

AffordanceNet: An End-to-End Deep Learning Approach for Object Affordance Detection

Thanh-Toan Do^{1,†}, Anh Nguyen^{2,†}, Ian Reid¹

Abstract—We propose AffordanceNet, a new deep learning approach to simultaneously detect multiple objects and their affordances from RGB images. Our AffordanceNet has two branches: an object detection branch to localize and classify the object, and an affordance detection branch to assign each pixel in the object to its most probable affordance label. The proposed framework employs three key components for effectively handling the multiclass problem in the affordance mask: a sequence of deconvolutional layers, a robust resizing strategy, and a multi-task loss function. The experimental results on the public datasets show that our AffordanceNet outperforms recent state-of-the-art methods by a fair margin, while its end-to-end architecture allows the inference at the speed of 150ms per image. This makes our AffordanceNet well suitable for real-time robotic applications. Furthermore, we demonstrate the effectiveness of AffordanceNet in different testing environments and in real robotic applications. The source code is available at <https://github.com/nqanh/affordance-net>.

I. INTRODUCTION

An object can be described by various visual properties such as color, shape, or physical attributes such as weight, volume, and material. Those properties are useful to recognize objects or classify them into different categories, however they do not imply the potential actions that human can perform on the object. The capability to understand functional aspects of objects or *object affordances* has been studied for a long time [1]. Unlike other visual or physical properties that mainly describe the object alone, affordances indicate functional interactions of object parts with humans. Understanding object affordances is, therefore, crucial to let an autonomous robot interact with the objects and assist humans in various daily tasks.

The problem of modeling object affordances can be considered in different ways. Castellini et al. [2] defined affordances in terms of human hand poses during the interaction with objects, while in [3] the authors studied object affordances in the context of human activities. In this work, similar to [4], we consider object affordances at *pixel* level from an image, i.e., a group of pixels which shares the same object functionality is considered as one affordance. The advantage of this approach is we can reuse the strong state of the art from the semantic segmentation field, while there is no extra information such as interactions with human is needed. Detecting object affordances, however, is a more difficult

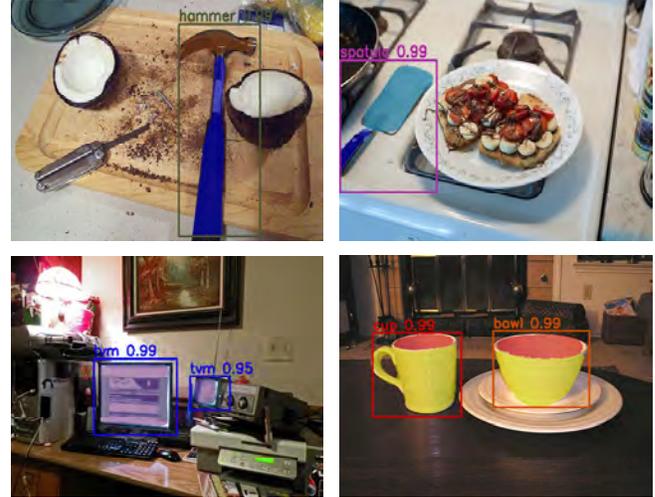


Fig. 1. Simultaneous affordance and object detection. Some example results of our AffordanceNet, which detects both objects and their multiple affordance classes using an end-to-end architecture.

task than the classical semantic segmentation problem. For example, two object parts with different appearances may have the same affordance label. It is because the affordance labels are based on the abstract concepts of human actions on the object. Furthermore, it is also essential for an affordance detection method to run in real-time and generalize well on unseen objects.

In many robotic applications, recognizing object affordances is essential, however the robot may still require more information to complete tasks. For example, to pour the water from a *bottle* into a *bowl*, the robot not only has to detect object affordances such as *grasp*, *contain*, but also be able to localize and recognize the relevant objects (i.e., *bottle*, *bowl*) [5]. In order to address this, the work in [6] proposed to use two sequential deep neural networks, one for object detection and one for affordance detection. However, by using two sequential deep networks, it is time consuming during testing, meaning that approach may not be applicable for real-time applications. In this work, we overcome this limitation by using an end-to-end architecture. Our proposed architecture jointly optimizes the object detection and the affordance detection using a multi-task loss function. We show that the proposed method reduces the complexity during training and testing while improves the overall affordance detection accuracy. Fig 1 shows some example results of our network, which can simultaneously detect the objects and their multiple affordance classes.

¹Thanh-Toan Do and Ian Reid are with the Australian Centre for Robotic Vision (ACRV), University of Adelaide. {Thanh-Toan.Do, Ian.Reid}@adelaide.edu.au

²Anh Nguyen is with the Department of Advanced Robotics, IIT, Italy. Anh.Nguyen@iit.it

[†] Both authors contributed equally to this work.

In computer vision, simultaneous object detection and object segmentation is becoming more popular [7]. Recent advances in deep learning allow training the detection branch and segmentation branch effectively together. The intuition is that although the detection branch uses object bounding boxes and the segmentation branch uses pixel labels, they can share the same feature maps from the convolutional backbone. The authors in [8] followed this methodology to build a network for instance segmentation problem and achieved state-of-the-art results. Our work is built upon the works of [6] and [8]. However, we differ from [6] by using an end-to-end architecture. We also differ from [8] by having new components, i.e., a new loss function, a sequence of deconvolutional layers, and a robust resizing strategy, for handling the problem of multiple affordance classes. We show that these new components are key factors to achieve high affordance detection accuracy. The experimental results on the public datasets show that our AffordanceNet outperforms recent state-of-the-art methods by a fair margin, while its end-to-end architecture allows inference on a test image in just 150ms. We also demonstrate the effectiveness of AffordanceNet in different testing environments, and in real robotic applications.

The remainder of this paper is organized as follows. We review the related work in Section II. We then describe our end-to-end architecture for jointly learning object detection and affordance detection in Section III. In Section IV, we present the extensive experimental results on the public datasets and the robotic demonstration on a full-size humanoid robot WALK-MAN. Finally, we conclude the paper in Section V.

II. RELATED WORK

The problem of understanding affordances at the pixel level has been termed “object part labelling” in the computer vision community, while it is more commonly known as “affordance detection” in robotics. In computer vision, the concept of affordances is not restricted to objects, but covers a wide range of applications, from understanding human body parts [9] to environment affordances [10] [11], while in robotics, researchers focus more on the real-world objects that the robot can interact with [4]. In [12], the authors used predefined primary tools to infer object functionalities from 3D point clouds. The work in [13] proposed to combine the global object poses with its local appearances to detect grasp affordances. In [14], the authors introduced a method to detect object affordances via object-action interactions from human demonstrations. In [4], the authors used hand-designed geometric features to detect object affordances at pixel level from RGB-D images.

With the rise of deep learning, recent works relied on deep neural networks for designing affordance detection frameworks. The work in [15] used two deep neural networks to detect grasp affordances from RGB images. The work in [16] used deep features from Convolutional Neural Networks (CNN) for detecting affordances from RGB-D images. It

gained a significant improvement over hand-designed geometric features [4]. Similar to [16], the work in [10] introduced multi-scale CNN to localize environment affordances. In [17], to avoid depending on costly pixel groundtruth labels, a weakly supervised deep learning approach was presented to segment object affordances. Recently, in [6], the authors proposed to use a deep learning-based object detector to improve the affordance detection accuracy on a real-world dataset. A limitation of that work is that its architecture is not end-to-end – i.e. two separate networks are used, one for object detection and one for affordance detection – and this is slow for both training and testing. Furthermore, by training two networks separately, the networks are not jointly optimal. In computer vision, the work of [18] introduced an end-to-end architecture to simultaneously detect and segment object instances. Recently, the authors in [8] improved over [18] by proposing a region alignment layer which effectively aligns the spatial coordinates of region of interests between the input image space and the feature map space.

The goal of this work is to simultaneously detect the objects (including the object location and object label) and their associated affordances. We follow the same concept in [6], however we use an end-to-end architecture instead of a sequential one. Our object affordance detection network can also be seen as a generalization of the recent state-of-the-art instance segmentation networks [8] [18]. In particular, our network can detect multiple affordance classes in the object, instead of binary class as in instance segmentation networks [8] [18].

III. JOINTLY AFFORDANCE AND OBJECT DETECTION

A. Problem Formulation

Inspired by [6] [8], our framework aims at simultaneously finding the object positions, object classes, and object affordances in images. Follow the standard design in computer vision, the object position is defined by a rectangle with respect to the top-left corner of the image; the object class is defined over the rectangle; the affordances are encoded at every pixel inside the rectangle. The region of pixels on the object that has the same functionality is considered as one affordance. Ideally, we want to detect all relevant objects in the image and map each pixel in these objects to its most probable affordance label.

B. AffordanceNet Architecture

We first describe three main components of our AffordanceNet: the Region of Interest (RoI) alignment layer (RoIAlign) [8] which is used to correctly compute the feature for an RoI from the image feature map; a sequence of convolution-deconvolution layers to upsample the RoI feature map to high resolution in order to obtain a smooth and fine affordance map; a robust strategy for resizing the training mask to supervise the affordance detection branch. We show that these components are the key factors to achieve high affordance detection accuracy. Finally, we present the whole AffordanceNet architecture in details. Fig. 2 shows an overview of our approach.

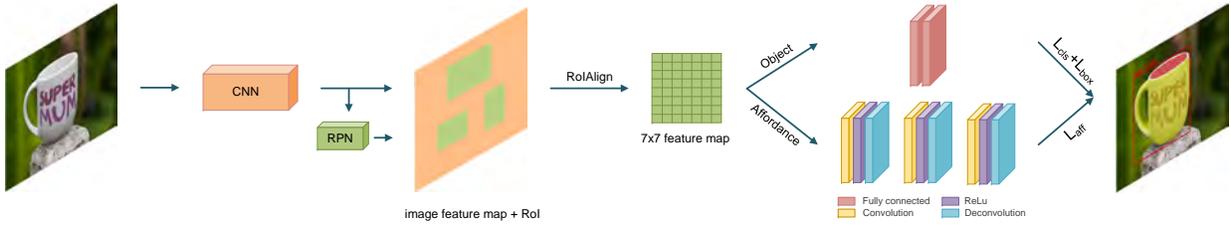


Fig. 2. An overview of our AffordanceNet framework. **From left to right:** A deep CNN backbone (i.e., VGG) is used to extract image features. The RPN shares weights with the backbone and outputs RoIs. For each RoI, the RoIAlign layer extracts and pools its features (from the image feature map, i.e., the *conv5_3* layer of VGG) to a fixed size 7×7 feature map. The object detection branch uses two fully connected layers for regressing object location and classifying object category. The object affordance detection branch consists of a sequence of convolutional-deconvolutional layers and ends with a softmax layer to output a multiclass affordance mask.

1) *RoIAlign*: One of the main components in the recent successful region-based object detectors such as Faster R-CNN [19] is the Region Proposal Network (RPN). This network shares weights with the main convolutional backbone and outputs bounding boxes (RoI / object proposal) at various sizes. For each RoI, a fixed-size small feature map (e.g., 7×7) is pooled from the image feature map using the RoIPool layer [19]. The RoIPool layer works by dividing the RoI into a regular grid and then max-pooling the feature map values in each grid cell. This quantization, however, causes misalignments between the RoI and the extracted features due to the harsh rounding operations when mapping the RoI coordinates from the input image space to the image feature map space and when dividing the RoI into grid cells.

In order to address this problem, the authors in [8] introduced the RoIAlign layer which properly aligns the extracted features with the RoI. Instead of using the rounding operation, the RoIAlign layer uses bilinear interpolation to compute the interpolated values of the input features at four regularly sampled locations in each RoI bin, and aggregates the result using max operation. This alignment technique plays an important role in tasks based on pixel level such as image segmentation. We refer the readers to [8] for a detailed analysis of the RoIAlign layer.

2) *Deconvolution for High Resolution Affordance Mask*: In recent state-of-the-art instance segmentation methods such Mask-RCNN [8] and FCIS [18], the authors used a small fixed size mask (e.g. 14×14 or 28×28) to represent the object segmentation mask. This is feasible since the pixel value in each predicted mask of RoI is binary, i.e., either foreground or background. We empirically found that using small mask size does not work well in the affordance detection problem since we have multiple affordance classes in each object. Hence, we propose to use a sequence of deconvolutional layers for achieving a high resolution affordance mask.

Formally, given an input feature map with size S_i , the deconvolutional layer performs the opposite operation of the convolutional layer to create a bigger output map with size S_o , in which S_i and S_o are related by:

$$S_o = s * (S_i - 1) + S_f - 2 * d \quad (1)$$

where S_f is the filter size; s and d are stride and padding parameters, respectively.

In practice, the RoIAlign layer outputs a feature map with size 7×7 . We use three deconvolutional layers to

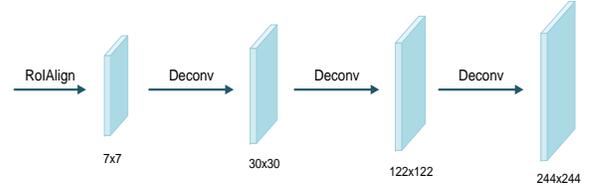


Fig. 3. A sequence of three deconvolutional layers to gradually upsample a 7×7 fixed size feature map to 244×244 .

upsample this map to higher resolution (see Fig. 3). The first deconvolutional layer has the padding $d = 1$, stride $s = 4$, and kernel size $S_f = 8$ to create the map with size 30×30 . Similarly, the second layer has the parameters ($d = 1$, $s = 4$, $S_f = 8$), and the third one has ($d = 1$, $s = 2$, $S_f = 4$) to create the final high resolution map with the size of 244×244 . It is worth noting that before each deconvolutional layer, a convolutional layer (together with ReLu) is used to learn features which will be used for the deconvolution. This convolutional layer can be seen as an adaptation between two consecutive deconvolutional layers. We analyze the effect of the affordance map size in Section IV-C.

3) *Robust Resizing Affordance Mask*: Similar to Mask-RCNN [8] and FCIS [18], our affordance detection branch requires a fixed size (e.g., 244×244) target affordance mask to supervise the training. During training, the authors in [8] [18] resized the original groundtruth mask of each RoI to the pre-defined mask size to compute the loss. This resizing step outputs a mask with values ranging from 0 to 1, which is thresholded (e.g., at 0.4) to determine if a pixel is background or foreground. However, using single threshold value does not work in our affordance detection problem since we have multiple affordance classes in each object. To address this problem, we propose a resizing strategy with multi-thresholding. Given an original groundtruth mask, without loss of generality, let $P = \{c_0, c_1, \dots, c_{n-1}\}$ be set of n unique labels in that mask, we first linearly map the values in P to $\hat{P} = \{0, 1, \dots, n-1\}$ and convert the original mask to a new mask using the mapping from P to \hat{P} . We then resize the converted mask to the pre-defined mask size and use the thresholding on the resized mask as follows:

$$\rho(x, y) = \begin{cases} \hat{p}, & \text{if } \hat{p} - \alpha \leq \rho(x, y) \leq \hat{p} + \alpha \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

where $\rho(x, y)$ is a pixel value in the resized mask; \hat{p} is one of values in \hat{P} ; α is the hyperparameter and is set to 0.005 in our experiments.

Finally, we re-map the values in the thresholded mask back to the original label values (by using the mapping from \hat{P} to P) to achieve the target training mask. Note that there is another way to achieve the fixed size target training mask. We can apply the resizing for each affordance label in the original groundtruth mask separately, i.e., when considering a label, that label is treated as foreground and other labels are treated as background. Then, we can combine the multiple resized masks to achieve the target training mask. However, from the practical point of view, this strategy is time consuming due to the multiple resizing for affordance classes in RoI.

4) *End-to-End Architecture*: Fig. 2 shows an overview of our end-to-end affordance detection network. The network is composed of two branches for object detection and affordance detection. Given an input image, we use the VGG16 [20] network as the backbone to extract deep features from the image. A RPN that shares the weights with the convolutional backbone is then used to generate candidate bounding boxes (RoIs). For each RoI, the RoIAlign layer extracts and pools its corresponding features (from the image feature map — the *conv5_3* layer of VGG16) into a 7×7 feature map. In the object detection branch, we use two fully connected layers, each with 4096 neurons, followed by a classification layer to classify the object and a regression layer to regress the object location. In the affordance detection branch, the 7×7 feature map is gradually upsampled to 244×244 to achieve high resolution map. The affordance branch uses a softmax layer to assign each pixel in the 244×244 map to its most probable affordance class. The whole network is trained end-to-end using a multi-task loss function.

C. Multi-Task Loss

In our aforementioned end-to-end architecture, the classification layer outputs a probability distribution $p = (p_0, \dots, p_K)$ over $K + 1$ object categories, including the background. As in [19], p is the output of a softmax layer. The regression layer outputs $K + 1$ bounding box regression offsets (each offset includes box center and box size): $t^k = (t_x^k, t_y^k, t_w^k, t_h^k)$. Each offset t^k corresponds to each class k . Similar to [21] [19] we parameterize for t^k , in which t^k specifies a scale-invariant translation and log-space height/width shift relative to an anchor box of the RPN. The affordance detection branch outputs a set of probability distributions $m = \{m^i\}_{i \in RoI}$ for each pixel i inside the RoI, in which $m^i = (m_0^i, \dots, m_C^i)$ is the output of a softmax layer defined on $C+1$ affordance labels, including the background.

We use a multi-task loss L to jointly train the bounding box class, the bounding box position, and the affordance map as follows:

$$L = L_{cls} + L_{loc} + L_{aff} \quad (3)$$

where L_{cls} is defined on the output of the classification layer;

L_{loc} is defined on the output of the regression layer; L_{aff} is defined on the output of the affordance detection branch.

The prediction target for each RoI is a groundtruth object class u , a groundtruth bounding box offset v , and a target affordance mask s . The values of u and v are provided with the training datasets. The target affordance mask s is the intersection between the RoI and its associated groundtruth mask. For pixels inside the RoI which do not belong to the intersection, we label them as background. Note that the target mask is then resized to a fixed size (i.e., 244×244) using the proposed resizing strategy in Section III-B.3. Specifically, we can rewrite Equation 3 as follows:

$$L(p, u, t^u, v, m, s) = L_{cls}(p, u) + I[u \geq 1]L_{loc}(t^u, v) + I[u \geq 1]L_{aff}(m, s) \quad (4)$$

The first loss $L_{cls}(p, u)$ is the multinomial cross entropy loss for the classification and is computed as follows:

$$L_{cls}(p, u) = -\log(p_u) \quad (5)$$

where p_u is the softmax output for the true class u .

The second loss $L_{loc}(t^u, v)$ is *Smooth L1* loss [22] between the regressed box offset t^u (corresponding to the groundtruth object class u) and the groundtruth box offset v , and is computed as follows:

$$L_{loc}(t^u, v) = \sum_{i \in \{x, y, w, h\}} Smooth_{L1}(t_i^u - v_i) \quad (6)$$

where

$$Smooth_{L1}(x) = \begin{cases} 0.5x^2 & \text{if } |x| < 1 \\ |x - 0.5| & \text{otherwise} \end{cases}$$

The $L_{aff}(m, s)$ is the multinomial cross entropy loss for the affordance detection branch and is computed as follows:

$$L_{aff}(m, s) = \frac{-1}{N} \sum_{i \in RoI} \log(m_{s_i}^i) \quad (7)$$

where $m_{s_i}^i$ is the softmax output at pixel i for the true label s_i ; N is the number of pixels in the RoI.

In Equation (4), $I[u \geq 1]$ is an indicator function which outputs 1 when $u \geq 1$ and 0 otherwise. This means that we only define the box location loss L_{loc} and the affordance detection loss L_{aff} only on the positive RoIs. While the object classification loss L_{cls} is defined on both positive and negative RoIs.

It is worth noting that our loss for affordance detection branch is different from the instance segmentation loss in [8] [18]. In those works, the authors rely on the output of the classification layer to determine the object label. Hence the segmentation in each RoI can be considered as a binary segmentation, i.e., foreground and background. Thus, the authors use per-pixel *sigmoid* layer and binary cross entropy loss. In our affordance detection problem, the affordance labels are different from the object labels. Furthermore, the number of affordances in each RoI is not binary, i.e., it is always bigger than 2 (including the background). Hence, we rely on a per-pixel *softmax* and a multinomial cross entropy loss.

TABLE I
PERFORMANCE ON IIT-AFF DATASET

	ED-RGB [16]	ED-RGBD [16]	DeepLab [24]	DeepLab- CRF [24]	BB-CNN [6]	BB-CNN- CRF [6]	AffordanceNet (ours)
contain	66.38	66.00	68.84	69.68	75.60	75.84	79.61
cut	60.66	60.20	55.23	56.39	69.87	71.95	75.68
display	55.38	55.11	61.00	62.63	72.04	73.68	77.81
engine	56.29	56.04	63.05	65.11	72.84	74.36	77.50
grasp	58.96	58.59	54.31	56.24	63.72	64.26	68.48
hit	60.81	60.47	58.43	60.17	66.56	67.07	70.75
pound	54.26	54.01	54.25	55.45	64.11	64.86	69.57
support	55.38	55.08	54.28	55.62	65.01	66.12	69.81
w-grasp	50.66	50.42	56.01	57.47	67.34	68.41	70.98
Average	57.64	57.32	58.38	59.86	68.57	69.62	73.35

D. Training and Inference

We train the network in an end-to-end manner using stochastic gradient descent with 0.9 momentum and 0.0005 weight decay. The network is trained on a Titan X GPU for 200k iterations. The learning rate is set to 0.001 for the first 150k and decreased by 10 for the last 50k. The input images are resized such that the shorter edge is 600 pixels, but the longer edge does not exceed 1000 pixels. In case the longer edge exceeds 1000 pixels, the longer edge is set to 1000 pixels, and the images are resized based on this edge. Similar to [8], we use 15 anchors in the RPN (5 scales and 3 aspect ratios). Top 2000 RoIs from RPN (with a ratio of 1:3 of positive to negative) are subsequently used for computing the multi-task loss. An RoI is considered positive if it has IoU with a groundtruth box of at least 0.5 and negative otherwise.

During the inference phase, we select the top 1000 RoIs produced by the RPN and run the object detection branch on these RoIs, followed by a non-maximum suppression [23]. From the outputs of the detection branch, we select the outputted boxes that have the classification score higher than 0.9 as the final detected objects. In case there are no boxes satisfying this condition, we select the one with highest classification score as the only detected object. We use the detected objects as the inputs for affordance detection branch. For each pixel in the detected object, the affordance branch predicts $C + 1$ affordance classes. The output affordance label for each pixel is achieved by taking the maximum across the affordance classes. Finally, the predicted 244×244 affordance mask of each object is resized to the object (box) size using the resizing strategy in Section III-B.3. In case there is the overlap between detected objects, similar to [6], the final affordance label is decided based on the affordance priority. For example, the “contain” affordance is considered to have low priority than other affordances since there may have other objects laid on it.

IV. EXPERIMENTS

A. Dataset and Baseline

IIT-AFF Dataset The IIT-AFF dataset is recently introduced in [6] and consists of 8,835 real-world images. This dataset is suitable for deep learning methods and robotic applications since around 60% of the images are from

TABLE II
PERFORMANCE ON UMD DATASET

	HMP [4]	SRF [4]	DeepLab [24]	ED-RGB [16]	ED-RGBD [16]	ED-RGB HHA [16]	AffordanceNet (ours)
grasp	0.367	0.314	0.620	0.719	0.714	0.673	0.731
w-grasp	0.373	0.285	0.730	0.769	0.767	0.652	0.814
cut	0.415	0.412	0.600	0.737	0.723	0.685	0.762
contain	0.810	0.635	0.900	0.817	0.819	0.716	0.833
support	0.643	0.429	0.600	0.780	0.803	0.663	0.821
scoop	0.524	0.481	0.800	0.744	0.757	0.635	0.793
pound	0.767	0.666	0.880	0.794	0.806	0.701	0.836
Average	0.557	0.460	0.733	0.766	0.770	0.675	0.799

ImageNet dataset [25], while the rest images are taken by the authors from cluttered scenes. In particular, this dataset contains 10 object categories, 9 affordance classes, 14,642 object bounding boxes, and 24,677 affordance regions at pixel level. We use the standard split as in [6] to train our network (i.e. 70% for training and 30% for testing).

UMD Dataset The UMD dataset [4] contains around 30,000 RGB-D images of daily kitchen, workshop, and garden objects. The RGB-D images of this dataset were captured from a Kinect camera on a rotating table in a clutter-free setup. This dataset has 7 affordance classes and 17 object categories. Since there is no groundtruth for the object bounding boxes, we compute the rectangle coordinates of object bounding boxes based on the affordance masks. We use only the RGB images of this dataset and follow the split in [4] to train and test our network.

Baseline As the standard practice, we use the F_{β}^w metric [26] to evaluate the affordance detection results. We compare our AffordanceNet with the following state-of-the-art approaches: DeepLab [24] with and without post processing with CRF (denoted as DeepLab and DeepLab-CRF), CNN with encoder-decoder architecture [16] on RGB and RGB-D images (denoted as ED-RGB and ED-RGBD), CNN with object detector (BB-CNN) and CRF (BB-CNN-CRF) [6]. For the UMD dataset, we also report the results from the geometric features-based approach (HMD and SRF) [4] and a deep learning-based approach that used both RGB and depth images as inputs (ED-RGBHHA) [16]. Note that, all the deep learning-based methods use the VGG16 as the main backbone for a fair comparison.

B. Results

IIT-AFF Dataset Table I summarizes results on the IIT-AFF dataset. The results clearly show that AffordanceNet significantly improves over the state of the art. In particular, AffordanceNet boosts the F_{β}^w score to 73.35, which is 3.7% improvement over the second best BB-CNN-CRF. It is worth noting that AffordanceNet achieves this result using an end-to-end architecture, and no further post processing step such as CRF is used. Our AffordanceNet also achieves the best results for all 9 affordance classes. We also found that for the dataset containing cluttered scenes such as IIT-AFF, the approaches that combine the object detectors with deep networks to predict the affordances (AffordanceNet, BB-

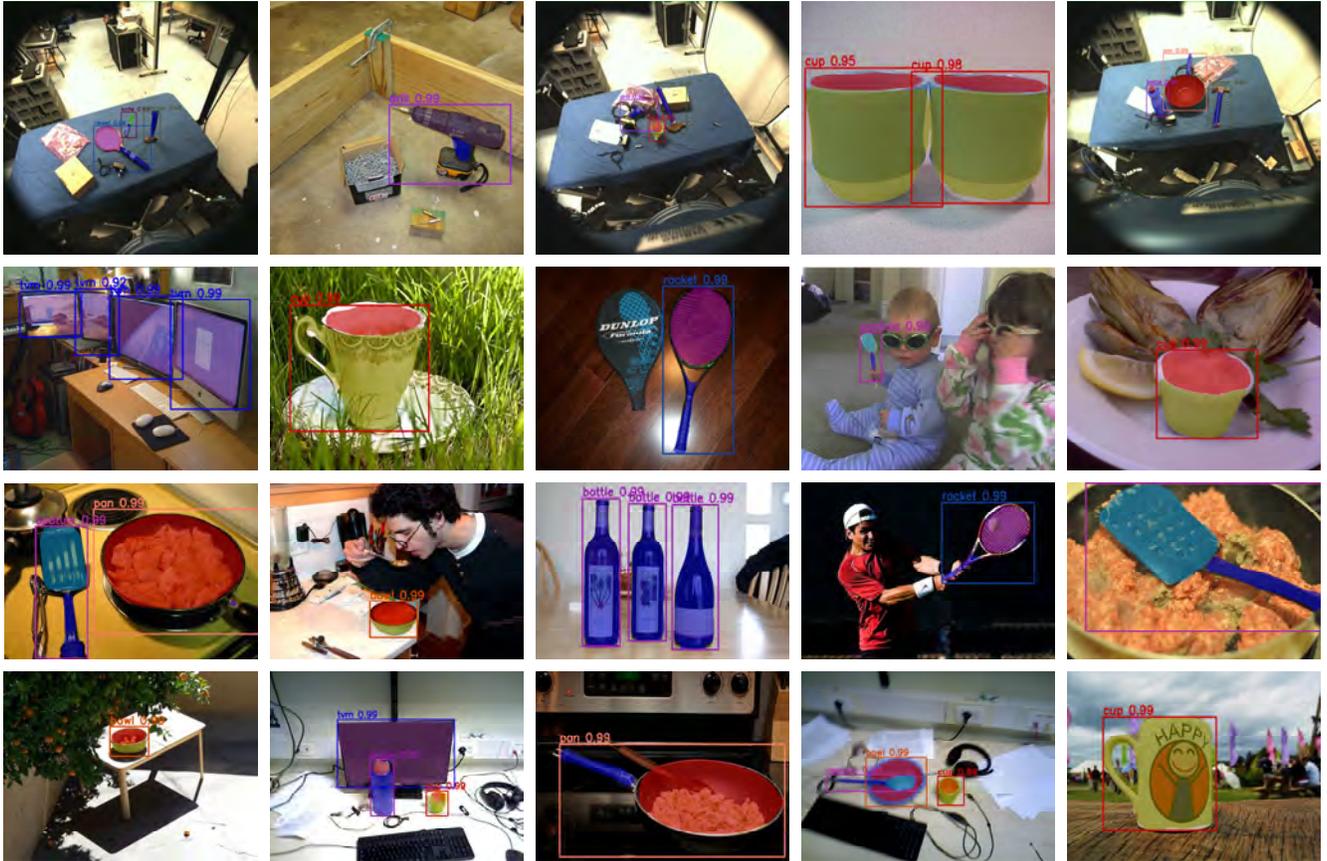


Fig. 4. Examples of affordance detection results by AffordanceNet on the IIT-AFF dataset.

CNN) significantly outperform over the methods that use deep networks alone (DeepLab, ED-RGB).

UMD Dataset Table II summarizes results on the UMD dataset. On the average, our AffordanceNet also achieves the highest results on this dataset, i.e., it outperforms the second best (ED-RGBD) 2.9%. It is worth noting that the UMD dataset only contains clutter-free scenes, therefore the improvement of AffordanceNet over compared methods is not as high as the one in the real-world IIT-AFF dataset. We recall that the AffordanceNet is trained using the RGB images only, while the second best (ED-RGBD) uses both RGB and the depth images. The Table II also clearly shows that the deep learning-based approaches such as AffordanceNet, DeepLab, ED-RGB significantly outperform the hand-designed geometric feature-based approaches (HMP, SRF).

To conclude, our AffordanceNet significantly improves over the state of the art, while it does not require any extra post processing or data augmentation step. From the robotic point of view, AffordanceNet can be used in many tasks since it provides all the object locations, object categories, and object affordances in an end-to-end manner. The running time of AffordanceNet is around $150ms$ per image on a Titan X GPU, making it is suitable for robotic applications. Our implementation is based on Caffe deep learning library [27]. The source code and trained models that allow reproducing the results in this paper will be released upon acceptance.

C. Effect of Affordance Map Size

In this section, we analyze the effect of the affordance map size. Follow the setup in Mask-RCNN, we use only one deconvolutional layer with parameters ($d = 1, s = 2, S_f = 4$) to create 14×14 affordance map from the 7×7 feature map (denoted as AffordanceNet14). Similarly, we change the parameters to ($d = 1, s = 4, S_f = 6$) to create the 28×28 affordance map (denoted as AffordanceNet28). Furthermore, we also setup networks which use two deconvolutional layers to create 56×56 affordance map (denoted as AffordanceNet56), and three deconvolutional layers to create 112×112 affordance map (denoted as AffordanceNet112). Finally, to check the effect of the convolutional layers, we also setup a network with 6 convolutional layers (together with ReLu), follow by a deconvolutional layer that up-sampling the 7×7 feature map to 14×14 (denoted as AffordanceNet14.6Conv).

Table III summarizes the average F_β^w score of the aforementioned networks on the IIT-AFF dataset. The results show that the affordance detection accuracy is gradually increasing when the bigger affordance map is used. In particular, the AffordanceNet14 gives very poor results since the map size of 14×14 is too small to represent multiclass affordances. The accuracy is significantly improved when we use the 28×28 affordance map. However, the improvement does not linearly increase with the affordance map size, it slows

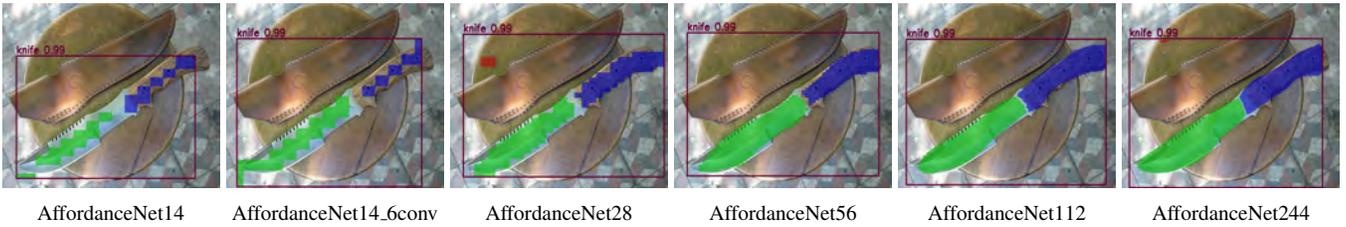


Fig. 5. Examples of predicted affordance masks using different mask sizes. The predicted mask is smoother and finer when a bigger mask size is used.

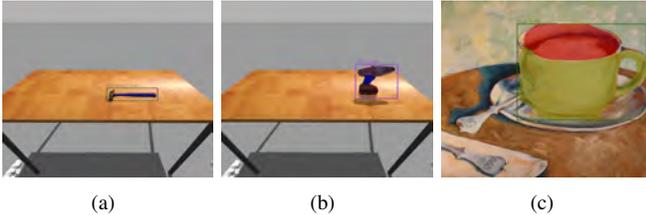


Fig. 6. Affordance detection in the wild. (a) and (b): We use AffordanceNet to detect the objects in Gazebo simulation. (c) AffordanceNet also performs well when the input is an artwork.

TABLE III
EFFECT OF MASK SIZE

	F_{β}^w
AffordanceNet14	57.71
AffordanceNet28	66.13
AffordanceNet56	71.54
AffordanceNet112	72.52
AffordanceNet14.6Conv	60.27
AffordanceNet	73.35

down when the bigger mask sizes are used. Note that using the big affordance map can improve the accuracy, but it also increases the number of parameters of the network. In our work, we choose the 244×244 map size for AffordanceNet since it both gives the good accuracy and can be trained with a Titan X GPU. We also found that using more convolutional layers (as in AffordanceNet14.6Conv) can also improve the accuracy, but it still requires to upsample the affordance map to high resolution in order to achieve good results. Fig 5 shows some example results when different affordance map sizes are used.

D. Affordance Detection in The Wild

The experimental results on the simple constrained environment UMD dataset and the real-world IIT-AFF dataset show that the AffordanceNet performs well on public research datasets. However, real-life images may be more challenging. In this study, we show some qualitative results to demonstrate that the AffordanceNet can generalize well in other testing environments. As illustrated in Fig 6, our AffordanceNet can successfully detect the objects and their affordances from artwork images or images from a simulated camera in Gazebo simulation [28]. Although this result is qualitative, it shows that AffordanceNet is applicable for wide ranges of applications, including in simulation environment which is crucial for developing robotic applications.



Fig. 7. WALK-MAN is performing a pouring task. The outputs of the object detection branch help the robot to recognize and localize the objects (i.e., bottle, pan) while the outputs of the affordance detection branch help the robot to perform the task (i.e., where on the bottle the robot should grasp and where on the pan the water should be poured).

E. Robotic Applications

Since the AffordanceNet can detect both the objects and their affordances at the speed of $150ms$ per image, it is quite suitable for robotic applications. To demonstrate that, we use the humanoid robot WALK-MAN [29] to perform different manipulation experiments. The robot is controlled in real-time using the XBotCore framework [30]. The whole-body motion planning is generated by OpenSoT library [31], while the AffordanceNet is used to provide visual information for the robot. Note that, from the 2D information outputted by AffordanceNet, we use the corresponding depth image to project it into 3D space, to be used in the real robot. Using this setup, the robot can perform different tasks such as grasping, pick-place, and pick-pouring. It is worth noting that all information produced by the AffordanceNet, i.e. the object locations, object labels, and object affordances are very useful for the tasks. For example, the robot knows where to grasp a *bottle* via the bottle's `grasp` affordance, and where to pour the water into a *pan* via the pan's `contain` affordance (see Fig. 7). Our experimental video can be found at the following link: <https://sites.google.com/site/affordancenetwork/>

V. CONCLUSION

We have proposed AffordanceNet, an end-to-end deep learning framework that can simultaneously detect the objects and their affordances. Different from state-of-the-art network architectures for instance segmentation, we proposed three components to address the problem of multiple affordance classes in affordance detection task: a sequence of deconvolutional layers, a robust resizing strategy, and a new loss function. We showed that these components are the key factors to achieve high affordance detection accuracy. The extensive experimental results show that our AffordanceNet not only achieves state-of-the-art results on public datasets, but can also be used in various robotic applications.

ACKNOWLEDGMENT

Thanh-Toan Do and Ian Reid are supported by the Australian Research Council through the Australian Centre for Robotic Vision (CE140100016). Ian Reid is also supported by an ARC Laureate Fellowship (FL130100102). Anh Nguyen is supported by the European Union Seventh Framework Programme (FP7-ICT-2013-10) under grant agreement no 611832 (WALK-MAN). The authors would like to thank Darwin G. Caldwell and Nikos G. Tsagarakis for the useful discussion.

REFERENCES

- [1] J. J. Gibson, *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin, 1979.
- [2] C. Castellini, T. Tommasi, N. Noceti, F. Odone, and B. Caputo, "Using object affordances to improve object recognition," *IEEE Transactions on Autonomous Mental Development*, 2011.
- [3] H. S. Koppula and A. Saxena, "Anticipating human activities using object affordances for reactive robotic response," *TPAMI*, 2016.
- [4] A. Myers, C. L. Teo, C. Fermüller, and Y. Aloimonos, "Affordance Detection of Tool Parts from Geometric Features," in *ICRA*, 2015.
- [5] A. Nguyen, D. Kanoulas, L. Muratore, D. G. Caldwell, and N. G. Tsagarakis, "Translating videos to commands for robotic manipulation with deep recurrent neural networks," in *ICRA*, 2018.
- [6] A. Nguyen, D. Kanoulas, D. G. Caldwell, and N. G. Tsagarakis, "Object-based affordances detection with convolutional neural networks and dense conditional random fields," in *IROS*, 2017.
- [7] B. Hariharan, P. A. Arbeláez, R. B. Girshick, and J. Malik, "Simultaneous detection and segmentation," in *ECCV*, 2014.
- [8] K. He, G. Gkioxari, P. Dollár, and R. B. Girshick, "Mask R-CNN," in *ICCV*, 2017.
- [9] G. Lin, A. Milan, C. Shen, and I. Reid, "RefineNet: Multi-path refinement networks for high-resolution semantic segmentation," in *CVPR*, 2017.
- [10] A. Roy and S. Todorovic, "A multi-scale cnn for affordance segmentation in rgb images," in *ECCV*, 2016.
- [11] T. Pham, T.-T. Do, N. Sünderhauf, and I. Reid, "Scenecut: Joint geometric and object segmentation for indoor scenes," in *ICRA*, 2018.
- [12] M. Schoeler and F. Wörgötter, "Bootstrapping the semantics of tools: Affordance analysis of real world objects on a per-part basis," *IEEE Transactions on Cognitive and Developmental Systems*, 2016.
- [13] H. O. Song, M. Fritz, D. Goehring, and T. Darrell, "Learning to detect visual grasp affordance," *IEEE Transactions on Automation Science and Engineering (TASE)*, 2015.
- [14] H. Kjellström, J. Romero, and D. Kragic, "Visual object-action recognition: Inferring object affordances from human demonstration," *CVIU*, 2011.
- [15] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," *IJRR*, 2015.
- [16] A. Nguyen, D. Kanoulas, D. G. Caldwell, and N. G. Tsagarakis, "Detecting Object Affordances with Convolutional Neural Networks," in *IROS*, 2016.
- [17] J. Sawatzky, A. Srikantha, and J. Gall, "Weakly Supervised Affordance Detection," in *CVPR*, 2017.
- [18] Y. Li, H. Qi, J. Dai, X. Ji, and Y. Wei, "Fully convolutional instance-aware semantic segmentation," in *CVPR*, 2017.
- [19] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks," in *NIPS*, 2015.
- [20] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," *CoRR*, vol. abs/1409.1556, 2014.
- [21] R. Girshick, J. Donahue, T. Darrell, U. C. Berkeley, and J. Malik, "R-CNN: Rich feature hierarchies for accurate object detection and semantic segmentation," in *CVPR*, 2014.
- [22] R. B. Girshick, "Fast R-CNN," in *ICCV*, 2015.
- [23] R. Girshick, F. Iandola, T. Darrell, and J. Malik, "Deformable part models are convolutional neural networks," in *CVPR*, 2015.
- [24] L. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs," *TPAMI*, 2017.
- [25] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "Imagenet large scale visual recognition challenge," *IJCV*, pp. 211–252, 2015.
- [26] R. Margolin, L. Zelnik-Manor, and A. Tal, "How to Evaluate Foreground Maps," in *CVPR*, 2014.
- [27] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. B. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional architecture for fast feature embedding," in *ACM MM*, 2014.
- [28] E. Mingo Hoffman, S. Traversaro, A. Rocchi, M. Ferrati, A. Settini, F. Romano, L. Natale, A. Bicchi, F. Nori, and N. G. Tsagarakis, *Yarp Based Plugins for Gazebo Simulator*. Springer, 2014, pp. 333–346.
- [29] N. G. Tsagarakis, D. G. Caldwell, F. Negrello, W. Choi, L. Baccelliere, V. Loc, J. Noorden, L. Muratore, A. Margan, A. Cardellino, L. Natale, E. Mingo Hoffman, H. Dallali, N. Kashiri, J. Malzahn, J. Lee, P. Kryczka, D. Kanoulas, M. Garabini, M. Catalano, M. Ferrati, V. Varricchio, L. Pallottino, C. Pavan, A. Bicchi, A. Settini, A. Rocchi, and A. Ajoudani, "WALK-MAN: A High Performance Humanoid Platform for Realistic Environments," *Journal of Field Robotics*, 2016.
- [30] L. Muratore, A. Laurenzi, E. Mingo Hoffman, A. Rocchi, D. G. Caldwell, and N. G. Tsagarakis, "Xbotcore: A real-time cross-robot software platform," in *IEEE International Conference on Robotic Computing*, 2017.
- [31] A. Rocchi, E. Mingo Hoffman, D. Caldwell, and N. Tsagarakis, "OpenSoT: A Whole-Body Control Library for the Compliant Humanoid Robot COMAN," in *ICRA*, 2015.