

RGB-D Grasp Detection via Depth Guided Learning with Cross-modal Attention

Ran Qin, Haoxiang Ma, Boyang Gao and Di Huang*

Abstract—Planar grasp detection is one of the most fundamental tasks to robotic manipulation, and the recent progress of consumer-grade RGB-D sensors enables delivering more comprehensive features from both the texture and shape modalities. However, depth maps are generally of a relatively lower quality with much stronger noise compared to RGB images, making it challenging to acquire grasp depth and fuse multi-modal clues. To address the two issues, this paper proposes a novel learning based approach to RGB-D grasp detection, namely Depth Guided Cross-modal Attention Network (DGCAN). To better leverage the geometry information recorded in the depth channel, a complete 6-dimensional rectangle representation is adopted with the grasp depth dedicatedly considered in addition to those defined in the common 5-dimensional one. The prediction of the extra grasp depth substantially strengthens feature learning, thereby leading to more accurate results. Moreover, to reduce the negative impact caused by the discrepancy of data quality in two modalities, a Local Cross-modal Attention (LCA) module is designed, where the depth features are refined according to cross-modal relations and concatenated to the RGB ones for more sufficient fusion. Extensive simulation and physical evaluations are conducted and the experimental results highlight the superiority of the proposed approach.

I. INTRODUCTION

Grasp detection, which aims to generate gripper configurations on objects, is essential to robotic manipulation in the real world, such as bin-picking and sorting. Early methods [1], [2] tackle this task by fitting the input to the templates with annotated grasps or analyzing shape characteristics in certain feature spaces. However, these methods either only work on known objects or suffer from high computational complexity, making them less practical for widespread use. Recently, with the rapid development of deep learning, data-driven methods [3]–[7] have dominated this field and shown great potential to fulfill a diversity of applications.

As a typical scenario, planar grasp detection has received increasing attention during the past decade, where the gripper configuration is usually represented as a 5-dimensional rotated rectangle on the image [8], and advanced object detection networks are adapted for prediction. Considering rich texture information in the color space, a large number of studies have been made in the literature [9]–[11], which take RGB images as input. Meanwhile, due to the innovation

This work is partly supported by the National Natural Science Foundation of China (62022011), the Research Program of State Key Laboratory of Software Development Environment (SKLSDE-2021ZX-04), and the Fundamental Research Funds for the Central Universities.

Ran Qin, Haoxiang Ma and Di Huang are with the State Key Laboratory of Software Development Environment, School of Computer Science and Engineering, Beihang University, Beijing, China.

Boyang Gao is with the Geometry Robotics and the School of Computer Science and Technology, Harbin Institute of Technology, Harbin, China.

*Corresponding author (email: dhuang@buaa.edu.cn).

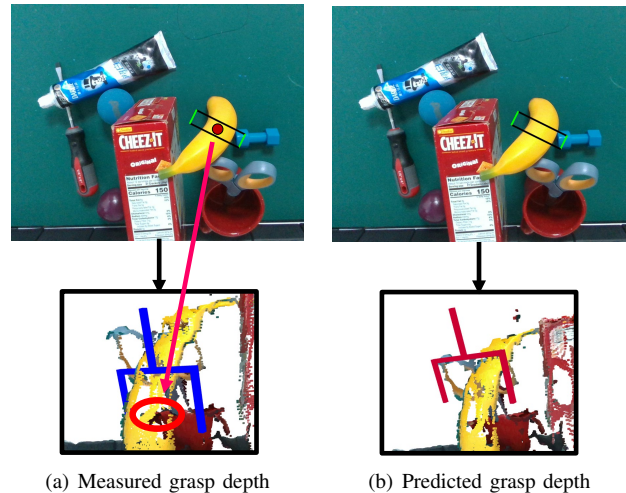


Fig. 1. Visualization of detected grasps which make use of depth clues in two different ways. The color of the grasp represents the score in terms of the force closure metric. The red indicates a grasp of a high quality while the blue denotes a grasp of a low quality. (a) shows the grasp which adopts the measured depth in the grasp center and (b) shows the grasp with the predicted depth achieved by the proposed method.

of depth sensors, a trend has appeared that object shapes are introduced to complement appearances [3], [12].

Despite promising performance reported in RGB-D grasp detection [3], [12], [13], two major issues still remain challenging, because for mainstream sensing devices, *e.g.* Intel RealSense and Microsoft Kinect, the data in the depth channel is of a relatively lower quality with much stronger noise compared to that in the RGB channel. One is how to effectively and efficiently acquire the grasp depth, and another is how to sufficiently combine the credits in the RGB and depth modalities.

For the former issue, a major manner is to directly employ the depth value of the raw data in the center of the predicted 5-dimensional rectangle [8], [14]. While efficient, the depth value is often not so accurate or even missing in such noisy images, which is prone to incur failed grasps colliding with objects, as shown in Fig. 1 (a). An alternative is to initially sample a set of patches and then estimate grasp depths from them [5], [15]–[17]. These methods indeed boost the accuracy to some extent, but the entire procedure is complicated and incurs huge computational cost, especially in cluttered scenes. Both the two strategies make the generalization to the real-world case problematic. For the latter issue, existing methods conduct the fusion of RGB and depth data by summing or concatenating at the early or middle stage of the networks [3], [13], [18], [19], where the two modalities are analogously treated. This straightforward way

confirms the complementarity of the two types of inputs; unfortunately, it does not fully take their properties into account. In particular, the depth images are generally coarse, and the extracted depth features tend to be misleading and thus bring difficulties in multi-modal clue integration. This dilemma is more serious when the networks become deeper [20], [21]. To the best of our knowledge, these two issues have not been well discussed and leave much space for improvement.

To address the issues aforementioned, in this paper, we propose a novel two-stage approach, namely Depth Guided Cross-modal Attention Network (DGCAN), to RGB-D grasp detection. Specifically, DGCAN builds a depth guided learning framework, where both the RGB and depth images are fed and their features are combined to generate grasp proposals. To better leverage the geometry clues in the noisy depth map, a complete 6-dimensional rectangle representation is adopted with the depth of the grasp dedicatedly considered besides its center, size, and orientation defined in the common 5-dimensional one. The prediction of the extra grasp depth substantially strengthens grasp features, thereby contributing to more accurate results. Furthermore, a Local Cross-modal Attention (LCA) module is designed to fuse the features which are separately encoded in the RGB and depth channels by two sub-networks. To reduce the negative impact caused by the discrepancy of data quality in two modalities, LCA conducts in an asymmetric way that depth features are refined according to cross-modal relations and concatenated to the RGB ones. In addition, a comprehensive RGB-D planar grasp dataset is produced based on GraspNet-1Billion [22] to facilitate this research. Extensive simulation and real-world evaluations are made and the experimental results highlight the superiority of the proposed approach.

II. RELATED WORK

A. RGB Grasp Detection

In the past decade, many efforts have been made for grasp detection on RGB images, and we can see its development from hand-crafted methods to learning based ones. Following a sample-based pipeline, [23] computes the edge and texture filter responses to decide whether an image patch contains a grasp point. Further, [24] extracts deep features by AlexNet [25] to estimate the graspable probabilities of different orientations. With the inspiration of the two-stage object detectors, such as Faster-RCNN [26] and RRPN [27], end-to-end trainable networks [10], [28], [29] are proposed to directly locate rotated grasp rectangles with consistently improved performance. By applying the backbones pretrained on ImageNet [25], they also alleviate the over-fitting risk, even with limited grasp annotations. Although RGB grasp detection methods achieve a great success, they are still not qualified enough for real-world applications due to the intrinsic ambiguity caused by 2D imaging techniques.

B. RGB-D Grasp Detection

Depth data convey geometry clues, which are necessary to complement appearance ones to build more accurate features for grasp detection. On the one hand, to acquire grasp depth,

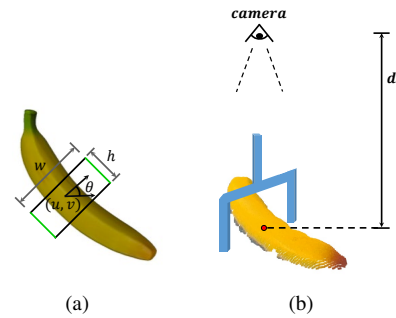


Fig. 2. (a) The 5-dimensional grasp configuration (u, v, w, h, θ) . (b) The grasp depth d in the 6-dimensional grasp configuration.

some attempts investigate the simple and direct strategy which uses the value of the depth image in the center of the predicted rectangle [8], [14], but the performance is sensitive to the quality of depth data, largely limited by current consumer-grade RGB-D sensors. Although some subsequent methods improve this by a sample-before-evaluate pipeline [5], [15]–[17], they suffer from high computational complexity, not easy to generalize to cluttered scenes. On the other hand, to jointly capture textural and geometric information, these methods conduct multi-modal combination in two major ways, *i.e.* early fusion and middle fusion. For early fusion, RGB and depth images are stacked before feeding to the network. [8] concatenates an RGB image and an aligned depth image as a 4-channel input. [3], [18] replace the blue channel in the RGB image with the depth channel to deliver a 3-channel matrix. However, as the distribution gap exists between modalities, it is difficult for a single network to achieve reconciliation. Regarding middle fusion, RGB and depth images are individually processed through two different networks with similar architectures and finally aggregated for prediction [13], [19], [30], [31]. These methods perform better than early fusion based ones, but they generally treat the two modalities in an analogous manner and ignore their quality difference, leading to sub-optimal results.

III. PROBLEM FORMULATION

Given an RGB image $I \in \mathbb{R}^{3 \times H \times W}$ and its corresponding depth image $D \in \mathbb{R}^{1 \times H \times W}$ as input, different from the common 5-dimensional rectangle representation for planar grasp detection proposed in [8], we highlight the necessity of grasp depth to RGB-D grasp detection and adopt a 6-dimensional rectangle with an additional variable for grasp depth, as Fig. 2 shows, which is formally defined as:

$$\mathbf{g} = (u, v, d, w, h, \theta) \quad (1)$$

where (u, v) , w , h and θ are the same as those defined in the 5-dimensional representation, denoting the center coordinate, width, height, and angle between the width and x -axis in the image of the grasp rectangle respectively. d is added to indicate the depth of the fingertips relative to the camera.

IV. METHODOLOGY

A. Overall Framework

With the 6-dimensional rectangle representation, DGCAN builds a depth guided learning framework. It follows the

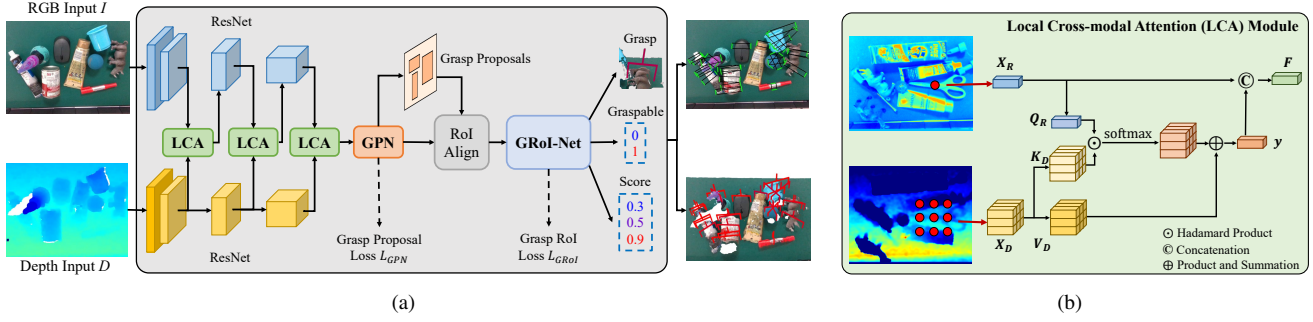


Fig. 3. (a) Overview of the proposed DGCAN framework. It consists of three main parts: Local Cross-modal Attention (LCA) based multi-modal fusion network, Grasp Proposal Network (GPN) and Grasp Region of Interest Network (GRoI-Net) with depth prediction. (b) Details of the LCA module. RGB features are treated as queries and depth features are treated as keys and values.

two-stage pipeline [3] but extends it by multi-modal fusion and grasp depth prediction. As in Fig. 3 (a), DGCAN is composed of three parts: *i.e.* Grasp Proposal Network (GPN), Grasp Region of Interest Network (GRoI-Net), and Local Cross-modal Attention (LCA) based fusion network.

Taking a 3-channel RGB image and a 1-channel depth image as input, two networks with the same architecture (*i.e.* ResNet [32]) are adopted to extract appearance and geometry features respectively, where the depth image is replicated to 3 channels before processing. To sufficiently utilize the RGB and depth information, the LCA module is designed and placed after each convolutional block of the network, which performs an asymmetric fusion by first refining the depth features through cross-modal relation learning and then concatenating them to the RGB ones. The multi-modal features in the last block are further fed into GPN to generate possible grasp regions and predict 6-dimensional grasps in GRoI-Net. In the inference stage, we apply Non-Maximum Suppression (NMS) after GPN to suppress duplicated grasp proposals and retain final grasps according to the graspability scores.

B. Grasp Proposal Network

As in [3], GPN is introduced as the first stage to estimate potential grasp candidates over the whole image. Concretely, given the multi-modal feature maps as input, it predicts the grasp bounding box $\hat{g} = (\hat{u}, \hat{v}, \hat{w}, \hat{h})$ and its probability of grasp proposal \hat{p} for each anchor.

Compared to the typical Region Proposal Network (RPN) [26] applied in the two-stage object detectors, GPN aims to find the optimal match from a dense collection of positive grasp labels with different scores. In this case, besides setting an Intersection over Union (IoU) threshold to filter Ground Truths (GTs) by positions and scales, only the grasp with the highest score s is assigned to the anchor. The loss of GPN is defined as follows:

$$L_{GPN}(\{(\hat{p}_i, \hat{t}_i)\}_{i=1}^N) = \frac{1}{N_{cls}} \sum_i L_{GPN}^{cls}(\hat{p}_i, \hat{p}_i^*) + \lambda_1 \frac{1}{N_{reg}} \sum_i \hat{p}_i^* L_{GPN}^{reg}(\hat{t}_i, \hat{t}_i^*) \quad (2)$$

where L_{GPN}^{cls} is the cross entropy loss of binary classification estimating whether an anchor is graspable, L_{GPN}^{reg} is the Smooth L1 loss for grasp proposal regression, and λ_1 is a

weight parameter. N_{cls} and N_{reg} indicate the number of sampled anchors and number of positive samples respectively. \hat{p}_i and \hat{p}_i^* denote the predicted probability and classification label of the i -th anchor. $\hat{p}_i^* = 1$ means a specific grasp and $\hat{p}_i^* = 0$ indicates that there is no grasp. \hat{t}_i and \hat{t}_i^* represent the predicted and GT of the 4-dimensional grasp vector. The predicted grasp proposal $(\hat{u}, \hat{v}, \hat{w}, \hat{h})$ is calculated as:

$$\begin{aligned} \hat{t}_u &= (\hat{u} - u_a)/w_a, & \hat{t}_v &= (\hat{v} - v_a)/h_a, \\ \hat{t}_w &= \log(\hat{w}/w_a), & \hat{t}_h &= \log(\hat{h}/h_a) \end{aligned} \quad (3)$$

where (u_a, v_a, w_a, h_a) denotes the grasp vector of an anchor. Then, RoI Align is employed to extract features from grasp proposals and feed them into the subsequent GRoI-Net.

C. Grasp Region of Interest Network

For each RoI, GRoI-Net further extracts the feature on it using the fourth block C_4 in ResNet [32] and predicts the 6-dimensional grasp configuration $\mathbf{g} = (u, v, d, w, h, \theta)$ with its grasp possibility q and grasp score s . Therefore, given M grasp proposals, the outputs of GRoI-Net are $M \times 7$, $M \times 2$ and $M \times 1$ for grasp target regression $t = (t_u, t_v, t_d, t_w, t_h, t_{sin}, t_{cos})$, binary classification q and grasp score estimation s , respectively. Particularly, for grasp depth prediction, we consider the depth measurement d_o captured by the RGB-D camera as the reference depth and regress the offset t_d defined as:

$$t_d = d - d_o. \quad (4)$$

Meanwhile, we encode the angle as two components $t_{sin} = \sin(2\theta)$ and $t_{cos} = \cos(2\theta)$ in the range $[-1, 1]$ and calculate the grasp orientation as:

$$\theta = \frac{1}{2} \arctan \frac{\sin(2\theta)}{\cos(2\theta)}. \quad (5)$$

With this definition, the target function of GRoI-Net is defined as follows:

$$L_{GRoI}(\{(q_i, t_i, s_i)\}_{i=1}^M) = \frac{1}{M_{cls}} \sum_i L_{GRoI}^{cls}(q_i, q_i^*) + \lambda_2 \frac{1}{M_{reg}} \sum_i L_{GRoI}^{reg}(t_i, t_i^*) + \lambda_3 \frac{1}{M_{reg}} \sum_i L_{GRoI}^{reg}(s_i, s_i^*) \quad (6)$$

TABLE I
SUMMARY OF THE PROPERTIES OF PUBLIC PLANAR GRASP DATASETS.

Dataset	Grasp Depth	Modality	Objects /Img	Grasps /Img	Num Imgs	Num Objects	Num Grasps
Levine <i>et al.</i> [33]	No	RGB-D	-	1	800K	-	800K
Pinto <i>et al.</i> [24]	No	RGB-D	-	1	50K	150	50K
Cornell [8]	No	RGB-D	1	~8	1035	240	8K
Mahler <i>et al.</i> [5]	Yes	D	1	1	6.7M	1500	6.7M
Jacquard [34]	No	RGB-D	1	~20	54k	11K	1.1M
VMRD [9]	No	RGB	~3	~20	4.7K	~200	100K
Multi-Object [3]	No	RGB-D	~4	~30	96	-	2904
REGRAD [35]	No	RGB-D	1~20	1.02K	900K	50K	100M
GraspNet-1Billion [22]	No	RGB-D	~10	3~9M	97K	88	~1.2B
GraspNet-Planar	Yes	RGB-D	~10	~30K	~11K	88	169M

where q_i^* , t_i^* and s_i^* denote the GTs for q_i , t_i and s_i . M_{cls} and M_{reg} are similar to N_{cls} and N_{reg} in Eq. 2. Here, the cross entropy loss is used for L_{GRoI}^{cls} , while the Smooth L1 loss is used for L_{GRoI}^{reg} . λ_2 and λ_3 are weight parameters.

D. Local Cross-modal Attention based Fusion Network

The previous RGB-D grasp detection networks [13], [19], [30] generally integrate RGB and depth features by summation or concatenation with the two modalities analogously treated. However, the data in the depth channel is of a lower quality than that in the RGB channel in existing RGB-D sensors and these methods do not fully take this gap into account, thus leading to sub-optimal fusion. To handle this, we dedicatedly design a new module, *i.e.* LCA, to combine the RGB and depth features in an asymmetric manner, where the relatively stable RGB features are used to refine the noisy depth ones. Furthermore, as local shape perception is critical to grasp detection, for the depth feature in each position, only the ones within a neighborhood are involved. As illustrated in Fig. 3 (b), LCA first refines depth features by weighting them according to the relation between the RGB and depth modalities in a local region and then concatenates the weighted depth features to those in the RGB stream.

Denote the RGB and depth feature as $X_R \in \mathbb{R}^{C \times H \times W}$ and $X_D \in \mathbb{R}^{C \times H \times W}$ respectively, the output of LCA F is formulated as:

$$F = f(X_R || r(X_R, X_D)) \quad (7)$$

where $||$ is the concatenation operation of the feature maps from the two modalities, f refers to 1×1 convolution with the number of channels modified from $2C$ to C , and r represents the depth refinement operator based on cross-modal attention. Concretely, X_R serves as the *query* and is embedded to $Q_R \in \mathbb{R}^{C' \times H \times W}$; X_D serves as the *key* and *value*, embedded to $K_D \in \mathbb{R}^{C' \times H \times W}$ and $V_D \in \mathbb{R}^{C' \times H \times W}$ by two different 1×1 convolutions. Inspired by [36], we design an attention module to capture the relationship between the RGB feature and the corresponding depth features within a neighborhood considering the importance of local geometric features to grasp detection. Given an RGB query at (i, j) , we extract depth keys and values in the neighboring region $N_k(i, j)$ with spatial size k and calculate the depth output v_{ij} as follows:

$$v_{ij} = \sum_{(m,n) \in N_k(i,j)} \text{softmax}(Q_{R,ij} K_{D,mn} / \sqrt{C'}) V_{D,mn} \quad (8)$$

where $(m, n) = (i + \Delta i, j + \Delta j)$, $\Delta i = -\lceil k/2 \rceil, -\lceil k/2 \rceil + 1, \dots, \lceil k/2 \rceil$, $\Delta j = -\lceil k/2 \rceil, -\lceil k/2 \rceil + 1, \dots, \lceil k/2 \rceil$. We then project the outputs v to the final depth features $y \in \mathbb{R}^{C \times H \times W}$ by 1×1 convolution. Following the self-attention [37] strategy, we implement the multi-head attention mechanism and add sine positional encodings in queries and keys.

E. Loss Function

During training, the whole network is optimized in an end-to-end manner by minimizing the following loss function:

$$L = L_{GPN}(\{\{\hat{p}_i, \hat{t}_i\}_{i=1}^N\}) + L_{GRoI}(\{\{q_i, t_i, s_i\}_{i=1}^M\}) \quad (9)$$

V. GRASPNET-PLANAR DATABASE

Current public RGB-D planar grasp data do not contain grasp depth annotations, we thus build a new benchmark based on a 6-DoF grasp dataset, *i.e.* GraspNet-1Billion [22], for evaluation. Following the same protocol as GraspNet-1Billion, we take 100 scenes for training and 90 scenes for testing, where the test set is divided into three parts for seen, similar and novel objects respectively. Because the camera pose of planar grasping is constrained to be perpendicular to the workspace plane from top to down, we only use the images where the angles between the camera's observation view and z -axis are smaller than 15° . In general, our benchmark consists of 5,513 RGB-D images from RealSense D435 and 5,341 images from Azure Kinect.

Although GraspNet-1Billion provides transformation matrices from 6-DoF grasps to planar grasps, the effect of grasp depths is ignored, which incurs inaccurate grasp annotations. To deal with this, we present a new pipeline. Specifically, for each object in the scene, we first select grasps with small angles between the grasp approaches and camera view. The approach directions of the selected 6-DoF grasps are set the same as the camera view. This projection may result in failed grasp labels due to the change of grasp poses; therefore, we re-evaluate these grasps by collision detection as well as force closure metric [38]. We assign new force-closure scores to the successful grasps as grasp scores and remove the failed samples. Finally, to obtain the gripper configurations of planar grasp detection, the grasp label \mathbf{G} in the $SE(3)$ space is projected to the image plane as follows:

$$\mathbf{G} = T_{CIg} \quad (10)$$

TABLE II

PERFORMANCE COMPARISON ON GRASPNET-PLANAR CAPTURED BY REALSENSE/KINECT. *CD* REPRESENTS COLLISION DETECTION.

Method	Seen			Similar			Novel		
	AP	AP _{0.8}	AP _{0.4}	AP	AP _{0.8}	AP _{0.4}	AP	AP _{0.8}	AP _{0.4}
GR-ConvNet [12]	19.97/17.72	25.42/22.23	11.43/10.60	13.24/13.04	16.93/16.39	6.79/7.67	6.32/4.41	7.67/5.41	1.69/1.52
GG-CNN2 [14]	27.45/23.73	35.31/30.20	16.60/14.90	21.49/18.02	27.87/23.44	11.27/10.12	9.76/7.12	11.66/8.77	2.95/3.14
Chu <i>et al.</i> [3]	29.71/25.69	35.98/30.70	23.76/20.96	24.17/20.92	30.00/25.44	16.86/15.87	10.92/6.70	13.38/8.23	6.69/3.88
DGCAN	49.85/47.32	59.67/57.27	42.24/38.55	41.46/35.73	50.31/44.22	33.69/26.99	17.48/16.10	21.83/20.01	7.90/7.81
DGCAN- <i>CD</i>	52.16/50.45	62.71/61.22	43.14/40.64	44.69/38.62	54.52/47.85	35.37/28.81	19.26/17.66	23.93/21.94	8.89/8.29

where T_{CI} is the intrinsic parameter matrix of the camera and \mathbf{g} is described in Eq. 1. Table I compares our GraspNet-Planar benchmark with major public counterparts. It provides not only multi-modal RGB-D images but also more complete annotations with grasp depths than others, which highlights its advantage for RGB-D planar grasp detection research.

VI. EXPERIMENTS

A. Protocols

The simulation evaluation on GraspNet-Planar follows the metric in the GraspNet-1Billion benchmark [22], *i.e.* Average Precision (AP). Given the predicted 6-dimensional planar grasps, we calculate AP_{μ} with different friction coefficient μ after grasp-NMS [22]. As in GraspNet-1Billion, the overall result AP is calculated by averaging of AP_{μ} , where μ ranges from 0.2 to 1.0 with the interval $\Delta\mu = 0.2$.

For physical evaluation, we choose 25 objects of various sizes, shapes and textures from the YCB Object Set [39], as shown in Fig. 4 (a). We conduct all real-world experiments on a 7-DoF robotic arm in two settings: single-object grasping and multi-object grasping. For single-object grasping, 25 objects are employed and each object is evaluated in three different poses. For multi-object grasping, we construct 10 cluttered scenes, each of which contains 5 different objects. Grasping methods are required to remove the objects in scenes as many as possible within 10 attempts. As for the metric, Grasp Success Rate (GSR) is used in both single-object grasping and multi-object grasping, and Scene Completion Rate (SCR) is additionally recorded for multi-object grasping.

B. Implementation Details

We build our network based on ResNet-50 [32]. For LCA in the three residual blocks, we adopt the local spatial size $k = (5, 3, 3)$, which is further discussed in Sec. VI-E. The anchor aspect ratios and scales in GPN are set to [0.5, 1, 2] and [32, 64, 128, 256], respectively.




Fig. 4. (a) The 25 household objects used in physical evaluation. (b) The hardware set-up with a 7-DoF Agile Diana-7 robot arm and an Intel RealSense D435i camera tied on the arm.




Fig. 5. Qualitative visualization on GraspNet-Planar captured by RealSense (the first 2 columns) and Kinect (the last column) after grasp-NMS [22]. The color in the second row denotes the predicted scores of grasp poses ranging from 0 to 1. The blue color represents 0 and the red color represents 1. Zoom for better views.

During training, the model is initialized by the weights pre-trained on ImageNet [25]. We choose 512 samples for both GPN and GRoI-Net, where the ratios of positive and negative samples are respectively 1:1 and 1:3. After GPN, we take 2,000 proposals whose IoUs are lower than 0.7. The weights in the loss function are set as $\lambda_1 = 1$, $\lambda_2 = 1$ and $\lambda_3 = 4$. During inference, we retain the predicted grasps with the graspable scores higher than 0.5 to deliver the final result.

The model is trained on four GTX1080Ti GPUs by the SGD optimizer with the momentum and weight decay set to 0.9 and 0.0001. The mini-batch size is 4 and the initial learning rate is 5×10^{-4} and decays at 43k iterations. Before training, grasp-NMS [22] is launched to filter out largely overlapped GTs. Random horizontal flipping and rotating are employed to augment input images. Besides, the hole-filling filter and bilateral filter are applied to depth images for smoothing.

In physical evaluation, we use a 7-DoF Agile Diana-7 robot arm with an Intel RealSense D435i camera tied on it. The RGB-D images are captured above the workspace, as shown in Fig. 4 (b). Our network and control algorithm are equipped on a desktop with an AMD Ryzen 5 2600 six-core processor and a single NVIDIA GeForce 1080 GPU.

C. Simulation Evaluation

Table II summarizes the results of different solutions on the GraspNet-Planar dataset. We compare the proposed approach to three representative end-to-end trainable planar grasp detectors [3], [12], [14]. [3] is a two-stage network which is built based on Faster-RCNN [26] and [12], [14] follow the one-stage pipeline which output dense predictions of grasp qualities, widths and orientations. It can be seen that our network outperforms the other state-of-the-art counterparts by a large margin, which demonstrates the effectiveness of grasp depth prediction and the LCA module. Some qualitative results are shown in Fig. 5, which indicate the validity of our approach in cluttered scenes.

TABLE III
PHYSICAL ROBOT EVALUATION RESULTS

Method	Single-Object	Multi-Object	
	GSR (%)	GSR (%)	SCR (%)
Dex-Net 4.0 [40]	70.67 (53/75)	67.14 (47/70)	94.00 (47/50)
FC-GQ-CNN [15]	73.33 (55/75)	67.61 (48/71)	96.00 (48/50)
DGCAN	88.00 (66/75)	84.75 (50/59)	100.00 (50/50)

TABLE IV
ABLATION STUDY OF GRASP DEPTH PREDICTION ON
GRASPNET-PLANAR CAPTURED BY REALSENSE.

Grasp Depth	Seen	Similar	Novel
Center	42.01	34.54	15.25
Classification	48.57	39.01	17.06
Ours	49.85	41.46	17.48

D. Physical Evaluation

In real-world experiments, Dex-Net 4.0 [40] and FC-GQ-CNN [15] which also involve the policies of acquiring grasp depths are employed for comparison. Dex-Net 4.0 takes measurements of grasp centers in depth images as grasp depths and FC-GQ-CNN uniformly samples depth values to several bins within the range of the whole scene and then evaluates them. As shown in Table III, Dex-Net 4.0 and FC-GQ-CNN achieve similar performance which fall behind ours in both settings.

E. Ablation Study

Influence of depth guided learning. We use two baselines, both of which deliver grasp depths. One is *Center*, where we predict 5-dimensional grasp rectangles and set the sum of the depth measurements of the rectangle centers and a fixed offset of 20 millimeters as the grasp depths. Another is *Classification*. Here, similar to [15], the prediction of grasp depths is modeled as a classification problem. We replace the regression head of grasp depths with a classification head. The grasp depths are uniformly divided into 40 bins from the maximum to the minimum depth value of scenes. The width of each bin is about 6 millimeters. It should be noted that for *Classification* we train the networks with different numbers of depth bins ranging from 25 to 50 by an interval of 5 and the one with 40 bins achieves the best performance. For fair comparison, we adopt very similar network architectures in all the models, which only differ in the head of GRoI-Net.

As shown in Table IV, the model that acquires grasp depths by *Center* reaches the lowest performance among the three, which indicates the significance of replacing the measured grasp depth to the predicted depth. When compared to *Classification*, the proposed approach has respectively gained 1.28%, 2.45% and 0.42% improvements on seen, similar and novel objects, proving that the grasp depth predicted by regression is more accurate than that by classification. Besides, *Classification* brings extra hyper-parameters, which increases the design complexity, such as the range of depth values and the number of depth bins.

Influence of LCA based multi-modal fusion. An unimodal model (*RGB*) taking only RGB images as input is evaluated as a baseline. We also employ an early fusion baseline following [3], *i.e.* *RGD*, which replaces the blue channel

TABLE V
ABLATION STUDY OF MULTI-MODAL FUSION MECHANISM ON
GRASPNET-PLANAR CAPTURED BY REALSENSE.

Fusion Method	Seen	Similar	Novel
RGB	40.83	27.76	9.94
RGD	47.87	38.05	16.07
RGB-D-Sum	48.08	38.95	15.43
RGB-D-Concat	48.61	40.46	16.82
LCA	49.85	41.46	17.48

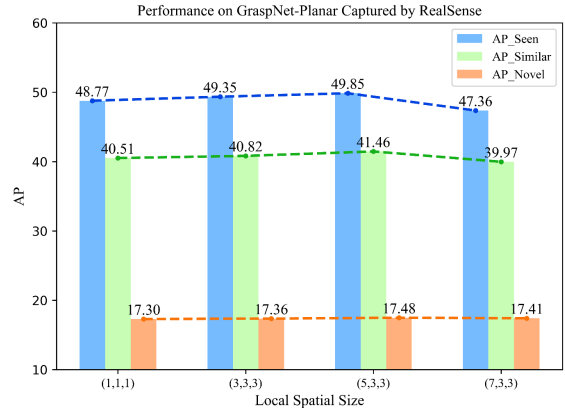


Fig. 6. AP under various local spatial sizes (k_1, k_2, k_3) for different stages (C_1, C_2, C_3) of ResNet.

with the depth channel. As shown in Table V, although jointly using RGB and depth data boosts the performance, the early fusion strategy is sub-optimal because of the distribution gap between two modalities. Furthermore, we evaluate another two middle fusion baselines, which exploit independent feature extractors to process RGB and depth data respectively and sum (*RGB-D-Sum*) or concatenate (*RGB-D-Concat*) the two types of features. The same as in LCA, multi-modal features are fused at the three residual blocks of the ResNet models. The results in Table V demonstrate that LCA outperforms the other middle fusion strategies, highlighting its superiority in RGB-D feature fusion for grasp detection.

To evaluate the influence of the local spatial extent in LCA, we train our network with different combinations of k_i in each block of ResNet, $i = 1, 2, 3$. The results are presented in Fig. 6 and as it shows, the LCA module achieves consistent improvements with different settings of local spatial sizes. However, when the size of the first block is set too large, there is a drop. Based on the ablations, we choose $(k_1, k_2, k_3) = (5, 3, 3)$ to reach the best performance.

VII. CONCLUSION

This paper proposes a novel two-stage approach, *i.e.* DGCAN, to RGB-D grasp detection. A complete 6-dimensional rectangle representation is adopted to emphasize the necessity of the grasp depth, which is not considered in the common 5-dimensional definition. The prediction of the grasp depth substantially improves grasp feature learning, thereby leading to more accurate results. Besides, an LCA module is designed, where the depth features are refined according to cross-modal relations and concatenated to the RGB ones for more sufficient fusion. The results of the simulation and physical experiments demonstrate its effectiveness.

REFERENCES

- [1] A. Bicchi and V. Kumar, "Robotic grasping and contact: A review," in *IEEE International Conference on Robotics and Automation*, 2000, pp. 348–353.
- [2] D. Buchholz, M. Futterlieb, S. Winkelbach, and F. M. Wahl, "Efficient bin-picking and grasp planning based on depth data," in *IEEE International Conference on Robotics and Automation*, 2013, pp. 3245–3250.
- [3] F. Chu, R. Xu, and P. A. Vela, "Real-world multiobject, multigrasp detection," *IEEE Robotics Autom. Lett.*, vol. 3, no. 4, pp. 3355–3362, 2018.
- [4] D. Morrison, J. Leitner, and P. Corke, "Closing the loop for robotic grasping: A real-time, generative grasp synthesis approach," in *Robotics: Science and Systems*, 2018.
- [5] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics," in *Robotics: Science and Systems*, 2017.
- [6] H. Ma and D. Huang, "Towards scale balanced 6-dof grasp detection in cluttered scenes," in *Conference on Robot Learning*, 2022.
- [7] Y. Wang, Y. Zheng, B. Gao, and D. Huang, "Double-dot network for antipodal grasp detection," in *IEEE International Conference on Intelligent Robots and Systems*, 2021, pp. 4654–4661.
- [8] Y. Jiang, S. Moseson, and A. Saxena, "Efficient grasping from RGBD images: Learning using a new rectangle representation," in *IEEE International Conference on Robotics and Automation*, 2011, pp. 3304–3311.
- [9] H. Zhang, X. Lan, S. Bai, X. Zhou, Z. Tian, and N. Zheng, "Roi-based robotic grasp detection for object overlapping scenes," in *IEEE International Conference on Intelligent Robots and Systems*, 2019, pp. 4768–4775.
- [10] X. Zhou, X. Lan, H. Zhang, Z. Tian, Y. Zhang, and N. Zheng, "Fully convolutional grasp detection network with oriented anchor box," in *IEEE International Conference on Intelligent Robots and Systems*, 2018, pp. 7223–7230.
- [11] D. Guo, F. Sun, H. Liu, T. Kong, B. Fang, and N. Xi, "A hybrid deep architecture for robotic grasp detection," in *IEEE International Conference on Robotics and Automation*, 2017, pp. 1609–1614.
- [12] S. Kumra, S. Joshi, and F. Sahin, "Antipodal robotic grasping using generative residual convolutional neural network," in *IEEE International Conference on Intelligent Robots and Systems*, 2020, pp. 9626–9633.
- [13] S. Kumra and C. Kanan, "Robotic grasp detection using deep convolutional neural networks," in *IEEE International Conference on Intelligent Robots and Systems*, 2017, pp. 769–776.
- [14] D. Morrison, P. Corke, and J. Leitner, "Learning robust, real-time, reactive robotic grasping," *Int. J. Robotics Res.*, vol. 39, no. 2-3, 2020.
- [15] V. Satish, J. Mahler, and K. Goldberg, "On-policy dataset synthesis for learning robot grasping policies using fully convolutional deep networks," *IEEE Robotics Autom. Lett.*, vol. 4, no. 2, pp. 1357–1364, 2019.
- [16] A. Gariépy, J. Ruel, B. Chaib-draa, and P. Giguère, "GQ-STN: optimizing one-shot grasp detection based on robustness classifier," in *IEEE International Conference on Intelligent Robots and Systems*, 2019, pp. 3996–4003.
- [17] W. Wang, W. Liu, J. Hu, Y. Fang, Q. Shao, and J. Qi, "Graspfusionnet: a two-stage multi-parameter grasp detection network based on RGB-XYZ fusion in dense clutter," *Mach. Vis. Appl.*, vol. 31, no. 7, p. 58, 2020.
- [18] J. Redmon and A. Angelova, "Real-time grasp detection using convolutional neural networks," in *IEEE International Conference on Robotics and Automation*, 2015, pp. 1316–1322.
- [19] A. Zeng, S. Song, K. Yu, E. Donlon, F. R. Hogan, M. Bauzá, D. Ma, O. Taylor, M. Liu, E. Romo, N. Fazeli, F. Alet, N. C. Dafle, R. Holladay, I. Morona, P. Q. Nair, D. Green, I. J. Taylor, W. Liu, T. A. Funkhouser, and A. Rodriguez, "Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching," in *IEEE International Conference on Robotics and Automation*, 2018, pp. 1–8.
- [20] W. Ji, J. Li, S. Yu, M. Zhang, Y. Piao, S. Yao, Q. Bi, K. Ma, Y. Zheng, H. Lu, and L. Cheng, "Calibrated RGB-D salient object detection," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2021, pp. 9471–9481.
- [21] S. Chen and Y. Fu, "Progressively guided alternate refinement network for RGB-D salient object detection," in *European Conference on Computer Vision*, 2020, pp. 520–538.
- [22] H. Fang, C. Wang, M. Gou, and C. Lu, "Graspnet-1billion: A large-scale benchmark for general object grasping," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11441–11450.
- [23] A. Saxena, J. Driemeyer, and A. Y. Ng, "Robotic grasping of novel objects using vision," *Int. J. Robotics Res.*, vol. 27, no. 2, pp. 157–173, 2008.
- [24] L. Pinto and A. Gupta, "Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours," in *IEEE International Conference on Robotics and Automation*, 2016, pp. 3406–3413.
- [25] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems*, 2012, pp. 1106–1114.
- [26] S. Ren, K. He, R. B. Girshick, and J. Sun, "Faster R-CNN: towards real-time object detection with region proposal networks," in *Advances in Neural Information Processing Systems*, 2015, pp. 91–99.
- [27] J. Ma, W. Shao, H. Ye, L. Wang, H. Wang, Y. Zheng, and X. Xue, "Arbitrary-oriented scene text detection via rotation proposals," *IEEE Trans. Multimed.*, vol. 20, no. 11, pp. 3111–3122, 2018.
- [28] A. Depierre, E. Dellandréa, and L. Chen, "Scoring graspability based on grasp regression for better grasp prediction," in *IEEE International Conference on Robotics and Automation*, 2021, pp. 4370–4376.
- [29] S. Ainetter and F. Fraundorfer, "End-to-end trainable deep neural network for robotic grasp detection and semantic segmentation from RGB," in *IEEE International Conference on Robotics and Automation*, 2021, pp. 13452–13458.
- [30] H. Zhu, Y. Li, F. Bai, W. Chen, X. Li, J. Ma, C. S. Teo, P. Y. Tao, and W. Lin, "Grasping detection network with uncertainty estimation for confidence-driven semi-supervised domain adaptation," in *IEEE International Conference on Intelligent Robots and Systems*, 2020, pp. 9608–9613.
- [31] Y. Song, J. Wen, D. Liu, and C. Yu, "Deep robotic grasping prediction with hierarchical rgb-d fusion," *International Journal of Control, Automation and Systems*, vol. 20, no. 1, pp. 243–254, 2022.
- [32] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [33] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, "Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection," *Int. J. Robotics Res.*, vol. 37, no. 4-5, pp. 421–436, 2018.
- [34] A. Depierre, E. Dellandréa, and L. Chen, "Jacquard: A large scale dataset for robotic grasp detection," in *IEEE International Conference on Intelligent Robots and Systems*, 2018, pp. 3511–3516.
- [35] H. Zhang, D. Yang, H. Wang, B. Zhao, X. Lan, J. Ding, and N. Zheng, "REGRAD: A large-scale relational grasp dataset for safe and object-specific robotic grasping in clutter," *IEEE Robotics Autom. Lett.*, vol. 7, no. 2, pp. 2929–2936, 2022.
- [36] P. Ramachandran, N. Parmar, A. Vaswani, I. Bello, A. Levskaya, and J. Shlens, "Stand-alone self-attention in vision models," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [37] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [38] V. Nguyen, "Constructing force-closure grasps," *Int. J. Robotics Res.*, vol. 7, no. 3, pp. 3–16, 1988.
- [39] B. Çalli, A. Walsman, A. Singh, S. S. Srinivasa, P. Abbeel, and A. M. Dollar, "Benchmarking in manipulation research: Using the yale-cmu-berkeley object and model set," *IEEE Robotics Autom. Mag.*, vol. 22, no. 3, pp. 36–52, 2015.
- [40] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg, "Learning ambidextrous robot grasping policies," *Sci. Robotics*, vol. 4, no. 26, 2019.