

End-to-end Trainable Deep Neural Network for Robotic Grasp Detection and Semantic Segmentation from RGB

Stefan Ainetter¹ and Friedrich Fraundorfer¹

Abstract—In this work, we introduce a novel, end-to-end trainable CNN-based architecture to deliver high quality results for grasp detection suitable for a parallel-plate gripper, and semantic segmentation. Utilizing this, we propose a novel refinement module that takes advantage of previously calculated grasp detection and semantic segmentation and further increases grasp detection accuracy. Our proposed network delivers state-of-the-art accuracy on two popular grasp dataset, namely Cornell and Jacquard. As additional contribution, we provide a novel dataset extension for the OCID dataset, making it possible to evaluate grasp detection in highly challenging scenes. Using this dataset, we show that semantic segmentation can additionally be used to assign grasp candidates to object classes, which can be used to pick specific objects in the scene.

I. INTRODUCTION

Automated grasping is a very active field of research in robotics. The process of having a robot manipulator to successfully grasp objects in a cluttered environment is still a challenging problem. Recent research on robotic grasping led to highly accurate grasp detection in scenes with multiple objects. Yang et al. [1] combine grasp detection with a CRF-based semantic model for task-oriented grasping and object detection. The works of [2], [3] combine grasp detection and object detection to predict grasp candidates that are assigned to specific objects in the scene. They showed that object detection provides accurate information about which objects are in the scene, however little information about the object's shape is provided. This can be especially critical in highly stacked scenes, where multiple objects overlap each other. In this work, we address grasp detection for parallel-plate grippers as well as semantic segmentation, by proposing an end-to-end trainable multi-task deep neural network. We show that one single network is able to provide high quality results for grasp detection and semantic segmentation in challenging scenes. Furthermore, we propose a novel refinement module which combines information about previously calculated grasp candidates and semantic segmentation to further increase the grasp detection accuracy. At last, we propose an extension to the OCID dataset [4], by adding hand-annotated grasp candidates and class labels to each object. Evaluation of our proposed model on this OCID dataset extension indicates high accuracy for grasp detection

¹All authors are with the Institute of Computer Graphics and Vision, Graz University of Technology, 8010 Graz, Austria {stefan.ainetter, fraundorfer}@icg.tugraz.at

The research leading to these results has received funding from the Austrian Ministry for Transport, Innovation and Technology (BMVIT) within the ICT of the Future Programme (4th call) of the Austrian Research Promotion Agency (FFG) under grant agreement n. 864807.



Fig. 1. Grasp detection and segmentation results of our proposed model for an image containing multiple graspable objects. **Top left** shows all predicted grasp candidates in the scene (blue lines denote parallel plates of the gripper, red lines denote opening width), **top right** shows the predicted semantic segmentation. **Bottom left** shows that combining both leads to valid grasp candidates for each object in the scene, which makes it possible to decide which object to pick. **Bottom right** shows the results for object detection represented by axis-aligned bounding boxes. Compared to the semantic segmentation, the bounding box representation is neither suitable to determine which grasp candidate belongs to which object, nor to determine the correct relationship between objects, as boxes highly overlap each other.

and segmentation in complex scenes with multiple objects. We also show that semantic segmentation can be used to assign grasp candidates to objects, which makes it possible to pick specific objects in the scene. Figure 1 shows results for our novel OCID dataset extension.

The main contributions of our paper are the following:

- 1) An end-to-end trainable deep neural network architecture for joint grasp detection and dense, pixel-wise semantic segmentation, which yields state-of-the-art performance for grasp detection.
- 2) A novel grasp refinement module, combining results from grasp detection and segmentation to further improve the overall grasp accuracy.
- 3) An extension of the OCID dataset [4] for robotic grasping, by adding ground truth grasp candidates as oriented bounding boxes and class information to each object.

II. RELATED WORK

Our work is related to prior research on object detection and semantic segmentation in robotic vision. More specific, it is related to robotic grasp detection and scene understanding for robotic manipulations.

A. Robotic Grasp Detection

The conventional method for grasp detection uses information about object geometry, physics models and force analytics [5]. With the rise of deep learning, data-driven methods [6] became more common. Methods like [7], [8] were the first to use deep neural networks and supervised learning to predict multiple grasp candidates for a single object. Methods like [9], [10], [11], [12], [13] perform grasp detection based on two-stage object detectors. This type of detectors consists of a region proposal network and a detector. After extracting features using proposals from the first stage, objects are detected in the second stage. Approaches like [3], [14] are single-stage detectors which divide the input image using a grid and perform detection on each cell. This usually reduces the computation time, with decreased prediction accuracy.

Other methods [15], [16] perform grasp detection on depth images as input, where accuracy is highly dependent on the quality of the input data.

Kumra et al. [17] use a generative architecture to predict pixel-wise grasps in the form of three images, namely quality, angle and width. Asif et al. [18], [19] use an ensemble of neural networks for grasp detection and segmentation, whereas using multiple networks is usually more expensive in terms of computational cost and memory consumption. DSGD [20] uses image information on different levels of the image hierarchy (global-, region-, pixel-level) for grasp detection. The works of [1], [2] perform the tasks of object detection with reasoning and grasp detection for complex scenes with multiple objects. Other approaches [21], [22] use Reinforcement Learning (RL) on a real or simulated robot to perform thousands of grasp attempts and use the feedback to improve grasp detection. RL has the advantage that no labeled data is necessary for training, but with the downside of being time and hardware consuming.

B. Semantic Segmentation

Several works use encoder/decoder based CNN architectures [23], [19] or dilated convolutions [24] for semantic segmentation. He et al. [25] performs the task of instance-specific semantic segmentation. In robotic vision, several methods [26], [27] predict semantic segmentation for unknown objects. Araki et al. [28] proposed a network for semantic segmentation and grasp detection for a suction cup using multi-task learning with a single deep neural network.

III. PROBLEM STATEMENT

Grasp detection. As proposed by Lenz et al. [8], we use the five-dimensional rectangular representation for robotic grasps. This representation consists of the location and

orientation of a parallel-plate gripper before it closes on the object. A grasp candidate \mathbf{g} is defined as

$$\mathbf{g} = (x, y, w, h, \theta), \quad (1)$$

whereas x and y describe the center of the grasp candidate, w and h describe the width and height, and θ describes the orientation of the rotated box representation.

Semantic Segmentation. Semantic segmentation is the task of assigning a class label to each pixel in the image. Note that for grasp detection of unseen objects, we define the set of semantic classes as $\{\textit{graspable}, \textit{non-graspable}\}$. If additional class information about the objects is available in the dataset, the set of semantic classes can be adjusted accordingly.

IV. PROPOSED METHOD

The proposed architecture is based on [29], a framework originally used for panoptic segmentation, which we adapted for our purpose. This section provides an overview about the specific network parts, which can be seen in Figure 2.

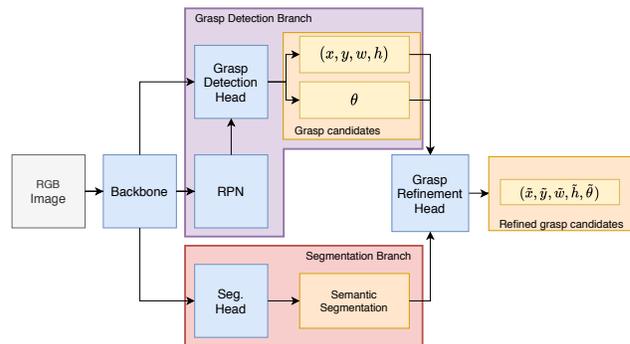


Fig. 2. Architecture of our proposed model. Both branches for grasp detection and segmentation share the backbone network as feature extractor. Both outputs (grasp candidates and semantic segmentation) are used as input for the grasp refinement head, which predicts refined grasp candidates with increased accuracy.

A. Shared Backbone

As feature extractor, we use a slightly modified ResNet-101 [30] with a Feature Pyramid Network (FPN) [31] on top. The modules $\{\text{conv2}, \text{conv3}, \text{conv4}, \text{conv5}\}$ of the ResNet-101 architecture are linked to the FPN. All Batch Normalization + ReLU layers in the original ResNet-101 structure are replaced by synchronized Inplace Activated Batch Normalization (InPlaceABNSync) [32], using LeakyReLU as activation function.

B. Grasp Detection Branch

The grasp detection branch is based on the state-of-the-art Faster R-CNN object detector [33], consisting of a Region Proposal Network (RPN) and a detection stage. This subsection provides details about how we modified this object detector for grasp detection.

1) *Region Proposal Network*: The RPN outputs rectangular region proposals using the features of the backbone network as input. These region proposals are defined as $\hat{\mathbf{r}} = (\hat{x}, \hat{y}, \hat{w}, \hat{h})$ where (\hat{x}, \hat{y}) denote the center of the proposal in pixel coordinates and (\hat{w}, \hat{h}) denote the width and height dimension, respectively. Note that the region proposals are axis-aligned, without information about a possible orientation.

2) *Grasp Detection Head*: The grasp detection head predicts grasp candidates, whereas one candidate \mathbf{g} is defined as described in Equation 1. Each previously calculated region proposal $\hat{\mathbf{r}}$ is used as input for the grasp detection head. Then, ROIAlign [25] is applied to extract feature maps with 14×14 spatial resolution that directly correspond to the region proposals. Afterwards, average pooling with kernel size 2 is applied to each feature map, before they are fed into two fully connected (fc) layers with 1024 neurons each. After each fc layer, an InPlaceABNSync normalization layer and Leaky ReLU activation with 0.01 is applied. The results are forwarded into two sub-networks:

Grasp orientation prediction. The first sub-network consists of a fc layer with 1024 neurons followed by $N_{classes} + 1$ output units, whereas $N_{classes}$ defines the number of orientation classes. Similar to [9], we discretized the grasp orientation θ into $N_{classes} = 18$ intervals with equal length, where each interval is represented by its mean value. The full set of possible orientation classes is defined as $\mathcal{C} = \{(1, \dots, N_{classes})\}$ and the additional class \emptyset to assess the possible invalidity of a proposal. The output units provide logits for a softmax layer that represents the probability distribution over all possible orientation classes. The probability associated to orientation class $c \in \mathcal{C}$ is used as score function s^c .

Bounding box prediction. The second sub-network consists of a fc layer with 1024 neurons followed by $4N_{classes}$ output units. The output units encode class-specific correction factors $(t_x^c, t_y^c, t_w^c, t_h^c)$ for each input proposal $\hat{\mathbf{r}}$. These correction factors, and the orientation information given by the score function s^c , can then directly be used to calculate a grasp candidate \mathbf{g} .

C. Segmentation Branch

The segmentation head takes the output of the first four scales of the FPN as input. Each scale of FPN features are fed to an independent Mini-DeepLab module [29], which makes it possible to capture global structures of the input image with relatively few memory consumption. The output of each Mini-DeepLab module is then up-sampled to $\frac{1}{4}$ of the input image size. Afterwards, all feature maps are concatenated and fed to a 1×1 convolution with $S_{classes}$ output channels, representing the probability distribution over all semantic classes.

D. Grasp Refinement Head

The grasp refinement head (see Figure 3) takes as input the previously predicted grasp candidates from the grasp detection branch and a semantic probability map from the

segmentation branch. First, we fuse both information by cropping the area of the grasp candidate from the probability map. The cropped probability map and the original probability map are then stacked together and used as input for a Multilayer Perceptron (MLP). Note that this approach enables us to combine geometric information of grasp candidates with information about object shape. The MLP is a two-layer fc network with five output neurons. After each fc layer, an InPlaceABNSync normalization layer and Leaky ReLU activation with 0.01 is applied. The output represents refined correction factors $(t_x^g, t_y^g, t_w^g, t_h^g, t_\theta^g)$ for each input grasp candidate \mathbf{g} , which are then used to calculate the refined grasp candidate $\bar{\mathbf{g}}$.

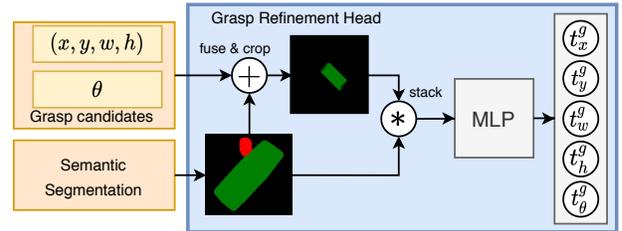


Fig. 3. Structure of the grasp refinement head. The Multilayer Perceptron (MLP) takes as input two stacked semantic probability maps with dimension $(H \times W)$, whereas (H, W) defines the height and width of the probability map. The output units of the MLP represent refined correction factors $(t_x^g, t_y^g, t_w^g, t_h^g, t_\theta^g)$ which can directly be used to calculate the refined grasp candidate. Note that this operation is performed simultaneously for N grasp candidates, leading to an $(N \times 2 \times H \times W)$ dimensional input for the MLP.

V. TRAINING LOSSES

For simultaneously learning the tasks of grasp detection, instance segmentation and grasp candidate refinement, we define the composite loss function \mathcal{L} as

$$\mathcal{L} = \lambda_{grasp} \mathcal{L}_{grasp} + \lambda_{sem} \mathcal{L}_{sem} + \lambda_{refine} \mathcal{L}_{refine}, \quad (2)$$

with the grasp detection loss \mathcal{L}_{grasp} , the semantic segmentation loss \mathcal{L}_{sem} and the grasp refinement loss \mathcal{L}_{refine} . All parts are weighted with a specific hyperparameter λ .

A. Grasp Detection Loss

The grasp detection loss \mathcal{L}_{grasp} is defined as

$$\mathcal{L}_{grasp} = \mathcal{L}_{RPN} + \mathcal{L}_{box} + \mathcal{L}_{rot}, \quad (3)$$

where \mathcal{L}_{RPN} defines the loss for training the RPN, \mathcal{L}_{box} defines the regression loss for the box coordinates and \mathcal{L}_{rot} defines the classification loss for the grasp orientation. We refer to [25] for additional information about \mathcal{L}_{RPN} .

The grasp orientation loss \mathcal{L}_{rot} is defined as

$$\mathcal{L}_{rot} = -\frac{1}{|\mathcal{R}|} \sum_{\hat{\mathbf{r}} \in \mathcal{R}_+} \log s_{\hat{\mathbf{r}}}^{c_r} - \frac{1}{|\mathcal{R}|} \sum_{\hat{\mathbf{r}} \in \mathcal{R}_-} \log s_{\hat{\mathbf{r}}}^{\emptyset}. \quad (4)$$

Note that $\mathcal{R} = \mathcal{R}_+ \cup \mathcal{R}_-$ is the set of valid and invalid region proposals generated using the RPN. The score function $s_{\hat{\mathbf{r}}}^{c_r}$

defines the probability that the region proposal $\hat{\mathbf{r}}$ belongs to the ground truth orientation class c_r , and $s_{\hat{\mathbf{r}}}^0$ defines the probability that the region proposal is invalid. For bounding box regression we use the loss \mathcal{L}_{box} defined as

$$\mathcal{L}_{box} = \sum_{i \in \{x, y, h, w\}} \text{smooth}_{L_1}(t_i^c - t_i^*), \quad (5)$$

with the smooth_{L_1} norm [33] defined as

$$\text{smooth}_{L_1}(z) = \begin{cases} 0.5z^2, & \text{if } |z| < 1 \\ |z| - 0.5, & \text{otherwise.} \end{cases} \quad (6)$$

Note that the parameters t_i^c are the output of the bounding box prediction. The correction factors t_i^* [34] are defined as:

$$\begin{aligned} t_x^* &= (x^* - \hat{x})/\hat{w}, & t_y^* &= (y^* - \hat{y})/\hat{h} \\ t_w^* &= \log(w^*/\hat{w}), & t_h^* &= \log(h^*/\hat{h}), \end{aligned} \quad (7)$$

whereas variables $(\hat{x}, \hat{y}, \hat{w}, \hat{h})$ and (x^*, y^*, w^*, h^*) are parameters from the region proposal $\hat{\mathbf{r}}$ and ground truth grasp candidate respectively.

B. Segmentation Loss

We denote $\mathcal{Y} = \{1, \dots, S_{classes}\}$ as the set of semantic segmentation classes. The per-image segmentation loss is a weighted per-pixel loss [29] defined as

$$\mathcal{L}_{sem} = - \sum_{j,k} w_{j,k} \log P_{j,k}(Y_{j,k}), \quad (8)$$

where (j, k) correspond to the pixel position in the image. Let $Y_{j,k} \in \mathcal{Y}$ be the semantic segmentation ground truth and $P_{j,k}(s)$ the predicted probability for the same pixel to be assigned to a semantic class $s \in \mathcal{Y}$, respectively. The weights $w_{j,k}$ are computed to implement a pixel-wise hard negative mining, which selects the 25% of the lowest predicted probabilities $P_{j,k}$ for all (j, k) using $w_{j,k} = \frac{4}{WH}$, and $w_{j,k} = 0$ otherwise. The spatial resolution of the input image is given as $(W \times H)$.

C. Grasp Refinement Loss

The grasp refinement loss \mathcal{L}_{refine} is defined as

$$\mathcal{L}_{refine} = \sum_{i \in \{x, y, h, w, \theta\}} \text{smooth}_{L_1}(t_i^g - \tilde{t}_i^*). \quad (9)$$

Compared to \mathcal{L}_{box} in Equation 5, the smooth L1 loss is additionally calculated for the orientation parameter θ . Note that the parameters t_i^g are the output of the grasp refinement head. The correction factors \tilde{t}_i^* were adapted based on [35], to take the orientation of the previously predicted grasp candidate \mathbf{g} into account:

$$\begin{aligned} \tilde{t}_x^* &= \frac{1}{w}((x^* - x) \cos \theta + (y^* - y) \sin \theta) \\ \tilde{t}_y^* &= \frac{1}{h}((y^* - y) \cos \theta - (x^* - x) \sin \theta) \\ \tilde{t}_w^* &= \log(w^*/w), & \tilde{t}_h^* &= \log(h^*/h) \\ \tilde{t}_\theta^* &= \frac{1}{\pi}((\theta^* - \theta) \text{ mod } \pi). \end{aligned} \quad (10)$$

Note that variables (x, y, w, h, θ) and $(x^*, y^*, w^*, h^*, \theta^*)$ are parameters from the grasp candidate \mathbf{g} and the ground truth grasp candidate respectively. The *mod* operation is used to adjust the orientation offset \tilde{t}_θ^* to $[0, \pi)$, which corresponds to the interval of possible values for the grasp orientation.

VI. EXPERIMENTS AND RESULTS

We assess the benefits of our proposed method on well-known grasping datasets, namely Cornell [8] and Jacquard [36]. Additionally, we evaluate our approach on our novel OCID [4] extension for robotic grasping. We compare our proposed architecture to state-of-the-art methods which also only use RGB data as input modality, to ensure fair comparison. Due to the fact that our grasp detection and segmentation branches are independent, we evaluate several model configurations throughout our experiments, with the following terminology: Detection (ours) refers to the model using only the grasp detection branch, Det.Seg (ours) consists of both grasp detection and segmentation branch, and Det.Seg.Refine (ours) consists of the full architecture (see Figure 2). We initialize the backbone network with pre-trained ImageNet [37] weights and freeze the parameters of the first two network modules {conv1, conv2} during all training runs. All training and evaluation runs were performed using a single Nvidia GeForce RTX 2080 Ti graphics card. Unless otherwise stated, results from other methods are taken from the corresponding paper.

A. Evaluation Metric

For grasp detection, we use the popular Jaccard index as accuracy measurement. A grasp candidate is reported correct if:

- 1) the angle difference between predicted grasp candidate g_p and ground truth grasp candidate g_{gt} is within 30° and
- 2) the Intersection over Union (IoU) of them is greater than 0.25 which means

$$IoU = \frac{|g_p \cap g_{gt}|}{|g_p \cup g_{gt}|} > 0.25 \quad (11)$$

For semantic segmentation, we report the IoU between predicted and ground truth segmentation.

B. Cornell dataset

The Cornell grasp dataset [8] contains 885 RGB-D images of 240 graspable objects, whereas each image has a dimension of $640 \times 480px$. Correct grasp candidates are hand-annotated using the rectangular grasp representation. Like previous work, we used 5-fold cross-validation and report the mean grasp detection accuracy. We performed an image-wise split into training and test set, whereas image-wise means that images are split randomly without considering which object is in the image. Due to the fact that this dataset does not provide ground truth segmentation data, we are only able to evaluate the model Detection (ours) in this experiment.

Data Preprocessing and Data Augmentation. Because of the relatively small size of this dataset, we used data augmentation to enlarge the dataset during training. Each image

was center cropped to $351 \times 351px$. Next, we applied random rotation between 0° and 360° , and randomly translated the image in x and y direction independently up to $50px$.

Training schedule. The network was trained end-to-end, using a learning rate of 0.04 with weight decay of 0.0001, a momentum factor of 0.9 with enabled nesterov momentum and SGD as optimizer. We used a batch size of 12 during training.

Quantitative Results. Table I shows the results on the Cornell grasp dataset. Note that our implementation achieves state-of-the-art results for accuracy compared to other methods using RGB images as input. Additionally, our model achieves a high FPS rate during inference, making it well suited for real-time applications.

TABLE I. Comparison of grasp accuracy and computational speed for the Cornell dataset. All results are produced with image-wise data split and RGB as input modality using 5-fold cross-validation.

Method	Input	Grasp Accuracy (%)	Speed (FPS)
Zhou [10], Resnet-50	RGB	97.7	9.9
Zhou [10], Resnet-101	RGB	97.7	8.5
Zhang [11], Resnet-101	RGB	93.6	25.2
Park [3], Darknet-19	RGB	97.7	140
Karaoguz [12], VGG-16	RGB	88.7	2
Chu [9], Resnet-50	RGB	94.4	8.3
Kumra [17], GR-ConvNet	RGB	96.6	52
Detection (ours)	RGB	98.2	63

C. Jacquard Dataset

The Jacquard dataset [36] contains 54k synthetic RGB-D images using 11k different objects. As ground truth, the dataset contains automatically generated grasp candidates, as well as ground truth semantic segmentation, which enables us to evaluate our multi-task, grasp detection and segmentation approach. We made an image-wise split of the dataset, using 95% of the data as training set, and the remaining images as test set.

Data Preprocessing. Due to the relatively large image size of $1024 \times 1024px$, we decided to downsample each image by the factor of 2. No data augmentation was needed due to the high number of samples in the dataset.

Training schedule. For this experiments, the networks Detection (ours) and Det_Seg (ours) were trained end-to-end using a learning rate of 0.023 with weight decay of 0.0001, a momentum factor of 0.9 with enabled nesterov momentum and SGD as optimizer. We used a batch size of 8 during training and weighting factors of $\lambda_{grasp} = 1.0$, $\lambda_{sem} = 0.8$ and $\lambda_{refine} = 0.8$ for loss balancing. For training Det_Seg_Refine (ours), we started from the pre-trained version Det_Seg (ours), freezing all backbone parameters while re-training all other parts. We changed the learning rate to 0.008 and the batch size to 4. Although Det_Seg_Refine (ours) could also be trained from scratch, we believe that starting from a well-known baseline makes it easier to demonstrate the effect of the proposed extension and lead to results that are easier to interpret.

Quantitative Results. Our experiments on the Jacquard dataset lead to several findings:

a) *All of our model configurations deliver state-of-the-art performances:* Table II shows the results on the Jacquard grasp dataset according to grasp accuracy, segmentation accuracy and computational speed. All our models achieve higher scores for grasp detection accuracy compared to other state-of-the-art methods using RGB images as input. Furthermore, the inference speed indicates that all our models are fast enough for real-time application.

b) *Multi-task learning is not automatically beneficial:* Focusing on the results of Det_Seg (ours), we see that multi-task learning for grasp detection and segmentation has in our case no beneficial effect on grasp accuracy. However, given the fact that these two tasks are quite different (grasp detection focuses on very small parts of the object whereas segmentation takes the whole scene into account) it is already remarkable that both models perform nearly equally well.

c) *The grasp refinement head significantly increases orientation accuracy:* Table II shows that using Det_Seg_Refine (ours) increases the grasp accuracy by a small margin compared to Det_Seg (ours). Due to the fact that the Jaccard index criteria are quite weak (see Subsection VI-A), Det_Seg (ours) is accurate enough to pass them in most cases. However, if we apply more strict criteria, we can see a significant benefit using our grasp refinement module. Table III shows the results for decreased orientation tolerance. Using our proposed refinement head, we were able to outperform our baseline model by more than 16%, which highlights a significant increase of orientation accuracy (see Table III, column 5°). Table IV shows the results for different IoU criteria. As there is no significant difference for our two models, this experiment indicates that multiple iterations of bounding box regression do not increase the performance (as stated in [34]).

TABLE II. Comparison of grasp accuracy, segmentation IoU and computational speed for the Jacquard dataset. All results are produced with image-wise data split and RGB as input modality.

Method	Grasp Accuracy (%)	Seg. IoU (%)	Speed (FPS)
Zhang[11], ROI-GD	90.4	-	-
Song [13], Resnet-101	91.5	-	-
Kumra [17], GR-ConvNet	91.8	-	-
Depierre [38]	85.7	-	-
Detection (ours)	92.69	-	53
Det_Seg (ours)	92.59	95.12	48
Det_Seg_Refine (ours)	92.95	94.86	31

D. OCID Dataset with Grasp-Extension

OCID [4] contains diverse settings of objects, background, context, sensor to scene distance, viewpoint angle and lighting conditions. The original purpose of this dataset is to evaluate semantic segmentation methods in scenes with increasing amount of complexity. OCID contains RGB-D data with automatically annotated segmentation. For our proposed dataset extension (we denote it as OCID_grasp), we manually

TABLE III. Comparison of grasp accuracy (in [%]) for the Jacquard dataset using different angle thresholds and an IoU threshold of 25%.

Results for the methods Zhou and Depierre are taken from [38].

Method	30 °	25 °	20 °	15 °	10 °	5 °
Zhou [10]	81.95	81.76	81.27	80.23	77.79	-
Depierre [38]	85.74	85.55	85.01	83.65	80.82	-
Detection (ours)	92.69	92.34	92.08	91.40	88.12	56.23
Det_Seg_Refine (ours)	92.95	92.88	92.42	91.52	88.92	72.79

TABLE IV. Grasp accuracy (in [%]) for the Jacquard dataset using different IoU thresholds and an angle threshold of 30°. Results for the methods Zhou and Depierre are taken from [38].

Method	IoU 25%	IoU 30%	IoU 35%
Zhou [10]	81.95	78.26	74.33
Depierre [38]	85.74	82.58	78.71
Song [13]	91.5	89.7	87.3
Detection (ours)	92.69	91.29	88.99
Det_Seg_Refine (ours)	92.95	91.33	88.96

annotated the ARID10 and ARID20 subsets of OCID with valid grasp candidates for each graspable object in the scene. Additionally, we added class information to each object, and assigned each grasp candidate to the corresponding object class. This allows us to evaluate grasp candidates for each object class in the scene. Overall, OCID_grasp consists of 1763 selected images with over 11.4k segmented object masks and more than 75k hand-annotated grasp candidates, and each object is classified into one of 31 classes. Figure 4 (left column) shows sample images of the dataset, indicating the high complexity of the scenes.

Data Preprocessing and Augmentation. We applied random rotation between 0° and 360° , and randomly translated the image in x and y direction independently up to $50px$ as data augmentation during training.

Training schedule. The network was trained end-to-end using a learning rate of 0.03 with weight decay of 0.0001, a momentum factor of 0.9 with enabled nesterov momentum and SGD as optimizer. We used a batch size of 8 during training and weighting factors of $\lambda_{grasp} = 1.0$ and $\lambda_{sem} = 0.75$ for loss balancing.

Evaluation metrics. Due to the fact that the scenes in this dataset contain multiple graspable objects, we decided to calculate the Jaccard index for each graspable object class in the scene, as well as the segmentation IoU for each object class. The Jaccard index at object class level has the additional criteria that the center of the grasp candidate has to be located on the predicted segmentation mask of this specific class. We used 5-fold cross-validation and report the mean grasp detection and segmentation accuracy.

Quantitative & Qualitative Results. We used Det_Seg (ours) as network architecture for this experiment. The network outputs possible grasp candidates for each graspable object in the scene, as well as semantic segmentation for each class. We selected the best grasp candidate for each object class according to the grasp confidence score, with

TABLE V. Results for grasp accuracy and semantic segmentation for the OCID_grasp dataset. For grasp accuracy, the Jaccard index was calculated on object class level, which corresponds to finding the best possible grasp for each class.

Method	Grasp Accuracy (%)	Seg. IoU (%)
Deg_Seg (ours)	89.02	94.05

the constraint that the center of the grasp candidate has to be located on the segmentation mask of this object class. Table V shows the results for grasp accuracy at object class level and mean IOU for instance segmentation. Both, the high grasp accuracy and segmentation IoU, indicate that our proposed model delivers highly accurate results for complex scenes. Qualitative results are shown in Figure 4.



Fig. 4. Qualitative results for OCID_grasp. Explanation of images from left to right: 1) raw input image; 2) predicted semantic segmentation, whereas each color represents a specific class; 3) best possible grasp for each class in the scene (blue lines denote parallel plates of the gripper, red lines denote opening width). If multiple instances of the same class are in the scene, e.g. class orange in the first and third row, only the one grasp candidate with the highest confidence will be shown. Each row refers to a separate input example.

VII. CONCLUSIONS

In this work, we have proposed a novel CNN architecture for jointly producing highly accurate grasp detection and semantic segmentation, using a shared backbone network as feature extractor. We showed that our proposed grasp refinement module can be successfully used to increase the accuracy of previously predicted grasp candidates, especially in terms of grasp orientation. Using our proposed OCID_grasp dataset extension, we showed high accuracy for grasp detection in complex scenes, and demonstrated how semantic segmentation can additionally be used to assign grasp candidates to specific objects. In the future, we plan to extend our approach using multi-modal input data to leverage the full extend of information provided by RGB-D grasp datasets.

REFERENCES

- [1] C. Yang, X. Lan, H. Zhang, and N. Zheng, "Task-oriented grasping in object stacking scenes with crf-based semantic model," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 6427–6434.
- [2] H. Zhang, X. Lan, S. Bai, L. Wan, C. Yang, and N. Zheng, "A multi-task convolutional neural network for autonomous robotic grasping in object stacking scenes," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 6435–6442.
- [3] D. Park, Y. Seo, D. Shin, J. Choi, and S. Y. Chun, "A single multi-task deep neural network with post-processing for object detection with reasoning and robotic grasp detection," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 7300–7306.
- [4] M. Suchi, T. Patten, D. Fischinger, and M. Vincze, "Easylab: A semi-automatic pixel-wise object annotation tool for creating robotic rgb-d datasets," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2019, pp. 6678–6684.
- [5] A. Bicchi and V. Kumar, "Robotic grasping and contact: A review," in *IEEE International Conference on Robotics and Automation (ICRA)*, vol. 1, 2000, pp. 348–353.
- [6] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-driven grasp synthesis—a survey," *IEEE Transactions on Robotics*, pp. 289–309, 2013.
- [7] J. Redmon and A. Angelova, "Real-time grasp detection using convolutional neural networks," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 1316–1322.
- [8] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," *The International Journal of Robotics Research*, vol. 34, no. 4-5, pp. 705–724, 2015.
- [9] F.-J. Chu, R. Xu, and P. A. Vela, "Real-world multiobject, multigrasp detection," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3355–3362, 2018.
- [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/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 7223–7230.
- [11] H. Zhang, X. Lan, S. Bai, X. Zhou, Z. Tian, and N. Zheng, "Roi-based robotic grasp detection for object overlapping scenes," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 4768–4775.
- [12] H. Karaoguz and P. Jensfelt, "Object detection approach for robot grasp detection," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2019, pp. 4953–4959.
- [13] Y. Song, L. Gao, X. Li, and W. Shen, "A novel robotic grasp detection method based on region proposal networks," *Robotics and Computer-Integrated Manufacturing*, vol. 65, p. 101963, 2020.
- [14] D. Park, Y. Seo, and S. Y. Chun, "Real-time, highly accurate robotic grasp detection using fully convolutional neural network with rotation ensemble module," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 9397–9403.
- [15] N. Gkanatsios, G. Chalvatzaki, P. Maragos, and J. Peters, "Orientation attentive robot grasp synthesis," *arXiv preprint arXiv:2006.05123*, 2020.
- [16] D. Morrison, P. Corke, and J. Leitner, "Closing the Loop for Robotic Grasping: A Real-time, Generative Grasp Synthesis Approach," in *Proc. of Robotics: Science and Systems (RSS)*, 2018.
- [17] S. Kumra, S. Joshi, and F. Sahin, "Antipodal robotic grasping using generative residual convolutional neural network," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020.
- [18] U. Asif, J. Tang, and S. Harrer, "EnsembleNet: Improving grasp detection using an ensemble of convolutional neural networks," in *The British Machine Vision Conference (BMVC)*, 2018, p. 10.
- [19] —, "Graspnet: An efficient convolutional neural network for real-time grasp detection for low-powered devices," in *International Joint Conferences on Artificial Intelligence (IJCAI)*, 2018, pp. 4875–4882.
- [20] —, "Densely supervised grasp detector (dsgd)," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 8085–8093.
- [21] 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 (ICRA)*, 2016, pp. 3406–3413.
- [22] 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," *The International Journal of Robotics Research*, vol. 37, no. 4-5, pp. 421–436, 2018.
- [23] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [24] L.-C. 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," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 4, pp. 834–848, 2017.
- [25] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 2961–2969.
- [26] M. Danielczuk, M. Matl, S. Gupta, A. Li, A. Lee, J. Mahler, and K. Goldberg, "Segmenting unknown 3d objects from real depth images using mask r-cnn trained on synthetic data," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2019, pp. 7283–7290.
- [27] C. Xie, Y. Xiang, A. Mousavian, and D. Fox, "The best of both modes: Separately leveraging rgb and depth for unseen object instance segmentation," in *Conference on Robot Learning*. PMLR, 2020, pp. 1369–1378.
- [28] R. Araki, T. Onishi, T. Hirakawa, T. Yamashita, and H. Fujiyoshi, "Mtdssd: Deconvolutional single shot detector using multi task learning for object detection, segmentation, and grasping detection," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 10487–10493.
- [29] L. Porzi, S. R. Bulo, A. Colovic, and P. Kotschieder, "Seamless scene segmentation," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 8277–8286.
- [30] K. He, X. Zhang, S. Ren, and J. Sun, "Identity mappings in deep residual networks," in *European Conference on Computer Vision (ECCV)*, 2016, pp. 630–645.
- [31] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 2117–2125.
- [32] S. Rota Bulò, L. Porzi, and P. Kotschieder, "In-place activated batchnorm for memory-optimized training of dnns," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 5639–5647.
- [33] S. Ren, K. He, R. 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.
- [34] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 580–587.
- [35] J. Ding, N. Xue, Y. Long, G.-S. Xia, and Q. Lu, "Learning roi transformer for oriented object detection in aerial images," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [36] A. Depierre, E. Dellandréa, and L. Chen, "Jacquard: A large scale dataset for robotic grasp detection," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 3511–3516.
- [37] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al., "Imagenet large scale visual recognition challenge," *International Journal of Computer Vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [38] A. Depierre, E. Dellandréa, and L. Chen, "Optimizing correlated graspability score and grasp regression for better grasp prediction," *arXiv preprint arXiv:2002.00872*, 2020.