

Affordance segmentation of hand-occluded containers from exocentric images

Tommaso Apicella^{1,2}, Alessio Xompero², Edoardo Ragusa¹,
Riccardo Berta¹, Andrea Cavallaro^{2,3,4}, Paolo Gastaldo¹

¹University of Genoa, Italy, ²Queen Mary University of London, U.K.

³Idiap Research Institute, ⁴École Polytechnique Fédérale de Lausanne, Switzerland

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

{edoardo.ragusa, riccardo.bertha, paolo.gastaldo}@unige.it

Abstract

Visual affordance segmentation identifies the surfaces of an object an agent can interact with. Common challenges for the identification of affordances are the variety of the geometry and physical properties of these surfaces as well as occlusions. In this paper, we focus on occlusions of an object that is hand-held by a person manipulating it. To address this challenge, we propose an affordance segmentation model that uses auxiliary branches to process the object and hand regions separately. The proposed model learns affordance features under hand-occlusion by weighting the feature map through hand and object segmentation. To train the model, we annotated the visual affordances of an existing dataset with mixed-reality images of hand-held containers in third-person (exocentric) images. Experiments on both real and mixed-reality images show that our model achieves better affordance segmentation and generalisation than existing models.

1. Introduction

The term *affordance* refers to the potential interactions between an agent and the environment [7]. Visual affordance segmentation, the task of identifying affordances in a scene observed by a camera [10], enables assistive technologies for robotics and prosthetic applications (e.g., grasping, object manipulation) or collaborative human-robot scenarios (e.g., handovers) [3,24,32]. In this scenario, examples of object affordances include *grasp* and *contain* for a cup and *grasp* and *cut* for a knife.

Visual affordance segmentation is challenging due to the variety of objects with different affordances; objects of the same category that can vary in their physical and appearance properties; occlusions by other objects that can be present in the image; lighting conditions that can affect the object appearance in the image; and affordance regions in the images that are not always easily identifiable by edges. Most of the




Figure 1: Predictions of the proposed visual affordance segmentation model on RGB images of hand-occluded containers. These images (background and object instances) are never seen during training and do not belong to any of the datasets. Key: ■ *graspable*, ■ *contain*, ■ *arm*.

methods for visual affordance segmentation modify existing models for semantic or instance segmentation to predict the affordances of objects placed on a table top [8, 18–22, 34]. Visual affordance segmentation is even more challenging when the object is hand-held by a person due to the occlusions caused by the hand and the different poses that the object can take. Only one method addresses this scenario, but the model does not explicitly consider the presence of the forearm and the hand [14]. This can result in inaccurate affordance segmentation. Moreover, the focus of the method on egocentric images from human perspective could be unsuitable for an assistive application, e.g. human-robot collaboration, or generalise to exocentric images.

In this paper, we propose a method that focuses the feature learning on the object and hand region separately, and predicts affordances of an object despite hand occlusions (see Fig. 1). We devise a UNet-like [23] multi-branch architecture to predict object and hand segmentation, and we develop a fusion module to learn separate sets of features in the hand and object region. To train our model, we extend the annotation of an existing mixed-reality dataset of

hand-held containers seen from a third-person view with visual affordances. To evaluate the generalisation performance in real settings, we select and manually annotate a set of images from two public datasets with hand-object interaction: HO-3D [9] and CORSMAL Containers Manipulation (CCM) [31]. Moreover, we re-train existing methods on the annotated dataset and perform a comparative analysis, showing that our proposed model can achieve more accurate segmentation and better generalisation¹.

2. Related work

In this section, we discuss existing methods for visual affordance segmentation from RGB images. We also compare existing datasets for this task and discuss their limitations.

2.1. Methods

Methods for visual affordance segmentation are usually based on models for semantic or instance segmentation, but considering affordance classes as semantic labels [19, 21, 35]. Some of these methods can be part of a two-stage approach that includes a detection step to locate objects of interest. Affordances are then predicted in the region cropped around the detected object or corresponding feature map [1, 3–5, 10, 20, 33]. The two steps can be tackled independently [1, 20], to extract the regions of interest, or learned in an end-to-end manner using multi-tasking [3, 5, 33]. For example, AffordanceNet [3, 5] replaces the instance segmentation branch of Mask R-CNN [12] to predict affordance classes of the regions of interest localised in the backbone. The dependence on the detection step can result in wrongly predicting the presence of an object or missing to localise the object, especially when the object is occluded by a hand.

Other methods use attention mechanisms to focus on the object area or select more relevant features for the affordance segmentation directly from the input image [8, 34–36]. For example, DRNAtt [8] uses Spatial Attention Module to model contextual information (similarity between features in each pixel position) and Channel Attention Module to model channel inter-dependencies (similarity among channels) [6]. However, DRNAtt could be sensitive to large portions of the image belonging to the background. Shared Gradient Attention [34] combines affordance segmentation with semantic edge detection. However, this model could be sensitive to edges or object transparencies.

All previous methods focus on scenes with objects that are placed on a tabletop, and are either fully visible or partially occluded due to clutter [8, 19, 34–36]. These objects are often opaque or textured and easily distinguishable from the background, causing the affordance segmentation models to fail when observed objects are in more challenging

poses or occluded, as in the case of hand-held objects when manipulated by a person.

There is only one close work that tackles the segmentation of affordances of hand-occluded objects [14]. This work uses a ResNet-FastFCN [30] with a pyramid parsing module to predict high-resolution outputs by using global contextual information (usually beneficial for segmenting general scenes). However, the model ignores the hand and the global contextual information does not help the model to capture fine affordance predictions on the object of interest. This can result in inaccurate segmentation of the affordances (e.g. predicted on the hand region).

2.2. Datasets

Existing datasets are designed to segment the functional (with respect to the design) regions of objects placed on a table [3, 4, 15, 16, 18, 20, 26]. UMD [18] has 28,843 images of objects placed on a blue rotating table in the same environment, and annotated with different affordance classes for each object category. Similarly, Multi-View [16] has 47,210 images of objects placed on a white rotating table in the same environment, but includes a larger number of affordance classes and object categories. TRANS-AFF [15] has 1,346 images of transparent objects on a table, increasing the difficulty to predict affordances. IIT-AFF [20] has 8,835 images of objects placed in a cluttered scene to better reflect a scenario with occlusions. CAD120-AFF [26] has 3,090 images, sampled from videos of humans performing activities in a realistic setting, e.g., kitchen, office. FPHA-AFF [14] has 4,300 images of hand-held objects acquired from an egocentric point of view. However, 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.

All the previous datasets are manually annotated but limited in their size, making the training of visual affordance segmentation model a problem. The need for thousands of labelled data has pushed towards the generation of synthetic data. UMD-Synth [4] simulates UMD in a synthetic manner and has 37,200 images of objects in different backgrounds and poses, but provides only one affordance per object region. AFF-Synth [3] has 30,245 images generated by using domain randomization to overcome the gap between simulated and real data. Unlike UMD-Synth, each image of AFF-Synth includes multiple objects. These synthetic datasets, however, are still small in size and do not contain occlusions, especially in the case of objects held by a person. Even though CAD120-AFF and FPHA-AFF address the scenario of hand-held objects, a suitable large-size dataset (synthetic or real) for segmenting affordances of objects with different properties (e.g., transparency, shape, size) and under hand-occlusions is missing.

¹Data, code, and trained models are available at <https://apicis.github.io/projects/acanet.html>

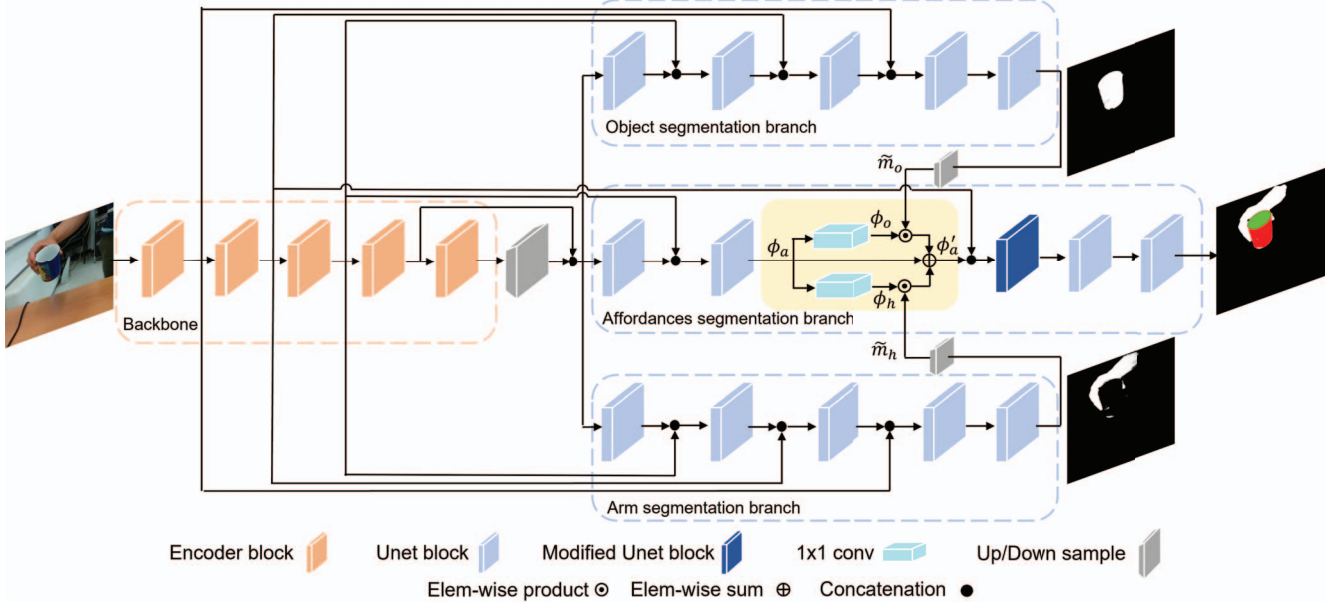


Figure 2: ACANet, our proposed model for Arm-Container Affordance segmentation of hand-held containers. The fusion block is highlighted in yellow.

3. The Arm-Container Affordance Network

In this section, we present our multi-branch architecture and the fusion module to perform affordance segmentation (see Fig. 2). We name the model Arm-Container Affordance Network (ACANet) as we consider containers for food and drinks, and *graspable* and *contain* as affordance classes. Assuming the input image has a correctly detected object of interest at the center, the model identifies the classes *graspable*, *contain*, *arm*, and *background*.

Let $I \in \mathbb{R}^{W \times H \times 3}$ be an exocentric RGB image, where W and H are the image width and height respectively, of a person holding an object, o . Let $S \in \{0, \dots, C-1\}^{W \times H}$ be a segmentation map where each pixel is one of the C classes that include *background*, *arm*, and the affordances on the visible surface of the object o . The objective is to design a model $f(I)$ that assigns each pixel in I to one of the corresponding C classes and that separates the arm from the object, while identifying the graspable regions of interest on the object from the regions to avoid.

3.1. Multi-branch architecture

We devise our multi-branch architecture starting from a UNet-like architecture [23] that uses skip connections between the encoder and the decoder to preserve the information and help the gradient flow during back-propagation. We replace the encoder with a ResNet-18 [13] to include residual connections within the convolutional layers of the encoder, easing the optimization problem [13]. Moreover, ResNet keeps the first pooling layer and replaces the others using a stride 2 in convolutional layers, unlike the origi-

nal UNet encoder that halves the resolution of tensors using only pooling layers. To double the resolution in the decoders, we replace the UNet convolutional layers with non-trainable up-sampling layers based on nearest interpolation. For each decoder, we also modified the last convolutional layer from 1×1 kernel and no padding (original UNet) to 3×3 kernel, stride 1, and padding. This allows the last convolution to consider neighboring pixels information.

This architecture outputs a tensor $S \in [0, 1]^{W \times H \times C}$ where each channel predicts one of the classes independently, and each pixel in a channel map is a probability such that $\sum_{c=0}^{C-1} S_c(i, j) = 1$, with $i \in \{1, \dots, W\}$ and $j \in \{1, \dots, H\}$ being the width and height indices, respectively. Including the class *arm* in S already handles the affordance segmentation of hand-occluded containers. However, we experimentally observed that the model first learns the classes with higher number of pixels in the annotation, affecting the prediction accuracy of the other classes.

We include two additional decoder branches that specialise in the segmentation of the arm and of the visible region of the object. Segmenting the object helps the model learn the area of the image where the affordances are. For simplicity, we refer to the three decoder branches as Arm, Object, and Affordance segmentation. The *Arm segmentation* branch predicts a probability map, $m_h \in [0, 1]^{W \times H}$, that separates the region associated to the arm (composition of forearm and hand) from all the rest. The *Object segmentation* branch predicts a probability map, $m_o \in [0, 1]^{W \times H}$, that separates the region associated to the visible object (held by the person hand) from all the rest. The *Affordance*

segmentation branch fuses the feature maps with the arm and object maps, and predicts the segmentation tensor \mathcal{S} .

3.2. Feature separation and fusion

Object and arm segmentation alone are insufficient to improve the segmentation accuracy under hand-occlusions. We therefore design a module that merges intermediate feature maps $\phi_a \in \mathbb{R}^{C' \times W' \times H'}$, with the down-sampled masks extracted by the *Arm segmentation* branch ($\tilde{m}_h \in [0, 1]^{W' \times H'}$) and the *Object segmentation* branch ($\tilde{m}_o \in [0, 1]^{W' \times H'}$), respectively. We compute the intermediate feature maps within the *Affordance segmentation* branch by using two UNet blocks that process the feature maps outputted by the backbone. The object and arm masks are down-sampled using bi-linear interpolation to match the size of the intermediate feature maps.

Instead of directly combining the features maps with the segmentation masks, we first learn specialised features in the object and arm regions. Specifically, we convolve ϕ_a with a set of 1×1 filters to obtain the feature map related to the object, $\phi_o \in \mathbb{R}^{C' \times W' \times H'}$ (where C' is the number of filters), and with another set of 1×1 filters to obtain the feature map related to the hand, $\phi_h \in \mathbb{R}^{C' \times W' \times H'}$. We then perform a pixel-wise weighting of the feature maps ϕ_o and ϕ_h with the corresponding segmentation mask \tilde{m}_o and \tilde{m}_h . This highly penalises the features outside of the predicted object (or arm) region. The merged feature maps are aggregated with the initial intermediate features as

$$\phi'_a = \phi_a + (\phi_h \odot \tilde{m}_h) + (\phi_o \odot \tilde{m}_o), \quad (1)$$

where \odot is the Hadamard product, i.e., the element-wise product between each feature map in ϕ_o, ϕ_h and the down-sampled segmentation masks \tilde{m}_o, \tilde{m}_h ².

3.3. Predicting object affordances and the hand

The *Affordance segmentation* branch uses the fused feature maps, ϕ'_a , as input to three UNet blocks to predict the output segmentation tensor \mathcal{S} . Note that the feature maps ϕ'_a are concatenated via skip connection with the corresponding intermediate feature maps in the encoder before the first UNet block (see Fig. 2). We modify the first UNet block after the fusion module to improve the processing of the feature maps. Specifically, we increase the number of output channels of the first convolutional filter from 64 to 128, and we add another convolutional layer to decrease the channels from 128 to 64. Furthermore, we avoid concatenating low-level information (e.g., edges) from the backbone to the last three UNet blocks via skip connections. This design choice helps the model preserve the changes in the feature maps ϕ'_a and predict affordances in the object region. The final segmentation map is: $\mathcal{S} = \operatorname{argmax}_c \mathcal{S}$.

²Mathematical simplification as \tilde{m}_o and \tilde{m}_h should be repeated for each channel of the feature maps ϕ_o and ϕ_h .

3.4. Loss functions

To train ACANet, we use a linear combination of a Dice loss [17] and two binary cross-entropy losses as:

$$\mathcal{L} = \mathcal{L}_a + \lambda_o \mathcal{L}_o + \lambda_h \mathcal{L}_h, \quad (2)$$

where the Dice loss, \mathcal{L}_a , operates on the affordance branch outputs, the binary cross-entropy, \mathcal{L}_o , operates on the object branch output, and the binary cross-entropy, \mathcal{L}_h , operates on the arm branch output. The hyper-parameters λ_o and $\lambda_h \in \mathbb{R}$ control the impact of object and hand segmentation losses, respectively. \mathcal{L} allows each branch to specialize for their segmentation task and influences the backbone to learn a common representation for all the branches.

The Dice loss, \mathcal{L}_a , addresses the imbalance problem between classes, as the majority of pixels can be labelled as *background* [27]. Given a batch of B predicted segmentation and corresponding annotations, the Dice loss is³:

$$\mathcal{L}_a = 1 - \frac{1}{C} \sum_{c=0}^{C-1} \frac{2 \sum_{b=1}^B \sum_{l=1}^{WH} y_{l,b}^c \hat{y}_{l,b}^c}{\epsilon + \sum_{n=1}^B \sum_{l=1}^{WH} \hat{y}_{l,b}^c + y_{l,b}^c}, \quad (3)$$

where $\hat{y} \in [0, 1]^{WH \times C}$ and $y \in \{0, 1\}^{WH \times C}$ are the reshaped predictions and annotations, respectively, with $\sum_{c=0}^{C-1} y_i^c = 1$.

The binary cross-entropy loss is used in binary classification and semantic segmentation task, considering each pixel as independent from the others. Given a batch of B predicted object segmentation masks and corresponding annotations, the binary cross-entropy loss for the object is:

$$\mathcal{L}_o = -\frac{1}{B} \sum_{b=1}^B \sum_{l=1}^{WH} v_{l,b} \log(\hat{v}_{l,b}) + (1 - v_{l,b}) \log(1 - \hat{v}_{l,b}), \quad (4)$$

where $\hat{v} \in [0, 1]^{WH}$ is the reshaped vector of m_o and $v \in \{0, 1\}^{WH}$ is the reshaped vector of the corresponding annotation. The binary cross-entropy for the hand, \mathcal{L}_h , is similarly computed using the reshaped vector of m_h and corresponding reshaped annotation.

3.5. Training with mixed-reality images

Training ACANet requires a large dataset with (exocentric) images of hand-occluded objects and segmentation annotation of both arm, object, and affordances. Such a dataset was not available, and collecting and manually annotating a new dataset is challenging, expensive, and time-consuming. We therefore complement an existing dataset, which has mixed-reality images of hand-occluded containers for object pose estimation [29], with visual affordance annotations. Using mixed-reality datasets can easily scale the generation of a larger number of images under different realistic backgrounds. Moreover, some existing works

³The margin $\epsilon = 10^{-7}$ avoids numerical issues when $\hat{y} = y = 0$.

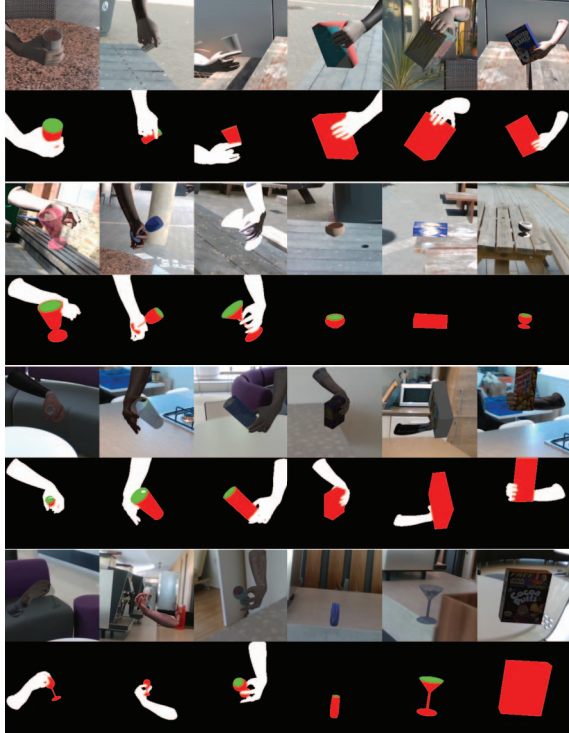


Figure 3: Samples of cropped RGB images and segmentation maps of arms and object affordances from the annotated mixed-reality dataset, CHOC-AFF. Key: \blacksquare *background*, $\color{red}\blacksquare$ *graspable*, $\color{green}\blacksquare$ *contain*, \square *arm*

on hand-object reconstruction or object pose estimation achieved good performance when training on mixed-reality datasets despite the domain gap [11, 28].

We use the publicly available CORSMAL Hand-Occluded Containers (CHOC) dataset that has 138,240 RGB images of 48 synthetic containers pseudo-realistically rendered on top of 30 different real backgrounds [29]. The dataset has 8,640 images with objects placed on top of a flat surface and 129,600 images of hand-held objects rendered in various locations and poses above the flat surface in the scene. The 48 containers are evenly distributed among 3 categories (box, stem, non-stem) and vary in their physical properties, such as size and shape, and appearance (textures, transparency). Hand-held objects were generated using synthetic forearms that hold the synthetic containers with three different visually plausible grasps and are orientated towards the pointing direction of the camera to simulate a potential offering of the object.

These characteristics and the available generation pipeline allows us to easily extend CHOC with the annotation of the *graspable* and *contain* affordances in addition to the existing annotations of the segmentation masks for the *arm* and *object*. We annotate the affordances on the surfaces

of the 3D CAD models of the 48 containers using Blender⁴. We then use the object poses annotated in CHOC to project the models with affordances on the image plane and render the affordance maps. We add the hand mask to the affordance map by pixel replacement. For simplicity, we refer to this version of the dataset with annotations of the affordances as CHOC-AFF. Fig. 3 shows sampled images and annotations of CHOC-AFF. We use CHOC-AFF for training and testing ACANet and other models, and we evaluate the generalisation of the models to real images.

4. Validation

4.1. Methods under comparison

We compare ACANet against a baseline and two state-of-the-art methods: ResNet18-UNet (RN18-U), ResNet50-FastFCN (RN50-F) [14], and DRNAtt [8]. RN18-U is the single-branch baseline behind ACANet and has a UNet-like architecture with a ResNet-18 based encoder. DRNAtt is the best performing model on the UMD dataset [2]. We implemented RN18-U and re-implemented DRNAtt to directly segment both the object affordances and the arm. We changed the channels of the last layer in RN50-F⁵ to segment only affordances of objects as per the original implementation in case of egocentric data [14].

4.2. Experimental setup

For our experiments, we split CHOC-AFF into *training* set and *validation* set to train the models, and two *testing* sets to evaluate the models generalisation to different backgrounds and different object instances. We also select and annotate images from two existing public datasets for hand-object pose estimation or reconstruction to evaluate the models in real conditions.

For CHOC-AFF, the training set has 89,856 images with 26 out of 30 backgrounds and 36 out of 48 containers (12 per object category). The validation set has 17,280 images with all 30 backgrounds and 6 container instances (2 per object category) different from the ones in training set. The first testing set has 13,824 images and evaluates the generalisation performance of the models to the same training object instances in 4 backgrounds not seen during training. The second testing set has 17,280 images and evaluates the generalisation performance of the models to 6 object instances (2 per object category) not seen during training in all 30 backgrounds [29].

For the two testing sets in real conditions, we consider HO-3D [9] and CCM [31] due to the presence of various challenges, such as presence of the human body, real interactions, and different object instances and hand-object

⁴<https://www.blender.org/>

⁵Original RN50-F implementation [30] is available at <https://github.com/wuhuikai/FastFCN>

poses. HO-3D is a multi-view video dataset of people manipulating different types of objects. We selected 150 frames of mugs and boxes as containers from the lateral and frontal cameras (with respect to the arm), keeping a diversity in object and hand poses. We used the segmentation provided by the authors as annotation for the classes *arm*⁶ and *graspable*, whereas we manually annotated the class *contain*. CCM is a dataset of multi-view sequences of people manipulating containers with different contents, and then offering the objects to a fixed robot arm. The offering phase allows us to evaluate the models under more realistic human grasps and object poses in a human-robot collaboration scenario, and more challenging conditions caused by different background and lighting settings. Moreover, containers can vary in their physical appearance (e.g., transparency, texture) or be affected by the presence of content. We selected the last frame (offering phase) of 150 sequences from a side perspective, diversifying objects, hand poses, and scene color settings. We manually annotated the affordance classes *contain* and *graspable*, and the class *arm* of only the hand(s) in contact with the offered container.

4.3. Training details and parameters settings

During training, we fine-tune all encoders, pre-trained on ImageNet [25], to start from a better initialisation than random. For ACANet, we set the hyper-parameters λ_o and λ_h to 1. For all models, we set the batch size to 2, the initial learning rate to 0.001, and we use the mini-batch Gradient Descent algorithm as optimizer with a momentum of 0.9 and a weight decay of 0.0001. For all models except RN50-F, we schedule the learning rate to decrease 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. For RN50-F, we set the learning rate schedule following the original setup. We use a Dice Loss to penalise the errors for RN18-U. We use the standard cross-entropy loss for DRNAtt and RN50-F (with auxiliary weight set to 0.2). We use early stopping with a patience of 10 epochs to reduce overfitting, and set the maximum number of epochs to 100.

Images can be of different resolutions and therefore we apply a cropping square window of fixed size to avoid distortions or adding padding. Assuming a perfect object detector, we crop a $W \times W$ window around the center of the bounding box obtained from the object mask annotation to restrict the visual field and obtain an object centric view. However, the cropping window can go out of the support of the image if the bounding box is close to the image border. In this case, we extend the side of the window that is inside the image support to avoid padding. In case the bounding box is bigger than the cropping window, we crop the image inside the bounding box and resize it to the window size. We apply this cropping procedure to all the images both in

⁶Note that the forearm is not annotated.

training and testing phases.

During training, we also use data augmentation to further increase the diversity of the images. Specifically, for each input image, we apply the following sequence of transformations: resize by a factor randomly sampled in the interval $[1, 1.5]$ to avoid degrading quality; center crop the resized image with a $W \times H$ window to restore the original image resolution; and horizontal flip with a probability of 0.5 to simulate the other arm. We set the window size to $W = H = 480$ (minimum resolution for RN50-F), as higher resolutions degrade the image quality.

4.4. Performance measures

To evaluate and compare the models, we compute the per-class precision, recall, and Jaccard Index as percentages across all images of a given dataset $\mathcal{D} = \{I_n | n = 1, \dots, N\}$. Precision measures the percentage of true positives among all positive predicted pixels. Recall measures the percentage of true positive pixels with respect to the total number of positive pixels. The Jaccard Index measures how much two regions with the same support are comparable (Intersection over Union or IoU).

To obtain these performance measures, we first compute true positives (TP), false positives (FP), and false negatives (FN) between prediction S_n and annotation G_n for each RGB image $I_n \in \mathcal{D}$ and for each class c . A true positive is a pixel $x \in I_n$ that is predicted as class c in S_n and the corresponding pixel in G_n is annotated as c . A false positive is a pixel $x \in I_n$ that is predicted as class c in S_n , but not annotated as c in G_n . A false negative is a pixel $x \in I_n$ that is not predicted as class c in S_n , but the corresponding pixel in G_n is annotated as c .

We therefore compute the per-class precision, P , the per-class recall, R , and the per-class Jaccard Index, J , as:

$$P = \frac{\sum_{n=1}^N \sum_{x \in I_n} TP}{\sum_{n=1}^N \sum_{x \in I_n} TP + FP}, \quad (5)$$

$$R = \frac{\sum_{n=1}^N \sum_{x \in I_n} TP}{\sum_{n=1}^N \sum_{x \in I_n} TP + FN}, \quad (6)$$

$$J = \frac{\sum_{n=1}^N \sum_{x \in I_n} TP}{\sum_{n=1}^N \sum_{x \in I_n} TP + FP + FN}. \quad (7)$$

Note that we keep the notation simple and we did not include the index of the class c in the above equations, as all the results will refer to the per-class measures.

4.5. Results and discussion

Table 1 compares the performance of the models on the mixed-reality and real testing sets. For the discussion and ranking of the methods, we consider J as the reference performance measure. Overall, ACANet outperforms the other

Table 1: Comparison of the affordance and arm segmentation results between the models on the two mixed-reality testing sets and on the two real testing sets.

Testing set	Model	<i>graspable</i>			<i>contain</i>			<i>arm</i>		
		<i>P</i>	<i>R</i>	<i>J</i>	<i>P</i>	<i>R</i>	<i>J</i>	<i>P</i>	<i>R</i>	<i>J</i>
CHOC-B	RN50-F	97.33	95.72	93.27	89.83	91.94	83.27	-	-	-
	RN18-U	96.79	96.44	93.45	84.94	93.16	79.95	96.55	96.46	93.24
	DRNAtt	96.38	97.04	93.63	91.84	90.63	83.88	96.94	97.19	94.30
	ACANet	97.09	96.60	93.88	89.46	94.67	85.17	96.48	96.52	93.24
CHOC-I	RN50-F	96.55	95.35	92.20	90.20	74.27	68.73	-	-	-
	RN18-U	96.33	96.35	92.94	88.97	74.32	68.04	96.67	96.91	93.78
	DRNAtt	95.85	96.74	92.85	90.48	71.08	66.13	97.00	96.88	94.07
	ACANet	96.36	96.51	93.11	88.72	76.68	69.86	96.94	96.77	93.90
HO-3D	RN50-F	95.61	18.29	18.14	90.69	79.57	73.56	-	-	-
	RN18-U	85.85	72.53	64.79	88.21	87.61	78.42	61.80	41.03	32.73
	DRNAtt	75.42	44.08	38.54	87.26	18.75	18.25	50.23	0.32	0.32
	ACANet	89.72	80.78	73.93	79.20	90.43	73.07	61.95	53.02	40.00
CCM	RN50-F	6.14	87.87	6.09	13.51	33.11	10.61	-	-	-
	RN18-U	13.69	78.69	13.20	31.92	42.44	22.28	44.21	42.53	27.68
	DRNAtt	6.37	95.09	6.35	0.00	0.00	0.00	4.47	0.24	0.23
	ACANet	10.22	86.50	10.06	45.40	37.46	25.83	49.47	45.35	31.00

Highlighted in **bold** the proposed model, and the best performing results.
 KEY – RN50-F: ResNet50-FastFCN [14], RN18-U: ResNet18-UNET, DRNAtt [8],
J: per-class Jaccard Index, CHOC-B: the CHOC-AFF testing set with new backgrounds, CHOC-I: the CHOC-AFF testing set with new instances.

models on all datasets and for most of the classes. This means that ACANet achieves better generalisation to other backgrounds and object instances in the testing sets.

The performance on the mixed-reality testing sets is similar among the models and *J* is higher than 90% for the classes *graspable* and *arm*. Models predict a low number of false positives and false negatives, resulting in a high value for precision and recall ($R, P > 95\%$). The class *contain* is the most challenging and *J* drops to the interval [79%, 85%] for the first testing set with unseen backgrounds (CHOC-B) and to the interval [66%, 70%] for the second testing set with unseen object instances (CHOC-I). ACANet outperforms the other models for the classes *graspable* and *contain* on both testing sets, whereas DRNAtt has the highest Jaccard Index for the class *arm*. RN50-F has the lowest performance, except for *J* in *contain* and *P* in *graspable*. This result shows that adding the *arm* class helps improve the performance. In CHOC-B, models except DRNAtt predict more false positives than negatives for the class *contain* ($P \in [84\%, 91\%]$, $R > 90\%$). On the contrary, the number of false negatives for the class *contain* is higher than false positives ($R \in [71\%, 76\%]$, $P > 88\%$) in CHOC-I.

The HO-3D and CCM testing sets allow us to assess the generalisation capabilities of the models to images acquired in real scenarios, given the known problem of the gap between synthetic and real data. As expected, performance of the models is lower in the real testing sets than in the mixed reality testing sets due to the domain shift. This can be observed especially for the classes *graspable* and *arm*.

In HO-3D, ACANet outperforms the other models for

the classes *graspable* and *arm*. However, all models tend to predict the wrong class in the *graspable* and *arm* regions ($P > R$), and even ACANet has a high number of false positives and false negatives for the class *arm* ($J = 40\%$). The performance for the class *arm* is penalised for all models due to the lack of annotation of the forearm, and the presence of the human body and challenging arm poses. For the class *contain*, DRNAtt predicts a high number of false negatives that affect the final performance ($J = 18.25\%$), whereas the other models predict a lower number of false positives than DRNAtt, resulting in a higher Jaccard index ($J \in [73.07\%, 78.42\%]$).

In CCM, the tablecloth and the presence of the human body are the main challenges for the models, causing a performance drop compared to the other datasets. In the presence of the tablecloth, models tend to predict *graspable* in most regions of the image. This results in a large difference between *P* and *R* (e.g., 76 percentage points for ACANet). ACANet achieves the best performance for the classes *contain* ($J = 25.83\%$) and *arm* ($J = 31\%$). DRNAtt does not generalise to the real images of CCM with $J \leq 1\%$ for the class *arm* and $J = 6.35\%$ for the class *graspable*, whereas the class *contain* is not predicted. This is caused by a large number of false positives towards the class *graspable* ($P = 6.37\%$) and a large number of false positives and false negatives for the class *arm* ($P = 4.47\%$, $R = 0.24\%$). The higher performance of ACANet compared to DRNAtt and RN18-U shows that learning arm and object features separately is better than learning affordances directly.

Fig. 4 shows and compares the affordance segmentation results of the models on sample images from the four testing sets. We chose images with objects and hand poses that are challenging and never seen in training, e.g., holding a box from the bottom; and with different backgrounds and lighting or different object appearances. Visually, ACANet achieves the most accurate segmentation for the arm and the object affordances (see also Fig. 1). For CHOC-B (columns 1 and 2), the predictions have more false positives in the background, whereas there are false positives or false negatives in the object region for CHOC-I (columns 3 and 4). For HO-3D, ACANet predictions show better affordance segmentation than other models, but the number of false positives of the class *arm* increases in presence of the human face (column 5). All models show a high number of false positives for the class *graspable* in the CCM testing set, especially when there is a colorful tablecloth (7th column). However, the predictions of RN18-U and ACANet are close to the annotation when the setting is similar to the training one, i.e., there is no colorful tablecloth nor the human body/face, but just the arm holding the container (columns 6 and 8). In the last column, we also chose to show the results for a transparent cup as transparency can be a challenge for the models due to the not clearly defined




Figure 4: Comparison of the predicted affordance and hand masks of the models on sampled images from the four testing sets. The segmentation masks are overlaid on the RGB images. KEY - GT: ground-truth, ■ *graspable*, ■ *contain*, ■ *arm*.

Table 2: Models size and computational cost.

	ACANet	RN18-U	RN50-F	DRNAtt
# parameters	20.82	14.32	66.40	17.38
GFLOPs	85.17	38.28	239.33	149.90

KEY: # parameters: number of parameters (in millions), GFLOPs: Giga Floating-point operations, RN50-F: ResNet50-FastFCN [14], RN18-U: ResNet18-UNET, DRNAtt [8].

borders. For example, the background environment or the content in the container can influence the segmentation predicted by the models. In the sampled image, the cup is filled with a content, and the models are able to correctly predict both the *contain* and *graspable* regions, except for DRNAtt.

We also briefly discuss the complexity of the models measured in number of parameters (in millions) and number of operations (as Giga Floating-point or GFLOPs). Table 2 shows that RN50-F has the highest amount of parameters and operations, but the generalisation performance is worse than ACANet. DRNAtt has 3 millions fewer parameters than ACANet, but $1.7\times$ more operations than ACANet and the performance on real-data is worse in all classes. Finally, ACANet has an increased number of parameters (6 millions) and GFLOPs (2.2 times) compared to the base-

line RN18-U, thus contributing to the higher performance in almost all classes and testing sets.

5. Conclusion

We tackled the problem of visual affordance segmentation of hand-occluded containers and proposed ACANet, a multi-branch convolutional neural network that fuses object and hand segmentation mask with the affordance features to specialise filters for object and hand regions. Training ACANet on an annotated dataset with mixed-reality images of hand-held containers leads to better generalisation to real images with containers not seen in training and with new backgrounds (e.g., on the HO-3D dataset or on-the-fly acquired images). Moreover, ACANet outperforms alternative methods when segmenting the *graspable* area and the person’s *arm*, and can achieve a more accurate segmentation even when the objects are in more challenging poses caused by how the person holds the object.

As future work, we will use a larger set of objects, perform an ablation study on the fusion module components, reduce the computational costs necessary to run the model, and validate it in a human-robot collaboration scenario.

References

- [1] T. Apicella, A. Cavallaro, R. Berta, P. Gastaldo, F. Bellotti, and E. Ragusa. An affordance detection pipeline for resource-constrained devices. In *IEEE Int. Conf. Electronics, Circuits, and Systems*, 2021. 2
- [2] D. Chen, D. Kong, J. Li, S. Wang, and B. Yin. A survey of visual affordance recognition based on deep learning. *IEEE Trans. Big Data*, pages 1–20, 2023. 5
- [3] A. D. Christensen, D. Lehotský, M. W. Jørgensen, and D. Chrysostomou. Learning to segment object affordances on synthetic data for task-oriented robotic handovers. In *Brit. Mach. Vis. Conf.*, 2022. 1, 2
- [4] F. Chu, R. Xu, and P. Vela. Learning affordance segmentation for real-world robotic manipulation via synthetic images. *IEEE Robotics Autom. Lett.*, 4(2):1140–1147, 2019. 2
- [5] T. Do, A. Nguyen, and I. Reid. AffordanceNet: An end-to-end deep learning approach for object affordance detection. In *IEEE Int. Conf. Robotics Autom.*, 2018. 2
- [6] J. Fu, J. Liu, H. Tian, Y. Li, Y. Bao, Z. Fang, and H. Lu. Dual attention network for scene segmentation. In *Conf. Comput. Vis. Pattern Recognit.*, 2019. 2
- [7] J. J. Gibson and L. Carmichael. *The senses considered as perceptual systems*, volume 2. Houghton Mifflin Boston, 1966. 1
- [8] Q. Gu, J. Su, and L. Yuan. Visual affordance detection using an efficient attention convolutional neural network. *Neurocomputing*, 440:36–44, 2021. 1, 2, 5, 7, 8
- [9] S. Hampali, M. Rad, M. Oberweger, and V. Lepetit. Honnotate: A method for 3D annotation of hand and object poses. In *Conf. Comput. Vis. Pattern Recognit.*, 2020. 2, 5
- [10] M. Hassanin, S. Khan, and M. Tahtali. Visual affordance and function understanding: A survey. *ACM Computing Surveys*, 54(3):1–35, 2021. 1, 2
- [11] Y. Hasson, G. Varol, D. Tzionas, I. Kalevatykh, M. J. Black, I. Laptev, and C. Schmid. Learning joint reconstruction of hands and manipulated objects. In *Conf. Comput. Vis. Pattern Recognit.*, 2019. 5
- [12] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask R-CNN. In *IEEE Int. Conf. Comput. Vis.*, 2017. 2
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Conf. Comput. Vis. Pattern Recognit.*, 2016. 3
- [14] S.M Hussain.S, L. Liu, W. Xu, and C. Lu. FPHA-Afford: A domain-specific benchmark dataset for occluded object affordance estimation in human-object-robot interaction. In *IEEE Int. Conf. Image Process.*, 2020. 1, 2, 5, 7, 8
- [15] J. Jiang, G. Cao, T. Do, and S. Luo. A4t: Hierarchical affordance detection for transparent objects depth reconstruction and manipulation. *IEEE Robotics Autom. Lett.*, 7(4):9826–9833, 2022. 2
- [16] Z. O. Khalifa and S. A. A. Shah. Towards visual affordance learning: A benchmark for affordance segmentation and recognition. *arXiv:2203.14092v2 [cs.CV]*, 2022. 2
- [17] F. Milletari, N. Navab, and S. Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *Int. Conf. 3D Vision*, 2016. 4
- [18] A. Myers, C. L. Teo, C. Fermüller, and Y. Aloimonos. Affordance detection of tool parts from geometric features. In *IEEE Int. Conf. Robotics Autom.*, 2015. 1, 2
- [19] A. Nguyen, D. Kanoulas, D. Caldwell, and N. Tsagarakis. Detecting object affordances with convolutional neural networks. In *IEEE Int. Conf. Intell. Robot Syst.*, 2016. 1, 2
- [20] A. Nguyen, D. Kanoulas, D. Caldwell, and N. Tsagarakis. Object-based affordances detection with convolutional neural networks and dense conditional random fields. In *IEEE Int. Conf. Intell. Robot Syst.*, 2017. 1, 2
- [21] E. Ragusa, C. Gianoglio, S. Dosen, and P. Gastaldo. Hardware-aware affordance detection for application in portable embedded systems. *IEEE Access*, 9:123178–123193, 2021. 1, 2
- [22] S. Rezapour Lakani, A. Rodríguez-Sánchez, and J. Piater. Towards affordance detection for robot manipulation using affordance for parts and parts for affordance. *Auton. Robots*, 43:1155–1172, 2019. 1
- [23] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Int. Conf. Med. Image Comput. Comput. Assist. Interv.*, 2015. 1, 3
- [24] P. Rosenberger, A. Cosgun, R. Newbury, J. Kwan, V. Ortenzi, P. Corke, and M. Grafinger. Object-independent human-to-robot handovers using real time robotic vision. *IEEE Robotics Autom. Lett.*, 6(1):17–23, 2020. 1
- [25] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *Int. J. Comput. Vis.*, 115(3):211–252, 2015. 6
- [26] J. Sawatzky, A. Srikantha, and J. Gall. Weakly supervised affordance detection. In *Conf. Comput. Vis. Pattern Recognit.*, 2017. 2
- [27] C. Sudre, W. Li, T. Vercauteren, S. Ourselin, and M. Jorge Cardoso. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In *Int. Conf. Med. Image Comput. Comput. Assist. Interv., DLMIA and ML-CDS Workshops*. 2017. 4
- [28] H. Wang, S. Sridhar, J. Huang, J. Valentin, S. Song, and L. J. Guibas. Normalized object coordinate space for category-level 6D object pose and size estimation. In *Conf. Comput. Vis. Pattern Recognit.*, 2019. 5
- [29] X. Weber, A. Xompero, and A. Cavallaro. A mixed-reality dataset for category-level 6D pose and size estimation of hand-occluded containers. *arXiv:2211.10470v1 [cs.CV]*, 2022. 4, 5
- [30] H. Wu, J. Zhang, K. Huang, K. Liang, and Y. Yu. FastFCN: Rethinking dilated convolution in the backbone for semantic segmentation. *arXiv:1903.11816v1 [cs.CV]*, 2019. 2, 5
- [31] A. Xompero, R. Sanchez-Matilla, R. Mazzon, and A. Cavallaro. CORSMAL Containers Manipulation, 2020. (1.0) [Data set]. Queen Mary University of London. <https://doi.org/10.17636/101CORSMAL1>. 2, 5
- [32] W. Yang, C. Paxton, A. Mousavian, Y. Chao, M. Cakmak, and D. Fox. Reactive human-to-robot handovers of arbitrary objects. In *IEEE Int. Conf. Robotics Autom.*, 2021. 1

- [33] C. Yin and Q. Zhang. Object affordance detection with boundary-preserving network for robotic manipulation tasks. *Neural. Comput. Appl.*, 34(20):17963–17980, 2022. [2](#)
- [34] C. Yin, Q. Zhang, and W. Ren. A new semantic edge aware network for object affordance detection. *J. Intelligent & Robotic Systems*, 104(1):1–16, 2022. [1](#), [2](#)
- [35] Y. Zhang, H. Li, T. Ren, Y. Dou, and Q. Li. Multi-scale fusion and global semantic encoding for affordance detection. In *Int. Joint Conf. on Neural Networks*, 2022. [2](#)
- [36] X. Zhao, Y. Cao, and Y. Kang. Object affordance detection with relationship-aware network. *Neural. Comput. Appl.*, 32(18):14321–14333, 2020. [2](#)