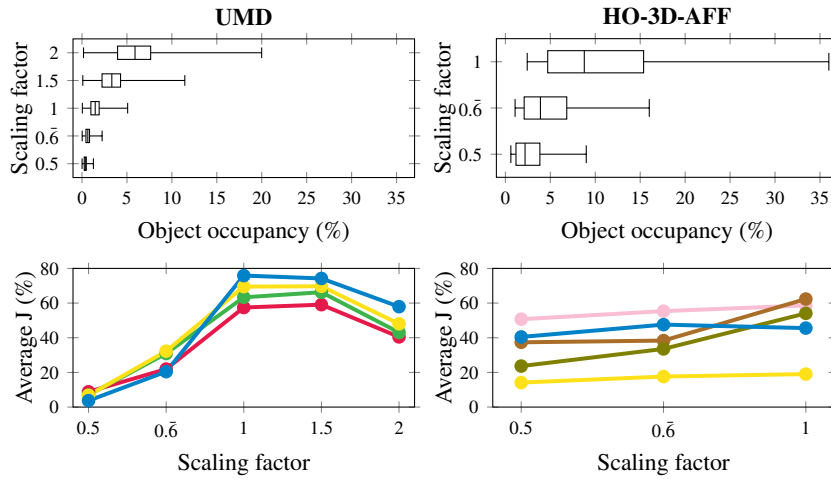
**Fig. 4:** Comparison between the ground truth, ACANet and M2F-AFF predictions on the testing sets of HO-3D-AFF and CCM-AFF. Legend: ■ *graspable*, ■ *contain*,  *arm*.

cant performance differences across the models on the two testing sets with real images. In HO-3D-AFF, the best-performing model is still ACANet, whereas M2F-AFF performance is mostly affected by the missing segmentation of *graspable* regions of the boxes close to the camera. This result suggests that the model segments better affordances of objects with low pixel occupancy. CCM-AFF is more challenging than HO-3D-AFF due to the presence of the human body and transparent containers. M2F-AFF outperforms existing models on this testing set.

Fig. 4 visually compares the segmentation results of ACANet and M2F-AFF on images from HO-3D-AFF and CCM-AFF, selected with object and hands in challenging conditions, e.g., a mug kept downward, a white cup held over a colored tablecloth. M2F-AFF has a more complete *arm* segmentation for a challenging hand-mug pose (1st column). M2F-AFF cannot segment the *graspable* region of the red box when the object is close to the camera (2nd column). However, M2F-AFF segments a more complete *graspable* and *arm* region than ACANet for the yellow and white box when farther from the camera (4th column). For CCM-AFF samples, M2F-AFF predicts more complete *arm* masks than ACANet. An incomplete region of the box is segmented as *graspable* when there is the colored table-cloth (4th column). Compared to ACANet, M2F-AFF reduces the number of false positives on regions such as the colored tablecloth and the human clothes correctly segmenting them as background.

## 5 Robustness to scale variations

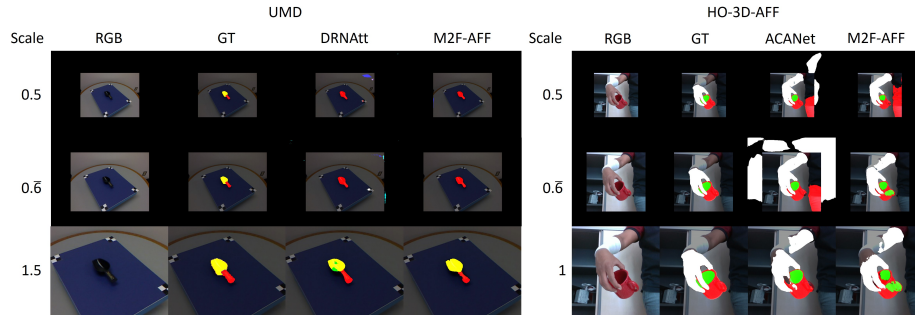
To analyse the performance of models when varying object scale, we simulate a zoom-out in the images dividing height and width by 1.5, 2, and zero-padding images to respect the input resolution of models. A zoom-in effect is applied by multiplying height and width by 1.5, 2 and then center cropping the images to the model input resolution. We



**Fig. 5:** Statistics of object pixels varying occupancy in UMD (left) and HO-3D-AFF (right) testing sets, and relative Average Jaccard Index ( $J$ ). The whiskers plot highlights the minimum and maximum values as well as the quartiles of the occupancy distribution in each testing set. When the scaling factor is greater than 1, images are center cropped, when it is lower than 1, images are zero-padded. Legend: — CNN [39], — AffordanceNet [14], — DRNAtt [17], — RN18-U, — ACANet [1], — ACANet50 [1], — M2F-AFF.

assess the performance in UMD for the tabletop setting and HO-3D-AFF for the hand-occluded setting. In UMD the object is framed in the image center and the distance with the camera is fixed, allowing for both scale variations without losing image quality. HO-3D-AFF is the testing set with the highest object scale (occupancy with a maximum over 35%, see Supp. Mat.) and where has average performance lower than other methods. Therefore, we zoom-out the images to shift the distribution closer to the one of other datasets with hand-occluded objects.

Fig. 5 (top left) shows the statistics of object occupancy in UMD testing set. The original testing set of UMD has a low statistic of the closest object occupancy being approximately 5%. As expected, by zooming-out, the statistics is shifted towards values under 5% and conversely the zooming-in causes the statistics to increase. Fig. 5 (bottom left) shows the average Jaccard Index of re-implemented methods on the UMD testing sets when varying object occupancy. Methods achieve a higher average Jaccard Index in zoomed-in testing sets than zoomed-out, due to the fact that the scaling in [1, 1.5] is present in the training augmentation. M2F-AFF outperforms other methods when the testing set images are zoomed-in. When images are zoomed-out, the best performing models are AffNet and CNN. Fig. 5 (top right) shows the object occupancy of images in HO-3D-AFF, progressively shifting to lower values when zooming-out images to reduce object scale. As shown in Fig. 5 (bottom right), RN18-U outperforms all models when the scaling factor is lower than 1, while M2F-AFF is the second best-performing model. Both ACANet and ACANet50 have an increase of more than 20 percentage points from the scaling factor of 0.6 to 1. DRNAtt is the only model that keeps low generalisation,



**Fig. 6:** Comparison between the segmentation ground-truths, and predictions of DRNAtt [17], ACANet [1], M2F-AFF, on UMD and CCM-AFF testing sets varying object scale. Legend: — *graspable*, — *cut*, — *scoop*, — *contain*, — *pound*, — *support*, — *wrap-grasp*, □ *arm*.

not only to HO-3D-AFF testing set but also to scale variation. These results suggest that the generalisation to different object scales remains an open challenge for affordance segmentation, providing a potential direction for future contributions.

Fig. 6 shows the comparison between models predictions when varying object scale. In UMD, when objects are scaled down, both DRNAtt and M2F-AFF produce false positives also outside the object region (first two rows), when the objects are large scaled up, the false positive are inside the object region (DRNAtt last row). In HO-3D-AFF, when objects are scaled down, ACANet produces false positives for the *arm* class in the additional zero-padding and for the *graspable* class on the human body, while M2F-AFF has similar predictions and produces false positives for *graspable* on the human body with the lowest object scale.

## 6 Conclusion

We identified reproducibility issues with affordance segmentation methods. These issues range from unknown training setups to unfair comparative evaluations. We re-implemented and evaluated a range of selected methods under a uniform training setup and common testing sets across two datasets. By re-training Mask2Former, a recent architecture based on latent vectors that was not yet considered for affordance segmentation, we showed that this model outperforms other methods on the testing set with objects on a tabletop and on most testing sets in the hand-occluded settings. Moreover, we analysed and showed how the methods are not robust to scale variations despite the use of scaling in the data augmentation during training. As future work, we will extend the benchmarking evaluation to datasets with multiple objects in an image and where occlusions are caused by other objects.

## References

1. Apicella, T., Xompero, A., Ragusa, E., Berta, R., Cavallaro, A., Gastaldo, P.: Affordance segmentation of hand-occluded containers from exocentric images. In: IEEE Int. Conf. Comput.

- Vis. Workshops (2023)
2. Badrinarayanan, V., Handa, A., Cipolla, R.: Segnet: A deep convolutional encoder-decoder architecture for robust semantic pixel-wise labelling. *arXiv:1505.07293v1 [cs.CV]* (2015)
  3. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: *Eur. Conf. Comput. Vis.* (2020)
  4. Caselles-Dupré, H., Garcia-Ortiz, M., Filliat, D.: Are standard object segmentation models sufficient for learning affordance segmentation? *arXiv:2107.02095v1 [cs.LG]* (2021)
  5. Chen, D., Kong, D., Li, J., Wang, S., Yin, B.: Adosmnet: a novel visual affordance detection network with object shape mask guided feature encoders. *Multimedia Tools and Applications* pp. 1–25 (2023)
  6. Chen, D., Kong, D., Li, J., Wang, S., Yin, B.: A survey of visual affordance recognition based on deep learning. *IEEE Trans. Big Data* pp. 1–20 (2023)
  7. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell.* **40**(4), 834–848 (2017)
  8. Cheng, B., Misra, I., Schwing, A.G., Kirillov, A., Girdhar, R.: Masked-attention mask transformer for universal image segmentation. In: *Conf. Comput. Vis. Pattern Recognit.* (2022)
  9. Cheng, B., Schwing, A., Kirillov, A.: Per-pixel classification is not all you need for semantic segmentation. In: *Adv. Neural Inf. Process. Syst.* (2021)
  10. Christensen, A.D., Lehotský, D., Jørgensen, M.W., Chrysostomou, D.: Learning to segment object affordances on synthetic data for task-oriented robotic handovers. In: *Brit. Mach. Vis. Conf.* (2022)
  11. Chu, F., Xu, R., Vela, P.: Learning affordance segmentation for real-world robotic manipulation via synthetic images. *IEEE Robotics Autom. Lett.* **4**(2), 1140–1147 (2019)
  12. Cui, L., Chen, X., Zhao, H., Zhou, G., Zhu, Y.: Strap: Structured object affordance segmentation with point supervision. *arXiv:2304.08492v1 [cs.CV]* (2023)
  13. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *Conf. Comput. Vis. Pattern Recognit.* (2009)
  14. Do, T., Nguyen, A., Reid, I.: AffordanceNet: An end-to-end deep learning approach for object affordance detection. In: *IEEE Int. Conf. Robotics Autom.* (2018)
  15. Fu, J., Liu, J., Tian, H., Li, Y., Bao, Y., Fang, Z., Lu, H.: Dual attention network for scene segmentation. In: *Conf. Comput. Vis. Pattern Recognit.* (2019)
  16. Gibson, J.J., Carmichael, L.: *The senses considered as perceptual systems*, vol. 2. Houghton Mifflin Boston (1966)
  17. Gu, Q., Su, J., Yuan, L.: Visual affordance detection using an efficient attention convolutional neural network. *Neurocomputing* **440**, 36–44 (2021)
  18. Guo, A., Wen, B., Yuan, J., Tremblay, J., Tyree, S., Smith, J., Birchfield, S.: Handal: A dataset of real-world manipulable object categories with pose annotations, affordances, and reconstructions. In: *IEEE Int. Conf. Intell. Robot Syst.* (2023)
  19. Hampali, S., Rad, M., Oberweger, M., Lepetit, V.: Honnotate: A method for 3D annotation of hand and object poses. In: *Conf. Comput. Vis. Pattern Recognit.* (2020)
  20. Hassan, M., Ghosh, P., Tesch, J., Tzionas, D., Black, M.J.: Populating 3d scenes by learning human-scene interaction. In: *Conf. Comput. Vis. Pattern Recognit.* (2021)
  21. Hassanin, M., Khan, S., Tahtali, M.: Visual affordance and function understanding: A survey. *ACM Computing Surveys (CSUR)* **54**(3), 1–35 (2021)
  22. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Conf. Comput. Vis. Pattern Recognit.* (2016)
  23. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: *IEEE Int. Conf. Comput. Vis.* (2017)
  24. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: *Conf. Comput. Vis. Pattern Recognit.* (2018)

25. Hu, Y., Chen, Y., Li, X., Feng, J.: Dynamic feature fusion for semantic edge detection. In: International Joint Conference on Artificial Intelligence (2019)
26. Hussain, S., Liu, L., Xu, W., Lu, C.: FPHA-Afford: A domain-specific benchmark dataset for occluded object affordance estimation in human-object-robot interaction. In: IEEE Int. Conf. Image Process. (2020)
27. Jaegle, A., Gimeno, F., Brock, A., Vinyals, O., Zisserman, A., Carreira, J.: Perceiver: General perception with iterative attention. In: Int. Conf. on Machine Learning (2021)
28. Jiang, J., Cao, G., Do, T., Luo, S.: A4t: Hierarchical affordance detection for transparent objects depth reconstruction and manipulation. *IEEE Robotics Autom. Lett.* **7**(4), 9826–9833 (2022)
29. Khalifa, Z., Shah, S.A.A.: A large scale multi-view RGBD visual affordance learning dataset. In: IEEE Int. Conf. Image Process. (2023)
30. Kjellström, H., Romero, J., Kragić, D.: Visual object-action recognition: Inferring object affordances from human demonstration. *Comput. Vis. Image Understanding* **115**(1), 81–90 (2011)
31. Kuhn, H.W.: The hungarian method for the assignment problem. *Naval research logistics quarterly* **2**(1-2), 83–97 (1955)
32. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: International Conference on Learning Representation (2019)
33. Margolin, R., Zelnik-Manor, L., Tal, A.: How to evaluate foreground maps? In: Conf. Comput. Vis. Pattern Recognit. (2014)
34. Minh, C.N.D., Gilani, S.Z., Islam, S.M.S., Suter, D.: Learning affordance segmentation: An investigative study. In: Digital Image Computing: Techniques and Applications (2020)
35. Morrison, D., Milan, A., Antonakos, E.: Uncertainty-aware instance segmentation using dropout sampling. In: Robotic Vision Probabilistic Object Detection Challenge (CVPR Workshop) (2019)
36. Mur-Labadia, L., Martinez-Cantin, R., Guerrero, J.J.: Bayesian deep learning for affordance segmentation in images. *IEEE Int. Conf. Robotics Autom.* (2023)
37. Myers, A., Teo, C.L., Fermüller, C., Aloimonos, Y.: Affordance detection of tool parts from geometric features. In: IEEE Int. Conf. Robotics Autom. (2015)
38. Nauata, N., Hu, H., Zhou, G.T., Deng, Z., Liao, Z., Mori, G.: Structured label inference for visual understanding. *IEEE Trans. Pattern Anal. Mach. Intell.* **42**(5), 1257–1271 (2019)
39. Nguyen, A., Kanoulas, D., Caldwell, D., Tsagarakis, N.: Detecting object affordances with convolutional neural networks. In: IEEE Int. Conf. Intell. Robot Syst. (2016)
40. Nguyen, A., Kanoulas, D., Caldwell, D., Tsagarakis, N.: Object-based affordances detection with convolutional neural networks and dense conditional random fields. In: IEEE Int. Conf. Intell. Robot Syst. (2017)
41. Osiurak, F., Rossetti, Y., Badets, A.: What is an affordance? 40 years later. *Neuroscience & Biobehavioral Reviews* **77**, 403–417 (2017)
42. Pang, Y., Xompero, A., Oh, C., Cavallaro, A.: Towards safe human-to-robot handovers of unknown containers. In: IEEE Int. Conf. Robot & Human Interactive Comm. (2021)
43. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: Int. Conf. Med. Image Comput. Comput. Assist. Interv. (2015)
44. Rosenberger, P., Cosgun, A., Newbury, R., Kwan, J., Ortenzi, V., Corke, P., Grafinger, M.: Object-independent human-to-robot handovers using real time robotic vision. *IEEE Robotics Autom. Lett.* **6**(1), 17–23 (2020)
45. Sanchez-Matilla, R., Chatzilygeroudis, K., Modas, A., Duarte, N.F., Xompero, A., Frossard, P., Billard, A., Cavallaro, A.: Benchmark for human-to-robot handovers of unseen containers with unknown filling. *IEEE Robotics and Automation Letters* **5**(2), 1642–1649 (2020)
46. Sawatzky, J., Srikantha, A., Gall, J.: Weakly supervised affordance detection. In: Conf. Comput. Vis. Pattern Recognit. (2017)



47. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representation* (2015)
48. Strudel, R., Garcia, R., Laptev, I., Schmid, C.: Segmenter: Transformer for semantic segmentation. In: *IEEE Int. Conf. Comput. Vis.* (2021)
49. Sun, K., Xiao, B., Liu, D., Wang, J.: Deep high-resolution representation learning for human pose estimation. In: *Conf. Comput. Vis. Pattern Recognit.* (2019)
50. Tang, M., Perazzi, F., Djelouah, A., Ben Ayed, I., Schroers, C., Boykov, Y.: On regularized losses for weakly-supervised cnn segmentation. In: *Eur. Conf. Comput. Vis.* (2018)
51. Thakur, S., Beyan, C., Morerio, P., Murino, V., Del Bue, A.: Leveraging next-active objects for context-aware anticipation in egocentric videos. In: *WACV* (2024)
52. Weber, X., Xompero, A., Cavallaro, A.: A mixed-reality dataset for category-level 6D pose and size estimation of hand-occluded containers. *arXiv:2211.10470v1 [cs.CV]* (2022)
53. Wen, B., Tremblay, J., Blukis, V., Tyree, S., Müller, T., Evans, A., Fox, D., Kautz, J., Birchfield, S.: Bundlesdf: Neural 6-dof tracking and 3d reconstruction of unknown objects. In: *Conf. Comput. Vis. Pattern Recognit.* (2023)
54. Wu, H., Zhang, J., Huang, K., Liang, K., Yu, Y.: FastFCN: Rethinking dilated convolution in the backbone for semantic segmentation. *arXiv:1903.11816v1 [cs.CV]* (2019)
55. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: *Conf. Comput. Vis. Pattern Recognit.* (2017)
56. Xompero, A., Donaher, S., Iashin, V., Palermo, F., Solak, G., Coppola, C., Ishikawa, R., Nagao, Y., Hachiuma, R., Liu, Q., et al.: The corsmal benchmark for the prediction of the properties of containers. *IEEE Access* **10**, 41388–41402 (2022)
57. Xompero, A., Sanchez-Matilla, R., Mazzon, R., Cavallaro, A.: CORSMAL Containers Manipulation (2020), [https://doi.org/10.17636/101CORSMAL1.\(1.0\)](https://doi.org/10.17636/101CORSMAL1.(1.0)) [Data set]. Queen Mary University of London.
58. Yang, W., Paxton, C., Mousavian, A., Chao, Y., Cakmak, M., Fox, D.: Reactive human-to-robot handovers of arbitrary objects. In: *IEEE Int. Conf. Robotics Autom.* (2021)
59. Yin, C., Zhang, Q.: Object affordance detection with boundary-preserving network for robotic manipulation tasks. *Neural. Comput. Appl.* **34**(20), 17963–17980 (2022)
60. Yin, C., Zhang, Q., Ren, W.: A new semantic edge aware network for object affordance detection. *J. Intelligent & Robotic Systems* **104**(1), 1–16 (2022)
61. Yu, F., Koltun, V., Funkhouser, T.: Dilated residual networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 472–480 (2017)
62. Zhang, H., Dana, K., Shi, J., Zhang, Z., Wang, X., Tyagi, A., Agrawal, A.: Context encoding for semantic segmentation. In: *Conf. Comput. Vis. Pattern Recognit.* (2018)
63. Zhang, H., Wu, C., Zhang, Z., Zhu, Y., Lin, H., Zhang, Z., Sun, Y., He, T., Mueller, J., Manmatha, R., et al.: Resnest: Split-attention networks. In: *Conf. Comput. Vis. Pattern Recognit.* (2022)
64. Zhang, H., Xue, J., Dana, K.: Deep ten: Texture encoding network. In: *Conf. Comput. Vis. Pattern Recognit.* (2017)
65. Zhang, Y., Li, H., Ren, T., Dou, Y., Li, Q.: Multi-scale fusion and global semantic encoding for affordance detection. In: *Int. Joint Conf. on Neural Networks* (2022)
66. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: *Conf. Comput. Vis. Pattern Recognit.* pp. 2881–2890 (2017)
67. Zhao, X., Cao, Y., Kang, Y.: Object affordance detection with relationship-aware network. *Neural. Comput. Appl.* **32**(18), 14321–14333 (2020)
68. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection (2020)