



King's Research Portal

DOI:

[10.1109/LRA.2022.3191231](https://doi.org/10.1109/LRA.2022.3191231)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Jiang, J., Cao, G., Do, T. T., & Luo, S. (2022). A4T: Hierarchical Affordance Detection for Transparent Objects Depth Reconstruction and Manipulation. *IEEE Robotics and Automation Letters*, 7(4), 9826-9833.
<https://doi.org/10.1109/LRA.2022.3191231>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

A4T: Hierarchical Affordance Detection for Transparent Objects Depth Reconstruction and Manipulation

Jiaqi Jiang^{1,2}, Guanqun Cao¹, Thanh-Toan Do³ and Shan Luo^{1,2}

Abstract—Transparent objects are widely used in our daily lives and therefore robots need to be able to handle them. However, transparent objects suffer from light reflection and refraction, which makes it challenging to obtain the accurate depth maps required to perform handling tasks. In this paper, we propose a novel affordance-based framework for depth reconstruction and manipulation of transparent objects, named *A4T*. A hierarchical AffordanceNet is first used to detect the transparent objects and their associated affordances that encode the relative positions of an object’s different parts. Then, given the predicted affordance map, a multi-step depth reconstruction method is used to progressively reconstruct the depth maps of transparent objects. Finally, the reconstructed depth maps are employed for the affordance-based manipulation of transparent objects. To evaluate our proposed method, we construct a real-world dataset *TRANS-AFF* with affordances and depth maps of transparent objects, which is the first of its kind. Extensive experiments show that our proposed methods can predict accurate affordance maps, and significantly improve the depth reconstruction of transparent objects compared to the state-of-the-art method, with the Root Mean Squared Error in meters significantly decreased from 0.097 to 0.042. Furthermore, we demonstrate the effectiveness of our proposed method with a series of robotic manipulation experiments on transparent objects. See supplementary video and results at <https://sites.google.com/view/affordance4trans>.

Index Terms—Robotics and Automation in Life Sciences, Computer Vision for Automation.

I. INTRODUCTION

TRANSSPARENT objects such as plastic bottles, glass dishes, and windows are widely seen in our daily lives. Many of the containers in biomedical and chemical laboratories are also transparent, e.g., Petri dishes, glass flasks, and vials. Humans can expertly identify these transparent objects and interact with them at ease, e.g., picking up a glass cup placed on a table and filling it with water. In recent years, domestic robots have been developed and deployed to

Manuscript received March 1, 2022; Revised June 4, 2022; Accepted June 23, 2022. This paper was recommended for publication by Editor Hyungpil Moon upon evaluation of the Associate Editor and Reviewers’ comments. This work was partly supported by a University of Liverpool and China Scholarship Council Award, and the EPSRC project “ViTac: Visual-Tactile Synergy for Handling Flexible Materials” (EP/T033517/1).

¹smARTLab, Department of Computer Science, University of Liverpool, Liverpool L69 3BX, United Kingdom. Emails: {jiaqi.jiang, g.cao, shan.luo}@liverpool.ac.uk

²Jiaqi Jiang and S. Luo are also with Department of Engineering, King’s College London, London WC2R 2LS, United Kingdom. E-mails: {jiaqi.1.jiang, shan.luo}@kcl.ac.uk

³T.-T. Do is with Department of Data Science and AI, Monash University, Clayton, VIC 3800, Australia. E-mail: toan.do@monash.edu.

Digital Object Identifier (DOI): see top of this page.




Fig. 1: Our goal is to obtain accurate depth maps of transparent objects by leveraging affordance detection, so as to facilitate the manipulation of them, i.e., stacking two plastic cups in the above example. **Top left:** A plastic cup is in the robot’s gripper and another is placed on the table, viewed by an RGB-D camera on its side. **Bottom left:** A depth map of the cup on the table is obtained from the camera, and its affordance map is predicted, in which the red and blue colours represent the affordance regions with deep cavities to hold liquid (“contain”) and that can be held (“wrap-grasp”), respectively. **Bottom right:** The affordance map is used to improve the depth map and predict the center of the gripper for stacking. **Top right:** With the improved depth maps and the predicted gripping center, the robot can stack the plastic cups successfully.

assist daily living tasks, and mobile robot chemists have also been created for materials discovery in laboratories [1]. As transparent objects are widely used in these environments, it is essential that domestic and laboratory robots are able to perceive the transparent objects that are in their surrounding environments and interact with transparent objects safely and dexterously.

Many vision-based techniques have been proposed for perceiving poses and shapes of objects in recent decades [2], [3]. However, most of the objects in the previous research have been opaque and the perception of transparent objects remains a challenging problem. Compared to opaque objects, the properties of transparent objects such as reflection, refraction, and the lack of salient features, pose additional challenges to the design of vision algorithms. Their transparent materials violate the Lambertian assumption that optical 3D sensors (e.g., LiDAR and RGB-D cameras) are based on. Hence, most of the depth data from transparent objects are invalid or contain unpredictable noise, which leads to missing or noisy depth approximations of their surfaces. Due to these challenges,

most of the current manipulation methods that highly rely on accurate depth information from cameras cannot be applied to the manipulation of transparent objects directly.

The existing works divide the manipulation of transparent objects into two steps: depth reconstruction and manipulation planning. In [4], the depth maps of transparent objects were first reconstructed and then objects were grasped with predicted action possibilities of object regions. However, the depth reconstruction in [4] requires reliable initial depth maps of contact edges, i.e., points of contact between transparent objects and tables. This requirement cannot be met for some areas such as regions with deep cavities to contain content and as a result depth maps for such regions cannot be determined.

In this paper, we propose a novel approach named *A4T*, i.e., Affordance for Transparent object depth reconstruction and manipulation. It couples depth reconstruction and manipulation planning for the manipulation of transparent objects via an affordance map that represents the object functionality of each pixel. The affordance map can improve the depth reconstruction and also output the functionalities of object regions for manipulation. As shown in Fig. 1, a pixel-wise affordance map of object functionalities is first predicted with a hierarchical AffordanceNet that fuses the high-level affordance classification with the low-level fine-grained affordance predictions to avoid the noisy fragments in predicted affordance maps. With the relative positions of affordance regions encoded, the pixel-wise affordance maps are then used to guide the multi-step depth reconstruction of transparent objects. Furthermore, the reconstructed depth maps and affordance maps are used to improve the affordance-based manipulation.

To evaluate our proposed method, we collected a real-world dataset of transparent objects *TRANS-AFF*, annotated with their affordances and depth maps for opaque objects of the same geometry as the transparent objects to provide the ground truth, which is the very first of its kind. The extensive experimental results demonstrate that our proposed *A4T* can remove the noisy fragments in affordance maps and reconstruct depth maps of transparent objects accurately. Compared to the state-of-the-art method [4], the Root Mean Squared Error in meters is reduced from 0.097 to 0.042. We also conduct the real affordance-based manipulation experiments with a robotic arm. The results show that our proposed depth reconstruction method can significantly enhance the success rates of pouring into transparent objects from 15% to 80%.

The rest of paper is structured as follows: Section II reviews related works and Section III introduces our *A4T* framework in detail; Section IV details the *TRANS-AFF* dataset; Section V analyses the experimental results; Finally, Section VI summarises the paper and discusses the work.

II. RELATED WORKS

A. Depth Reconstruction for Transparent Objects Grasping

Depth reconstruction is an important step towards reliable manipulation of transparent objects as it mitigates the sensor failures in depth images caused by reflection and refraction of transparent objects. Most of the current depth reconstruction methods for transparent objects either rely on a specific capture

process or prior knowledge. Han et al. [5] developed a special fixed view capture system, in which parts of the object were immersed in a liquid to make light only reflects for once in objects. In [6] and [7], the shapes of transparent objects were recovered with known backgrounds and 3D models, respectively.

There are some other approaches reconstructing object shapes without the prior knowledge about geometries or 3D model. In [8], an approach was proposed that first matches pixels from Time-of-Flight images and then reconstructs surfaces with triangulating methods. Neural radiance fields [9] were used in [10] to infer the geometries of transparent objects. Sajjan et al. [4] used the global optimisation algorithm proposed in [11] to reconstruct the missing or noisy depth regions of transparent objects. In [12], Zhu et al. used a local implicit neural representation and an iterative depth refinement model to complete the depth information of transparent objects. However, all the above methods only focus on depth reconstruction, and have to be combined with other manipulation planning methods [2], [13] to manipulate transparent objects. In our work, we propose a novel approach that couples depth reconstruction and manipulation planning for manipulation of transparent objects via affordance maps that can improve the depth reconstruction and also output the functionalities of object regions for manipulation.

B. Affordance Detection

The study of affordance learning has attracted increasing attention from researchers in recent years. Currently, there are two main research directions for understanding the affordances of objects, i.e., behavior-grounded affordance learning and pixel-level affordance learning. The behavior-grounded affordance learning [14], [15], [16] learns object affordance via observing the effects of robot's actions performed on the objects. The tool affordance was defined in [14] with three tool functional features that are annotated by hand. In [15], 2D geometrical visual features were used to represent the functional features for tools. Parallel Self-Organising Maps (SOMs) were used in [16] to learn the tool affordance based on their 3D geometry.

The pixel-level affordance learning takes a group of pixels that share the same object functionality as an affordance. In [17], the affordances of tool parts were detected from local shapes and geometries. However, those hand-designed features can only capture the information of object properties. In recent years, deep learning approaches have been used to learn affordances and understand the relationship between different parts of one object [13], [18], [19], [20]. An encoder-decoder architecture and Fully Convolutional Networks (FCN) [21] were used to efficiently obtain dense affordances predictions in [19] and [13], respectively. Different from previous semantic segmentation methods, in [18], [20], objects were detected (i.e., the object location and the object label) simultaneously with their associated affordances. Instead of only using affordances to enhance the learning of visual representations [20], our hierarchical AffordanceNet fuses the high-level predicted affordance classes and the low-level fine-grained affordance maps to avoid noisy fragments in affordance map.

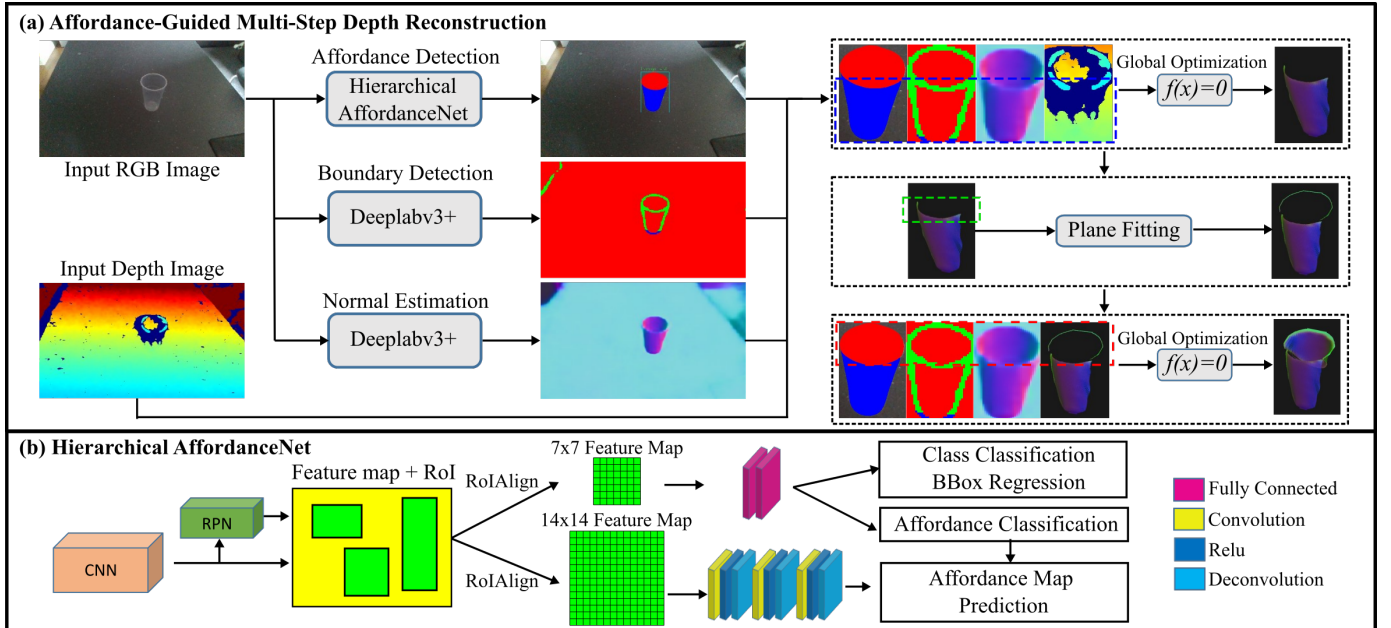


Fig. 2: (a): *From Left to Right*. Given an RGB-D image of a scene with transparent objects, **A4T** uses three networks to infer 1) affordance maps of transparent objects, 2) occlusion boundaries and contact edges, and 3) surface normals. Then based on the affordance map, the transparent object is progressively reconstructed with a global optimisation method and a plane fitting method. The colourful dash rectangles in the top right corner represent the input information for each reconstruction step. (b): *From Left to Right*. A deep Convolutional Neural Network (CNN) backbone is used to extract the features of RGB images. The Region Proposal Network (RPN) shares weights with the CNN backbone and outputs RoIs. For each RoI, two RoIAlign layers extract and pool its features to feature maps of a fixed size. Two fully connected layers are used for object classification, object location regression and affordance classification. Three convolutional-deconvolutional layers are used to obtain affordance maps that are fused with the affordance classification scores. Finally, a softmax layer is applied to output a multi-class affordance mask.

III. METHODOLOGY

In this work, we propose an affordance-based depth reconstruction framework that facilitates the robotic manipulation of transparent objects, as shown in Fig. 2. The transparent object in the image is first detected, with its affordance classes and affordance map, using a novel hierarchical AffordanceNet. Then an affordance-guided multi-step depth reconstruction method is used to progressively reconstruct the depth maps of the transparent object. Finally, using the reconstructed depth the robot executes the affordance-based manipulation tasks, i.e., picking, pouring, and stacking.

A. Hierarchical AffordanceNet

Our hierarchical AffordanceNet extends our previous AffordanceNet [18] by introducing an affordance classification module, as shown in Fig. 2-(b). By fusing the high-level affordance classes with the low-level fine-grained affordance map, our hierarchical AffordanceNet can make the predicted affordance map consistent with the predicted affordance classes. The RGB features are first extracted, with the ResNet50 [22] network as the backbone. A Region Proposal Network (RPN) that shares the same weights with the convolutional backbone is then used to generate the Regions of Interest (RoIs). They are followed by two kinds of RoIAlign [23] layers that extract and pool the corresponding features of candidate Bounding Boxes (BBboxes) into 7×7 feature maps (for the BBox regression and affordance classification) and 14×14 feature maps (for the affordance map prediction).

Affordance classification. The affordance classification is recognised as a multi-label classification problem, and parallel

to the object detection. We feed the 7×7 feature maps to two fully connected layers, each with 1024 neurons, followed by N binary classifiers, where N is the number of affordance classes. The N binary classifiers output a N -dimensional classification score A_c that represents the possibility of each affordance within the proposal. To train the affordance classification, a multinomial cross entropy loss L_{AFF-C} is used and defined as follows:

$$L_{AFF-C} = -\frac{1}{N} \sum_{i=0}^N (y_i * \log(A_c(i)) + (1 - y_i) * \log(1 - A_c(i))) \quad (1)$$

where y_i denotes the label of the i^{th} affordance class.

Affordance map prediction. As for the pixel-wise affordance map prediction, the 14×14 feature maps are upsampled to 112×112 with three convolutional-deconvolutional layers. Then a convolutional layer is used to generate the output affordance map A_m with $N + 1$ channels, where the plus one represents the additional channel for background. To train the affordance map, a multinomial cross entropy loss L_{AFF-M} is used and computed as follows:

$$L_{AFF-M} = -\frac{1}{K} \sum_{i \in RoI} \log(m_{s_i}^i) \quad (2)$$

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

Fusion of affordance classification. Different from [20] that only uses affordance classification to guide feature learning during training, we fuse the affordance classification scores

A_c and the predicted affordance map A_m to infer the final pixel-wise affordance map A_m^F :

$$A_m^F = [1, A_c] \otimes A_m \quad (3)$$

where A_m^F with $N + 1$ channels and \otimes is the element-wise product, also known as Hadamard product. Note that the value 1 in Equation 3 is used to keep the background channel of the pixel-wise affordance map. After obtaining the final affordance map A_m^F , a high-resolution affordance mask is then generated by a softmax layer.

B. Affordance-guided Depth Reconstruction

In this subsection, we propose a multi-step depth reconstruction method that embeds the relative positions of different affordance regions. Different from [4] that reconstructs the entire image at one step, we reconstruct every individual object with a multi-step method that reconstructs different affordance regions one by one. In this way, the reconstructed depth of the current affordance region can be recognised as a clue for the reconstruction of the next affordance region, so that even the areas enclosed by depth discontinued boundaries can be reconstructed. Firstly, the RGB-D image, the predictions of surface normals and occlusion boundaries are cropped for each detected instance. Then the depth map of one transparent object is progressively reconstructed with a global optimisation method and a RANSAC-based plane fitting method.

Cropping of Transparent Objects. Following [4], first two DeepLabv3+ networks [24] are used to predict surface normals and boundaries of transparent objects. The boundaries include both occlusion boundaries and contact edges where depths are not continued and objects are contacted with tables, respectively. To reconstruct every individual object instead of the entire image in [4], the depth image, as well as its corresponding surface normals image and occlusion boundary image are cropped. Then based on the predicted affordance map from Sec. III-A, the unreliable depth measurements of transparent objects are removed from the depth image.

Depth reconstruction. For each detected transparent object, we fill in their removed depth map with two methods, i.e., the global optimisation algorithm proposed by Zhang and Funkhouser [11] and a RANSAC-based plane fitting method [25]. Specifically, the global optimisation algorithm completes the depth information of the transparent objects by minimizing the integrated error E :

$$E = \lambda_D E_D + \lambda_S E_S + \lambda_N E_N B \quad (4)$$

where E_D measures the distance between the estimated depth and the observed raw depth, E_S measures the difference between the depths of neighboring pixels, and E_N measures the consistency between the estimated depth and the predicted surface normal, respectively; λ_D , λ_S , λ_N are chosen according to [4], [11], [26]. λ_D is set to 1000 to encourage the reconstructed depth as close as possible to the reliable raw depth of opaque objects; λ_S is set to 0.001 to encourage adjacent pixels to have the same depths; λ_N is set to 1 as a reference. B down-weights the normal terms based on the predicted probability that a pixel is on an occlusion boundary.

The RANSAC-based plane fitting method is used to complete the edge depth with part of known points on it. Firstly, it randomly selects three points and uses Singular Value Decomposition (SVD) to initialize the plane parameters [27]. Then it detects all points belonging to the calculated plane, according to a given threshold. Afterwards, it repeats these procedures I ($= 500$ in our case) times. In each time, the plane parameters will be updated if the current result is better than the last saved one. If the objects vertically are placed on a flat table, we constrain the fitted plane parallel to the table to obtain a more robust fitting result.

Firstly, we use the global optimisation method to complete the affordance regions with contact edges on their boundaries. Then the following reconstruction processes will be determined by the connected edges on first reconstructed regions (i.e., points of contact between two regions of different affordances). If the connected edge is depth discontinued, such as the connected edge between “wrap-grasp” and “contain” (the blue region and the red region on cropped affordance map in Fig. 2, respectively), the RANSAC-based plane fitting method would be used to complete the boundary depth of their neighbor affordance regions. After that, the global optimisation method is used to complete their neighbor affordance regions, with the optimised depth on boundaries taken as the reliable depth. Otherwise, if the connected edge is depth continued, such as the connected edge between “wrap-grasp” and “support” (the bottom of containers that can be supported by a table), the global optimisation would be used again to reconstruct the depth in their neighbor affordance regions.

C. Affordance-based Manipulation

To demonstrate that our method has the ability to manipulate the transparent objects, we take three manipulation tasks, picking, pouring and stacking as examples, as shown in Fig. 3.




Fig. 3: An example of the generated proposals for different tasks. Red, green, and blue arrows represent the x-axis, y-axis, and z-axis of the end-effector.

Pouring. In the pouring task, the “contain” affordance is used to generate the pouring proposal $P_{pour} = [R, T] \in SE(3)$. Specifically, the translation vector T is calculated by averaging the point clouds on “contain” region’s boundary. As for the rotation vector R , the z-axis of the end-effector is perpendicular to the plane where the “contain” boundary lies on and is in the same vertical plane with y-axis. If the robotic arm pours with a container, an offset of container’s length will be added to the y-axis.

Picking. In the picking task, the “wrap-grasp” affordance is used to generate the picking proposal $P_{pick} = [R, T] \in SE(3)$. Specifically, the translation vector T is calculated by averaging the point clouds of “wrap-grasp” region. As for the rotation vector R , the z-axis of the end-effector is parallel




Fig. 4: Sample objects from the **TRANS-AFF** Dataset. **Top row**: RGB images. **Middle row**: Depth images. **Bottom row**: Pixel-wise affordances label. Red, blue and green colour represent “contain”, “wrap-grasp” and “support” respectively.

to the predicted surface normal of the center pixel of “wrap-grasp” region and is in the same vertical plane with y-axis.

Stacking. In the stacking task, the “contain” and “support” affordances are used to generate the stacking proposal $P_s = [R, T] \in SE(3)$, where the stacking proposals for objects containing “support” affordance and “contain” affordance are generated with the same ways to generating the picking proposal and pouring proposal, respectively.

IV. DATASET COLLECTION

A. TRANS-AFF Dataset

To evaluate our proposed methods, we collect a new RGB-D Affordance dataset for transparent objects named as **TRANS-AFF**, with some samples shown in Fig. 4. We collect the dataset with a RealSense D435i camera and a RealSense D415 camera that have a resolution of 1280×720 pixels and a resolution of 1920×1080 pixels, respectively. To provide accurate ground-truth depth maps for the transparent objects, we collected a “twin” RGB-D pair for every scene that contains transparent objects. In the “twin” setup, the transparent objects in the original image are replaced with an identical spray-painted instance that can reflect light evenly and provide accurate depth information. In total, there are 1,346 pairs of RGB and depth images for 8 graspable containers as shown in Fig. 5. Compared to other affordance datasets in the literature [17], [19], our **TRANS-AFF** dataset is the first one that includes transparent objects and provides accurate depth maps. Furthermore, to our best knowledge, it is also the first dataset designed for both depth reconstruction and affordance detection.

In this dataset, we define three surfaces effective affordances by the way it comes in contact with the objects it affects. For example, a water cup standing upright on the table has two affordance parts, i.e., the inner surface and the outer surface. The inner surface of a cup has the effective affordance “contain”, as it comes in contact with the liquid that is contained. The outer surface of a cup has the affordance “wrap-grasp” as it can be held with the hand and the palm. In addition to the two above, we set the bottom of the cup as “support”, as it can make the cup stand upright on the table. Those three affordances are summarised in Table I.



Fig. 5: Objects in the **TRANS-AFF** Dataset. **First row** (from left to right): disposable cup, highball cup, rectangular cup, and square jar; **Second row** (from left to right): tumble cup, bowl, round jar, and measurement cup.

TABLE I: Description of the three affordance labels (see Fig. 4 for examples).

Affordance	Description
contain	With deep cavities to hold liquid or other content.
wrap-grasp	Regions that can be held with the hand.
support	A flat part that can be supported vertically on a table.

We split all the data in the dataset instead of the objects by a ratio of 7:2:1 for training, validation and testing. In this way, all the transparent objects used in this work are known a priori, which is consistent with the fact that the models of the objects used in laboratories and factories are available from their manufacturers. To reduce the influence of different shadows and caustics between opaque and transparent twins, only transparent objects are used for training the hierarchical AffordanceNet to focus on the features of transparent objects.

V. EXPERIMENT RESULTS

In this section, we conduct a series of experiments to evaluate our affordance detection for transparent object depth reconstruction and manipulation. The goal of the experiments are three-fold: 1) To evaluate the effectiveness of the proposed hierarchical AffordanceNet in our **TRANS-AFF** dataset; 2) To investigate how affordance detection can improve the performance of transparent object’s depth reconstruction; 3) To investigate how the reconstructed depth information can improve the success rate of affordance-based manipulation.

A. Multi-Affordance Prediction

To evaluate the affordance detection results, the weighted F-measures metric F_{β}^{ω} is used and can be computed as follows:

$$F_{\beta}^{\omega} = (1 + \beta^2) \frac{\text{Pr}^{\omega} \cdot \text{Rc}^{\omega}}{\beta^2 \cdot \text{Pr}^{\omega} + \text{Rc}^{\omega}} \quad (5)$$

where $\beta = 1$, Pr^{ω} and Rc^{ω} are the weighted precision and recall values, respectively. All the networks are trained for 20 epochs using a fixed learning rate of 0.001 with the SGD optimizer.

TABLE II: The ablation study of the hierarchical AffordanceNet

AFF-C	AFF-F	contain	support	grasp	Average
		83.80	85.08	95.26	88.04
✓		83.95	83.39	94.75	87.36
✓	✓	87.94	83.48	94.92	88.78

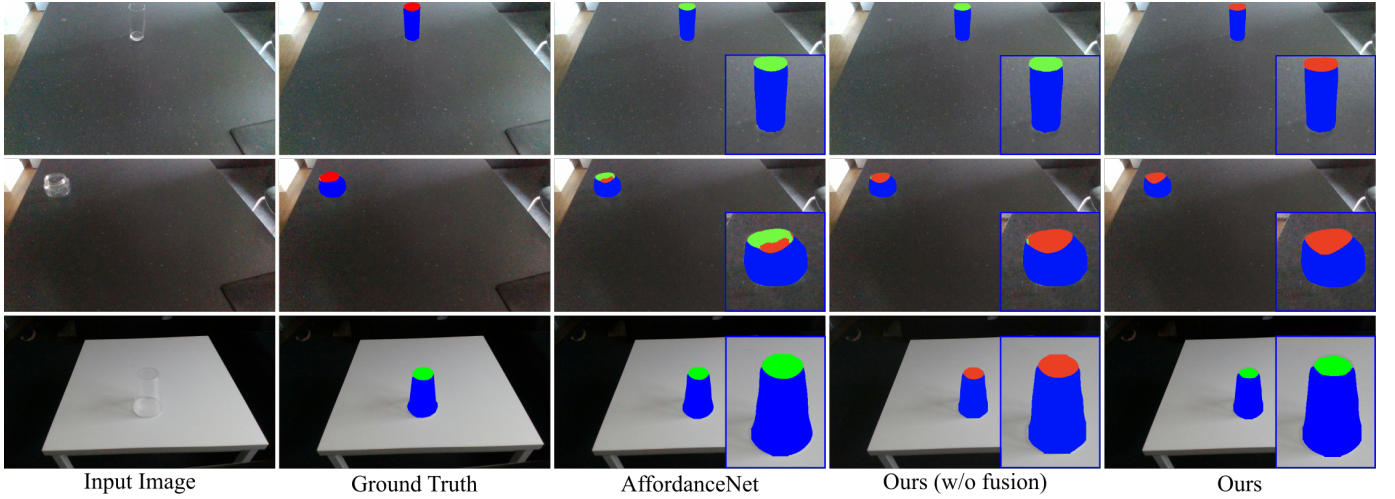


Fig. 6: Comparisons of affordance detection. The first two rows show that our method outperform the original AffordanceNet [18] and can avoid the noisy fragments in affordance maps. The third row shows that our hierarchical AffordanceNet may lose some fine-grained information as it involves the high-level affordance classification.

Table II summarizes the average F_{β}^{ω} score of the aforementioned networks when testing with the images that have a resolution of 1280×720 pixels in TRANS-AFF dataset. The “AFF-C” and “AFF-F” in Tab. II represent the affordance classification module and the fusion of affordance classification, respectively. The results show that our hierarchical model can significantly improve the performance on affordance “contain,” and has comparable performance to AffordanceNet.

We also show the qualitative results of affordance detection in Fig. 6. The “AFF-C” can enhance the learning of affordance representation but cannot fully avoid noisy fragments, as the second row of Fig. 6. Via fusing the high-level affordance classes with the low-level fine-grained affordance map, our hierarchical AffordanceNet can avoid the wrong affordance class and the noisy fragments as the first row and the second row of Fig. 6, respectively. However, there is a trade-off between affordance classification and pixel-wise affordance segmentation. The affordance classification module sometimes influences the fine-grained affordance segmentation, as the third row of Fig. 6, which limits the overall performance.

B. Depth Reconstruction

For the depth estimation, we use the standard metrics used in previous studies [28]: Root Mean Squared Error in meters (RMSE), the median error relative to the depth (Rel), Mean Absolute Error (MAE) and percentage of pixels with predicted depths falling within an interval ($[\delta = \max(\text{predicted}/\text{true}, \text{true}/\text{predicted})]$, where δ is 1.05, 1.10 or 1.25) [4].

We compare our proposed algorithm with the state-of-the-art algorithm ClearGrasp [4] and DeepCompletion [11] on the images within our TRANS-AFF dataset. The depth reconstruction results on different affordance regions are shown in Table III. It is noticed that our method has a similar performance on “wrap-grasp” regions and “support” regions to ClearGrasp method. However, there is a significant improvement on “contain” regions with the RMSE in meter decreased from 0.182 to 0.046, which contributes to the notable RMSE decrease for all regions from 0.097 to 0.046.

The reason why ClearGrasp method performs worst in “contain” regions is because the global optimisation method cannot determine the depth of the regions that are enclosed by occlusion boundaries. However, our algorithm can achieve a stable depth reconstruction on those challenging areas as shown in the fourth column of Fig. 7. It is also noticed that our method shows the worst accuracy on “support” regions compared to “contain” and “wrap-grasp,” this is because the DeepLabv3+ network incorrectly predicts the surface normals of some objects in the “support” regions as the surface normals of the “contain” regions.

As shown in Fig. 8, we also compare the reconstructed point clouds of the first picture of Fig. 6 when different affordance networks are used. Our method shows slightly better performance compared to using the original AffordanceNet [18], as the mispredicted affordance map leads to the wrong reconstruction of the “contain” region. More importantly, our hierarchical AffordanceNet can predict the affordance map without noisy fragments as the second row of Fig. 6, which makes the planning of our multi-step depth reconstruction easier. Objects other than graspable containers do not have “contain” affordance regions as they do not have cavities to contain content. This makes the plane fitting method in our multi-step reconstruction unneeded and hence our multi-step reconstruction will degenerate into ClearGrasp [4]. The objects such as bath bomb molds without “contain” affordance regions were tested in ClearGrasp [4] and could be reconstructed well with the vanilla global optimisation method in [11].

C. Real-world Manipulation

For the manipulation tasks, we use the average success rate $= \frac{\# \text{successful attempts}}{\# \text{total attempts}}$ as the evaluation metric. A picking attempt will be recognised as a successful case if the object is grasped and lifted up for 5 seconds. A pouring attempt will be recognised as a successful case if the object in the robotic hand or the content in the grasped container falls into the container. A stacking attempt will be recognised as a successful case if the object in robotic hand is stacked with the same object on the table.




Fig. 7: Samples of baseline comparison. We view those point clouds from different orientation to clearly compare the reconstruction of the “contain” regions. As shown in this figure, our method can achieve stable reconstruction of transparent objects, especially the inner surfaces of cups.

TABLE III: Baseline comparisons on different evaluation regions

	Evaluation Region	RMSE↓	Rel↓	MAE↓	$\delta_{1.05}$ ↑	$\delta_{1.10}$ ↑	$\delta_{1.25}$ ↑
DeepCompletion[11]	All	0.105	0.162	0.080	24.81	38.35	79.24
ClearGrasp[4]	All	0.097	0.139	0.072	53.51	64.44	79.46
Ours	All	0.042	0.064	0.033	63.33	86.47	95.78
DeepCompletion[11]	contain	0.106	0.185	0.096	15.40	24.83	70.67
ClearGrasp[4]	contain	0.182	0.329	0.172	14.20	19.70	41.77
Ours	contain	0.046	0.056	0.032	63.50	85.16	93.66
DeepCompletion[11]	wrap-grasp	0.089	0.141	0.069	28.31	43.03	84.89
ClearGrasp[4]	wrap-grasp	0.036	0.049	0.025	72.46	86.28	96.43
Ours	wrap-grasp	0.034	0.059	0.030	66.39	88.98	96.96
DeepCompletion[11]	support	0.139	0.192	0.090	35.97	49.81	70.60
ClearGrasp[4]	support	0.066	0.090	0.044	60.07	73.26	87.44
Ours	support	0.060	0.100	0.048	45.86	78.19	91.91



Fig. 8: Comparison using different AffordanceNet. Using the affordance prediction from AffordanceNet [18] results the failed reconstruction of the “contain” region.

TABLE IV: Comparison on success rate of different manipulation tasks.

Depth Reconstruction	picking	pouring	stacking
RealSense	20.0%	0.0%	25.0%
ClearGrasp	82.5%	15.0%	57.5%
Ours	82.5%	80.0%	87.5%

We conduct the affordance-based manipulation experiments with a robotic system that consists of a 6-DoF robotic manipulator UR5 from Universal Robots, a Robotiq-85 parallel two-finger gripper, and an Intel RealSense D415 camera. We use an ArUco Marker to achieve hand-eye calibration. Every manipulation task is tested with 40 attempts, and the results are summarised in Table IV. “RealSense” represents using the original depth information obtained from the RealSense camera to generate manipulation proposals. We set all objects vertically on the table and add a protection height to avoid the collision between the robotic arm and the table when manipulating the transparent objects.

The experimental results show that the robot arm can hardly manipulate the transparent objects using the original depth obtained from the RealSense camera. As for the picking task,

both our method and ClearGrasp method can significantly improve the success rate. The failure cases are mainly caused by inaccurate depth reconstruction, when no contact edges are detected or surface normals are predicted incorrectly. Nevertheless, constrained by the depth reconstruction performance, the success rates of the pouring task and the stacking task are not improved impressively with ClearGrasp method. Compared to ClearGrasp method, our approach with the ability to reconstruct “contain” regions accurately, can improve the success rates of the stacking task from 57.5% to 87.5%, and the pouring task from 15% to 80%. The snapshots of different manipulation tasks are shown in Fig. 9.

VI. CONCLUSION AND DISCUSSION

In this paper, we introduce a novel framework called A4T that utilizes affordance detection results to achieve multi-step depth reconstruction and manipulation of transparent objects. The extensive experiments show that our multi-step depth reconstruction method is able to reconstruct the depth of regions enclosed by depth discontinued boundaries i.e. the “contain” regions and shows better performance on our collected **TRANS-AFF** dataset. The robot manipulation experiments demonstrate that our method is superior to the state-of-the-art method [4] in the manipulation tasks related to the affordance “contain,” i.e., the pouring task and the stacking task.

Our approach is not limited to the three affordance classes used in this work and has potential to be extended to other




Fig. 9: Snapshots of transparent object manipulation. **First row:** The robotic arm picks up the glass cup. **Second row:** The robotic arm pours the candies into the disposable cup. **Third row:** The robotic arm stacks one disposable cup into the other cup.

affordance classes such as a new affordance class “squeeze” for a dropper. Compared to opaque objects that can be made of a wide variety of materials, transparent objects are limited to materials like glass and plastics. As a consequence, most of them are fragile and tend to have fewer affordance classes. The affordance classes used in our work, i.e., “contain,” “wrap-grasp” and “support” represent the majority of the functionalities of the transparent objects in our daily lives and scientific laboratories.

We validate our approach on structured scenarios that have simple backgrounds and clean tables, as the environments where glasses exist such as chemistry laboratories are normally structured. However, the limitation of scene diversity may constrain the generalisation of our method to unseen challenging environments with complex backgrounds and varying lighting conditions. We could improve the algorithm generalisation by augmenting the current dataset, e.g., collecting more real-world data with objects under different kinds of backgrounds and lighting conditions, and generating a diverse synthetic affordance dataset with a photo-realistic render engine.

Our approach is an initial step towards reconstructing and manipulating the challenging transparent objects, and is constrained by the prediction of surface normals and contact edges. In cluttered environments, the occlusions between different objects may lead to noisy predictions of surface normals and invisible contact edges, and therefore result in failed depth reconstruction. It could be alleviated by replacing the global optimisation method with an end-to-end reconstruction method that is not dependent on surface normals and contact edges.

REFERENCES

- [1] B. Burger, P. M. Maffettone, V. V. Gusev, C. M. Aitchison, Y. Bai, X. Wang, X. Li, B. M. Alston, B. Li, R. Clowes, *et al.*, “A mobile robotic chemist,” *Nature*, vol. 583, no. 7815, pp. 237–241, 2020.
- [2] J. Mahler, F. T. Pokorny, B. Hou, M. Roderick, M. Laskey, M. Aubry, K. Kohlhoff, T. Kröger, J. Kuffner, and K. Goldberg, “Dex-net 1.0: A cloud-based network of 3d objects for robust grasp planning using a multi-armed bandit model with correlated rewards,” in *Proc. IEEE Int. Conf. Robot. Autom.*, pp. 1957–1964, 2016.
- [3] A. Murali, A. Mousavian, C. Eppner, C. Paxton, and D. Fox, “6-dof grasping for target-driven object manipulation in clutter,” in *Proc. IEEE Int. Conf. Robot. Autom.*, pp. 6232–6238, 2020.
- [4] S. Sajjan, M. Moore, M. Pan, G. Nagaraja, J. Lee, A. Zeng, and S. Song, “Clear grasp: 3d shape estimation of transparent objects for manipulation,” in *Proc. IEEE Int. Conf. Robot. Autom.*, pp. 3634–3642, 2020.
- [5] K. Han, K.-Y. K. Wong, and M. Liu, “A fixed viewpoint approach for dense reconstruction of transparent objects,” in *Proc. IEEE Conf. Comput. Vis. pattern Recognit.*, pp. 4001–4008, 2015.
- [6] Y. Qian, M. Gong, and Y. Hong Yang, “3d reconstruction of transparent objects with position-normal consistency,” in *Proc. IEEE Conf. Comput. Vis. pattern Recognit.*, pp. 4369–4377, 2016.
- [7] C. J. Phillips, M. Lecce, and K. Daniilidis, “Seeing glassware: from edge detection to pose estimation and shape recovery,” in *Proc. Robot. Sci. Syst.*, vol. 3, 2016.
- [8] U. Klank, D. Carton, and M. Beetz, “Transparent object detection and reconstruction on a mobile platform,” in *Proc. IEEE Int. Conf. Robot. Autom.*, pp. 5971–5978, 2011.
- [9] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein, “Implicit neural representations with periodic activation functions,” *Adv. Neural Inf. Process. Syst.*, vol. 33, pp. 7462–7473, 2020.
- [10] J. Ichnowski, Y. Avigal, J. Kerr, and K. Goldberg, “Dex-nerf: Using a neural radiance field to grasp transparent objects,” in *Proc. Conf. Rob. Learn.*, pp. 526–536, PMLR, 2022.
- [11] Y. Zhang and T. Funkhouser, “Deep depth completion of a single rgb-d image,” in *Proc. IEEE Conf. Comput. Vis. pattern Recognit.*, pp. 175–185, 2018.
- [12] L. Zhu, A. Mousavian, Y. Xiang, H. Mazhar, J. van Eenbergen, S. Debnath, and D. Fox, “Rgb-d local implicit function for depth completion of transparent objects,” in *Proc. IEEE/CVF Conf. Comput. Vis. pattern Recognit.*, pp. 4649–4658, 2021.
- [13] A. Zeng, S. Song, K.-T. Yu, E. Donlon, F. R. Hogan, M. Bauza, D. Ma, O. Taylor, M. Liu, E. Romo, *et al.*, “Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching,” in *Proc. IEEE Int. Conf. Robot. Autom.*, pp. 1–8, 2018.
- [14] R. Jain and T. Inamura, “Learning of tool affordances for autonomous tool manipulation,” in *Proc. IEEE/SICE Int. Symp. Syst. Integration*, pp. 814–819, IEEE, 2011.
- [15] T. Mar, V. Tikhonoff, G. Metta, and L. Natale, “Self-supervised learning of grasp dependent tool affordances on the icub humanoid robot,” in *Proc. IEEE Int. Conf. Robot. Autom.*, pp. 3200–3206, 2015.
- [16] T. Mar, V. Tikhonoff, G. Metta, and L. Natale, “Self-supervised learning of tool affordances from 3d tool representation through parallel som mapping,” in *Proc. IEEE Int. Conf. Robot. Autom.*, pp. 894–901, 2017.
- [17] A. Myers, C. L. Teo, C. Fermüller, and Y. Aloimonos, “Affordance detection of tool parts from geometric features,” in *Proc. IEEE Int. Conf. Robot. Autom.*, pp. 1374–1381, 2015.
- [18] T.-T. Do, A. Nguyen, and I. Reid, “Affordancenet: An end-to-end deep learning approach for object affordance detection,” in *Proc. IEEE Int. Conf. Robot. Autom.*, pp. 1–5, 2018.
- [19] A. Nguyen, D. Kanoulas, D. G. Caldwell, and N. G. Tsagarakis, “Detecting object affordances with convolutional neural networks,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, pp. 2765–2770, 2016.
- [20] F.-J. Chu, R. Xu, C. Tang, and P. A. Vela, “Recognizing object affordances to support scene reasoning for manipulation tasks,” *arXiv preprint arXiv:1909.05770*, 2019.
- [21] J. Long, E. Shelhamer, and T. Darrell, “Fully convolutional networks for semantic segmentation,” in *Proc. IEEE Conf. Comput. Vis. pattern Recognit.*, pp. 3431–3440, 2015.
- [22] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. IEEE Conf. Comput. Vis. pattern Recognit.*, pp. 770–778, 2016.
- [23] K. He, G. Gkioxari, P. Dollár, and R. Girshick, “Mask r-cnn,” in *Proc. IEEE Conf. Comput. Vis. pattern Recognit.*, pp. 2961–2969, 2017.
- [24] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, “Encoder-decoder with atrous separable convolution for semantic image segmentation,” in *Proc. European Conf. Comput. Vis.*, pp. 801–818, 2018.
- [25] M. Y. Yang and W. Förstner, “Plane detection in point cloud data,” in *Proc. Int. Conf. Mach. Control Guidance*, vol. 1, pp. 95–104, 2010.
- [26] H. Liu, X. Tang, and S. Shen, “Depth-map completion for large indoor scene reconstruction,” *Pattern Recognit.*, vol. 99, p. 107112, 2020.
- [27] R. Vidal, R. Tron, and R. Hartley, “Multiframe motion segmentation with missing data using powerfactorization and gpca,” *Int. J. Comput. Vis.*, vol. 79, no. 1, pp. 85–105, 2008.
- [28] D. Eigen, C. Puhrsch, and R. Fergus, “Depth map prediction from a single image using a multi-scale deep network,” in *Adv. Neural Inf. Process. Syst.*, pp. 2366–2374, 2014.