

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

Effective deployment of CNNs for 3DOF pose estimation and grasping in industrial settings

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Effective deployment of CNNs for 3DOF pose estimation and grasping in industrial settings / de Gregorio D.; Zanella R.; Palli G.; Di Stefano L. - STAMPA. - (2021), pp. 9411912.7419-9411912.7426. (Intervento presentato al convegno 25th International Conference on Pattern Recognition, ICPR 2020 tenutosi a ita (ELECTR NETWORK) nel 2021) [10.1109/ICPR48806.2021.9411912].

Availability:

This version is available at: https://hdl.handle.net/11585/841335 since: 2021-12-10

Published:

DOI: http://doi.org/10.1109/ICPR48806.2021.9411912

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (https://cris.unibo.it/). When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

D. De Gregorio, R. Zanella, G. Palli and L. Di Stefano, "Effective Deployment of CNNs for 3DoF Pose Estimation and Grasping in Industrial Settings," 2020 25th International Conference on Pattern Recognition (ICPR), 2021.

The final published version is available online at: <u>https://dx.doi.org/10.1109/ICPR48806.2021.9411912</u>

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<u>https://cris.unibo.it/</u>)

When citing, please refer to the published version.

Effective Deployment of CNNs for 3DoF Pose Estimation and Grasping in Industrial Settings

Daniele De Gregorio EYECAN.ai Bologna, Italy daniele.degregorio@eyecan.ai Riccardo Zanella DEI - University of Bologna Bologna, Italy riccardo.zanella2@unibo.it Gianluca Palli DEI - University of Bologna Bologna, Italy gianluca.palli@unibo.it Luigi Di Stefano DISI - University of Bologna Bologna, Italy luigi.distefano@unibo.it

Abstract—In this paper we investigate how to effectively deploy deep learning in practical industrial settings, such as robotic grasping applications. When a deep-learning based solution is proposed, usually lacks of any simple method to generate the training data. In the industrial field, where automation is the main goal, not bridging this gap is one of the main reasons why deep learning is not as widespread as it is in the academic world. For this reason, in this work we developed a system composed by a 3-DoF Pose Estimator based on Convolutional Neural Networks (CNNs) and an effective procedure to gather massive amounts of training images in the field with minimal human intervention. By automating the labeling stage, we also obtain very robust systems suitable for production-level usage. An open source implementation of our solution is provided, alongside with the dataset used for the experimental evaluation.

I. INTRODUCTION

Among visual perception tasks, 2D Object Detection with category-level classification has achieved an effectiveness thanks to Convolutional Neural Networks (CNNs) [1],[2],[3]. Unfortunately, the more complex perception task of the 3D Pose estimation has not experienced the same strengthening, notwithstanding remarkable results [4],[5],[6] have endorsed CNNs also for this - more complex - task. We argue that between the key reasons for this state-of-affairs is the lack of training data: while for 2D Object Detection huge datasets like Pascal VOC [7] or COCO [8] define a reliable testbed for the community, the same can not be said for the 3D counterpart (with the exception of some small datasets e.g. like [9],[10]). Obviously, for relevant robotic applications, such as, for instance, a fully automatic pick-and-place, Pose Estimation is an essential stage of the overall pipeline, and, as stated before, the availability of training data hinders deployment of CNNs. Thus, the claim of our work is that for a real industrial application it is not sufficient to develop advanced data-driven models, like a convolutional neural networks, but - simultaneously - the data sourcing problem should be addressed. Thus, we propose to tackle the object detection and 3-DoF pose estimation task by an integrated framework based on CNNs wherein the required labeled training data are sourced autonomously, i.e. with negligible human intervention. We also show how the proposed framework enables development of a fully automated robotic grasping system.

The basic idea of our approach is to take advantage of the known difficulty of CNNs to learn rotational invariant image features. As shown in [11],[12] these networks redundantly learn multiple representation of sought objects when they exhibit multiple rotation in the training images. As opposed to approaches like [13], [11], which try to learn rotation-invariant representation, we leverage on classical CNN-based Object Detectors to formulate the angle estimation task as a classification problem by leading the network to interpret each single object orientation as a stand-alone class (for this reason we name the algorithm LOOP: Leveraging on a generic Object detector for Orientation Prediction). As previously mentioned, we endow our approach with an automated dataset generation technique that allows to label an entire video sequence easily, provided that the sequence features the same type of image acquisitions conditions as those in which the detector is exptected to operate at test time (e.g. a quite-planar scene in front of the camera). This approach enables collections of massive amounts of training data which, in turn, allow the creation of an almost perfect object detector (i.e. $\sim 0.99 \text{ mAP}$ according to our experiments).

Thus, the key contributions of our paper are:

- A novel approach to extend a generic CNN-based 2D Object Detector in order to predict oriented bounding boxes.
- A fast and reliable Labeling pipeline that allows to gather a labeled Dataset to train the above mentioned extended Object Detector with minimal human intervention (i.e. the user has to manually label only the first frame of a video sequence).
- A real proof-of-concept pick&place robotic application based on this approach.

An open source implementation of the proposed method is available online 1 .

II. RELATED WORKS

Nowadays, object detection and 3-DoF Pose Estimation in industrial settings is mainly addressed by classical computer vision approaches based on hand-crafted 2D features, which

This work was supported by the European Union's Horizon 2020 research and innovation programme under grant agreement No. 870133 as part of the RIA project REMODEL (Robotic tEchnologies for the Manipulation of cOmplex Deformable Linear objects).

¹ https://github.com/m4nh/loop



Fig. 1. The overall pipeline of LOOP. The starting point is the creation of a dataset of oriented bounding boxes (oriented labels) with minimal human intervention. The *oriented* labels are converted in *unoriented* labels, by means of the f_{o2u} function, suitable to train a classical object detector, which encodes objects orientation in the classification process. Hence the detector infers *unoriented* predictions (y_i) , with the same orientation encodings. The *unoriented* predictions are then transformed in *oriented* predictions (\tilde{y}_i) by means of the f_{u2o} function.

can effectively represent the orientation of salient local image structures. SIFT [14] is one of the most popular 2D feature detector and descriptor by which it is possible to implement a full Object Detection pipeline for textured objects. By matching multiple 2D local features it is possible to estimate the homography (or even a rigid transform) between a model image and the target scene. SIFT can be replaced by other popular alternatives, like SURF, KAZE, ORB, BRISK etc.. [15] presents a comprehensive evaluation of the main algorithms for 2D feature matching. As stated before, the aforementioned methods are suitable - mostly - for textured objects detection, but the same matching pipeline can be deployed replacing them with detectors/descriptor based on geometrical primitives (e.g. oriented segments) amenable to texture-less objects. One of the leading texture-less object detectors is BOLD [16], which was then followed by BORDER [17] (and its extension, referred to as BIND [18]). A popular alternative to features, instead, is Rotation-Invariant Template Matching like OST [19], OCM [20] or Line2D [21]. However, a tamplate-based approach may not be the most efficient solution for real-time applications.

The above-mentioned considerations have lead us to compare LOOP mainly with SIFT [14] and BOLD [16], undoubtedly two state-of-the-art approaches in textured and texture-less object detection, respectively. Our claim is to propose a real-time deep learning alternative able to cope with both textured as well as untextured models and, seamlessly, over plain and cluttered backgrounds. As stated in section I, as we can exploit any generic CNN Object Detector, we investigated here about well established approaches like YOLO [1] and SSD[2], which resolve the 2D object detection and recognition task with a single image analysis pass, as well as Faster R-CNN[3], which instead conceptually splits the detection and recognition parts. We found that YOLOv3 [22], the newest declension of the classical YOLO algorithm, represents a satisfactory test case for our experiments, though is worth pointing out that LOOP is detector-agnostic.

Another research line to attain robotic grasping systems concerns direct estimation of grasping points from images by



Fig. 2. Graphic representation of the conversion between *Oriented-to-Unoriented* bounding boxes. The real angle θ is converted in the index C_{θ} of the corresponding quantized bin among the $k = \lceil \frac{2\pi}{\alpha} \rceil$ possible bins.



Fig. 3. A graphical representation of the conversion between *Unoriented-to-Oriented* bounding boxes. It is important to know the *ratio* of the sought object in order to produce a commesurate oriented bounding box.

means of CNNs. One of the most used approach, conceptually similar to our method, is the 2D Rectangular Representation of grasp, as described in [23]. The authors demonstrated that a 2D representation of grasp is enough to perform a 3D manipulation with a robotic arm. Recent works like [24] or [25] estimate the position and the orientation of these 2D Rectangular Representation of grasp by means of a CNN, as either a regression or classificatiton problem, respectively. However, we believe that the full 2D Oriented Bounding Box of the sought objects yielded by our approach is a better representation for grasp in planar setting, because it allows to perform both obstacle avoidance as well as model-based grasp points computation.

III. HOW DOES LOOP WORK?

As illustrated in Figure 1, given an RGB image, the LOOP framework can produce a set of predictions $\breve{y}_i = \{\breve{b}_i, \theta_i, c_i\},\$ where $\check{b}_i = \{x, y, w, h\} \in \mathbb{R}^4$ represents the coordinates of the Oriented Bounding Box (OBB in short) clockwise-rotated by an angle of θ_i and $c_i \in \mathbb{Z}^+$ is the object class. As already mentioned, we leverage on a classical object detector, which outputs a set of simpler predictions $y_i = \{b_i, \hat{c}_i\}$ where $b_i = \{x, y, W, H\} \in \mathbb{R}^4$ represents the coordinates of the *unoriented* Bounding Box (BB in short) and $\hat{c}_i \in \mathbb{Z}^+$ encodes, with our formulation, both the object class as well as orientation information. In subsection III-A we will explain how to transform an oriented prediction \breve{y}_i into an un-oriented one, y_i , while in subsection III-B we will describe the inverse procedure. Finally, in subsection III-C we will explain how to generate Oriented Bounding Box labels for an entire video sequence by labeling just the first frame.

A. f_{o2u} : The Oriented-to-Unoriented Function

As stated before, our approach is an extension of a classical 2D Object Detector to make it capable of estimating also the orientation of a target object. We formulate the angle estimation problem as a classification task by simply quantizing the angular range into k bins and by expanding all the C categories, managed by the object detector, into C' = kC new classes. Thus, as shown in Figure 2, for each object instance, we can compute its real-valued angle with $\theta = \arccos(\mathbf{u}_x \cdot \mathbf{v}_x)$, as the angle between the unit vector \mathbf{v}_x , directed as the first edge of the corresponding OOB, and the x axis of the image \mathbf{u}_x . Obviously the theta angle so defined, only for illustrative intent, is limited to the range $[0, \pi]$, for this reason it is necessary to calculate it using

$$\theta = \begin{cases} \operatorname{atan2}(v_{x_y}, v_{x_x}), & \text{if } v_{x_y} >= 0\\ 2\pi + \operatorname{atan2}(v_{x_y}, v_{x_x}), & \text{if } v_{x_y} < 0 \end{cases}$$
(1)

where (v_{x_y}, v_{x_x}) are the components of the \mathbf{v}_x vector. The real-valued angle θ is then converted into the corresponding bin index C_{θ} , with $C_{\theta} = \lfloor \frac{\theta}{\hat{\theta}} \rceil \in \{0, ..., k\}$, where $\hat{\theta}$ is *quantization step* so that $k = \lceil \frac{2\pi}{\hat{\theta}} \rceil$. In order to build an unique formulation to obtain the final converted class we can write:

$$f_{o2u}(c_i, \theta, \hat{\theta}) = c_i k + C_{\theta} = c_i k + \lfloor \theta / \hat{\theta} \rceil = \hat{c}_i$$
(2)

where f_{o2u} (i.e. *o2u*: *Oriented-to-Unoriented*) is the function used to convert the original object class c_i in the expanded class \hat{c}_i which encodes not only the object type but also its quantized orientation. The corresponding BB is computed simply by applying the minimum bounding box algorithm to the 4 vertices of the original OOB.

B. f_{u2o} : The Unoriented-to-Oriented Function

Assuming $y_i = \{b_i, \hat{c}_i\}$ the generic prediction of the Object Detector, where \hat{c}_i is the predicted class (built with the Equation 2) and b_i the corresponding un-oriented bounding box, the purpose of the f_{u2o} function, as depicted in Figure 1, is to produce an equivalent oriented prediction $\breve{y}_i = \{\breve{b}_i, \theta_i, c_i\}$ where c_i is the original class, mapping the object of belonging only, and \check{b}_i is an oriented bounding box (when omitted in images, the angle θ_i is represented, for simplification, with the red arrow oriented as the longest axis of an OBB). This procedure can be thought as the inverse of that described in subsection III-A. The function to determine the original class, c_i , and the predicted angle θ_i is pretty simple:

$$f_{u2o}(\hat{c}_i, \hat{\theta}) = \begin{cases} c_i = \lfloor \frac{\hat{c}_i}{k} \rfloor \\ \theta_i = \hat{\theta}_i \cdot (\hat{c} \mod k) \end{cases}$$
(3)

where $\hat{\theta}$ is the same discretization step as used in the f_{o2u} counterpart. Conversely, the estimation of the OBB (b_i) given the simple BB (b_i) and the corresponding angle θ_i is not a trivial problem because of the infinite number of solutions with no constraints. However, for each object in the dataset we can compute the *ratio* r of its bounding box in a nominal condition (e.g. when it exhibit 0° in an image) and use it as a constraint to reduce the complexity of the procedure. If we define with $P_{\check{b}_i} = \{\mathbf{p}_0, \mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3\}$ the set of OBB's 4 expressed in the reference frame F_{b_i} centered in the bounding box b_i (the clockwise order, as depicted in the last frame of Figure 3, has to be reliable in order to have that $\mathbf{v}_x = \frac{\mathbf{p}_1 - \mathbf{p}_0}{||\mathbf{p}_1 - \mathbf{p}_0||}$), the ratio can be computed easily as $r = \frac{\mathbf{p}_1 - \mathbf{p}_0}{\mathbf{p}_3 - \mathbf{p}_0}$. Thanks to the object aspect ratio we can execute the pipeline depicted in Figure 3 to build an OBB: starting from the original BB and an angle θ , we superimpose a small version of the corresponding OBB, b_{sm_i} (built only using the aspect ratio information), in the center of the BB rotating it by the provided angle; we estimate as $d = \left|\frac{W}{2} + p_{lm_x}\right|$ the distance between the leftmost vertex of the OBB, p_{lm} , and the left edge of the original BB; we enlarge b_{sm_i} by a scaling factor s, in order to have d = 0. Therefore, taking the example of Figure 3, if $\mathbf{p}_{lm} \coloneqq \mathbf{p}_3$, we can simplify the above mentioned distance condition, obtaining the new desired x coordinate of \mathbf{p}_3 as $\hat{p}_{3_x} = -\frac{W}{2}$. And then reformulating it as a scaling problem $\hat{p}_{3_x} = s \cdot p_{3_x}$ we get that $s = -\frac{W}{2p_{3x}}$. The factor s can be used to generate a scaling matrix and transform all 4 vertices consistently.

As mentioned above, due to the error introduced by the Object Detector in the estimation of b_i instances, the construction of a BB from an OBB is not perfectly invertible, so the algorithm just proposed tries to minimize one of the many possible constraints (the proximity of the *left-most* point). Surely an optimization algorithm could take into account more than one constraint (e.g. the proximity of all 4 vertices) but in our case this approach is sufficiently precise and fast to test the rest of the pipeline.

C. Automatic Dataset Generation

In section I we underlined that it is important to provide a smart solution to collect training data for data-driven models to be deployed in real industrial applications. For this reason, we endowed LOOP with an automated labeling tool based on video sequences. The hypotheses allowing the tool to work well are:

1) the video sequence frames a tabletop scene with the image plane as parallel as possible to the supporting surface;



Fig. 4. For each pair of consecutive frames I_i , I_{i+1} it is possible to estimate the rigid transformation **A** such that $I_{i+1} = \mathbf{A}I_i$, by means of a 2D feature matching pipeline, like e.g. ORB[26]. The same transformation **A** can be applied to each OOB label.

- camera movements have to be, as much as possible, only rotational (*Rotate* around the camera z axis) and translational (*Lift* along the z axis);
- 3) the height of target objects has to be mostly uniform.

Figure 4 exemplifies the above requirements by showing two consecutive frames of a suitable video sequence. The figure shows clearly how restricted camera movements (i.e. lift and rotate) leads to a controlled rigid transformation **A** between the two consecutive images I_i , I_{i+1} such that $I_{i+1} = \mathbf{A}I_i$. The same rigid transformation A can be applied to each OOB \check{b}_i present in the image I_i so as to obtain a new set of OOB such that $\check{b}_{i+1} = \mathbf{A}\check{b}_i$. This procedure can be repeated for each consecutive pair of images in the video sequence, it is therefore clear how the sole human intervention is to create the OOB labels in the first frame I_0 . To estimate the rigid transformation **A** any 2D feature-based matching pipeline, described in section II, may be used. We decided to use ORB [26], a patent-free solution, to make our software completely open-source¹.

IV. RESULTS

In this section we will describe both data and models used in our experiments in order to maximize reproducibility. In subsection IV-A we will introduce a novel dataset, dubbed the *LOOP Dataset*, used in our quantitative and qualitative experiments. In subsection IV-B we will describe how the core object detector has been trained over the LOOP Dataset in order to obtain different declensions of the more complex LOOP Detector. To prove the effectiveness of our solution, in subsection IV-C we analyze the absolute LOOP's performances, while in subsection IV-D we present a detailed comparison between LOOP and the state-of-the-art approaches based on handcrafted features. Finally, in subsection IV-E, a qualitative evaluation of our approach within a typical robotic task is provided.

A. Experimental setup: LOOP Dataset

Thanks to the technique examined in subsection III-C, we created our an experimental –public– dataset with 12 objects equally divided into *Textured* and *Untextured* (as shown in Figure 5), submerged in several challenging scenes, in order to extensively compare our approach with classical ones. We

collected 15 tabletop scenes, with randomly arranged objects, featuring different backgrounds: 3 scenes with homogeneous background; 3 scenes with wood; 3 scenes with black background and 5 scenes with an high-clutter background (several prints of Pollock's painting). This increasing variability is used in order to achieve Domain Randomization [27] during training and very challenging scenes during the test. These 4 different scene categories are shown in Figure 6, where they are intentionally presented in increasing order of complexity. We collected a total of 7155 self-labeled images (strictly speaking, we manually labeled only 15 images, one per scene). Moreover, by using the ground truth coming from the whole dataset we also built a Synthetic version of it just by stitching one version of each model with the same arrangement proposed by the original dataset (right column in Figure 6 illustrates synthetic version of the real images in the left column). This version of the dataset is built in order to reproduce a version of the training data in which the distribution of the objects in terms of position and orientation is identical to the real one (also the backgrounds are intentionally similarly synthetized) but the variance of objects semblance, in the images space, is purposely kept low: i.e. stitching the same version of the object synthetically rotated is, under our hypothesis, less effective than produce real rotated viewpoints, especially dealing with deep neural models, not to mention that even the variations of light conditions and perspective are not correctly captured by a synthetic dataset generated this way.

B. Deep Object Detector

As reference for our benchmarks we used YOLOv3 [22], a state-of-the-art Object Detector based on CNNs. We fine-tuned the YOLOv3 model, pretrained on ImageNet [28], with the LOOP approach using 13 scenes of the LOOP Dataset (about 6200 labeled images) with a 80%/20% split for training and test. The training strategy is to freeze weights of the network's feature extractor (namely, the backend), for 2 epochs training only the object detection layers (namely, the frontend), then fine-tuning the whole architecture for 50 further epochs, with a learning rate of 0.001. The two remaining scenes of the dataset (i.e. one with wood background, thought as a simple testbench, and one with high clutter background thought as a complex one), never used during the training phase are used to evaluate the performance of the whole system. We will call these two scenes Simple Scene (477 images) and Hard Scene (477 images) in the rest of this section.

We trained several models using different $\hat{\theta}$ (i.e. angle discretization parameter, as described in subsection III-A), thereby producing different declensions of the detector useful in understanding how the discretization factor goes to affect performances. For the sake of convenience, we will adopt the short nickname LOOP_{α} to identify a LOOP model trained using the discretization angle $\hat{\theta} = \alpha$ (e.g. LOOP₁₀ identifies a model trained by quantizing the whole *angle turn* in 36 bins of 10° each). Moreover, we add the letter *S* to the above mentioned notation (e.g. LOOP^{*S*}_{α}) to represent the same model trained with the synthetic version of the LOOP Dataset.



Fig. 5. LOOP Dataset objects grouped in Textured and Untextured.



Fig. 6. In the left column: 4 samples coming from the LOOP Dataset one for each background category. The right column shows the synthetic version of each sample.

C. LOOP performances

In the claim of this work we said that with the LOOP approach it is possible to build effortlessly an Object Detector, based on CNNs suitable to real industrial applications. In this section, indeed, we evaluate LOOP over the LOOP Dataset to understand if its performances are good enough to embody it into a reliable industrial system.

Figure 7 groups together several Precision/Recall curves obtained by varying the threshold over the confidence output of the YOLOv3 model. Figure 7 (a), (b) and (c) depict the precision/recall of several detectors (i.e. each of which trained with a different angle discretization $\hat{\theta}$) over the *Simple Scene*, *Hard Scene* and both, respectively. In the same plots also the LOOP^S₁₀ model, thus trained only over synthetic data, is depicted in order to assess its performance compared to its counterparty trained on real data (i.e. LOOP₁₀). Table I resumes the *mean average precision* for the considered models: in this concise summary it is even clearer how deploying a synthetic dataset is significantly less effective than leveraging on real imagery, especially when it comes to very complex real situations.

From these results we can conclude that LOOP_{10} is the best version of our detector. The LOOP_5 and LOOP_{20} versions resulted slightly worse. This analysis shows, clearly, how the angle discretization parameter $\hat{\theta}$ affects performances in both direction: too small a value may cause a dramatic increase of detector categories C' (e.g. with $\hat{\theta} = 5$ we have $C' = \frac{360}{5}12 = 864$); on the other hand, if $\hat{\theta}$ is too high, the angle prediction is subject to a large discretization error producing low IoU scores. Moreover, training on synthetic data only (LOOP_{10}^S), although acceptable, is worse than training on real data. With regard to single objects precision/recall, as

Model	Simple	Hard	Overall							
LOOP ₅	0.96	0.96	0.96							
LOOP ₁₀	0.99	0.98	0.99							
$LOOP_{10}^S$	0.95	0.86	0.92							
$LOOP_{20}^{10}$	0.95	0.97	0.96							
LOOP ₃₀	0.90	0.88	0.89							
$LOOP_{45}$	0.69	0.70	0.69							
TABLE I										

The mean average precision (*mAP*) computed for each model, across all objects, for the *Simple Scene*, *Hard Scene* and both.

will be described in more detail later, symmetric objects are – unsurprisingly – somewhat confusing the algorithm.

D. Hand-crafted features vs Deep Learning

We used our LOOP Dataset to test the LOOP approach compared with the state-of-the-art algorithms based on handcrafted features. The first competitor is, certainly, SIFT [14], one of the most used feature-based approach for Textured objects detection. For the Untextured counterpart, we chose BOLD [16], which uses highly repeatable geometric primitives as 2D features in order to perform object detection also with monochrome targets. Both methods are implemented in the same object detection pipeline as described in [16]: 1) keypoints detection; 2) key-point description; 3) features correspondences validated through Generalized Hough Transform (as described in [14])and 4) Pose Estimation. The Pose Estimation in this case is modified in order to obtain a Rigid Transform instead of a Full Affine (homography). In this way we are sure that both detectors will yield predictions similar to $\check{y}_i = \{\check{b}_i, \theta_i, c_i\}$ without producing distorted bounding boxes, as it is the case of homographies.

We compare the output of the previous two pipelines with the output directly obtained with the LOOP approach. We measured the performances with the classical precision/recall metric computed taking into account the Intersection Over Union (IoU) of the predicted OBB: each detection, of a generic algorithm, is counted as *True Positive* if its IoU is > 0.5compared with the corresponding ground truth OBB. With this metric we can analyse: Precision, Recall, FScore and Average IoU for each algorithm. We introduce also an additional term called Oriented Intersection Over Union (OIoU) which measures the classical IoU commesurate to the effectiveness of the algorithm in predicting the correct angle. Formally $OIoU = IoU \times max \left(\frac{\mathbf{v}_x \cdot \hat{\mathbf{v}}_x}{\|\|\mathbf{v}_x\|\| \|\hat{\mathbf{v}}_x\|}, 0 \right)$, so the Oriented IoU is inversely proportional to the angle between the two \mathbf{v}_x of the ground truth and the predicted OBB; it is forced to be between 0 and 1, with the max operator, in order to discard a priori opposed angles. The OIoU term is introduced to measure the



Fig. 7. Precision/Recall curves for the LOOP approach. These plots have been produces by varying the threshold of the predictions *confidence* between 0.0 and 1.0 with a step of 0.05. The first column compares various discretization angles. The second column shows performances of the best version, $LOOP_{10}$, for each single object in the dataset.

capability of each algorithm in distinguishing quite symmetric objects: in the LOOP Dataset, for instance, the two objects *artifact_orange* and *box_brown* seems very symmetrical even thought they are not. In such cases, if the angle prediction is completely wrong, the IoU may be high but the OIoU is very low.

Under the assumption that, as deduced in subsection IV-C, the model with the best trade-off between precision and recall is $LOOP_{10}$, we used the latter to compete with the state-ofthe-art. In Table II and Table III several comprehensive results are shown. The first table compares SIFT and BOLD with $LOOP_{10}$ and $LOOP_{10}^S$ dealing with the *Simple Scene*, i.e. the scene with a low clutter background. The second table, instead, deals with the Hard Scene containing an highly cluttered background. Both tables contain performances dealing with each object, a summary for Textured and Untextured objects and an Overall index. As vouched by these experimental results, LOOP outperforms both SIFT and BOLD in both scenarios: our approach is able to cope with both Textured and Untextured objects seamlessly and is very robust to the high level of clutter in the Hard Scene. On the other hand, as pointed out in Table II, the OIoU index shows a slight fall of LOOP when dealing with symmetric objects (e.g. the box_brown OIoU is 0). This conceptual problem can be easily resolved treating symmetrical objects differently during the angle discretization

process, described in subsection III-A, with a formulation similar to the one introduced in [6]. A qualitative evaluation is present in Figure 9 featuring a real output of the LOOP detector on two samples randomly picked between the Simple and the Hard scenes subsets.

E. Real Robotic application

As a qualitative evaluation of our approach we designed a proof-of-concept pick&place robotic application based on the outcome of the LOOP detector. The setup consist in an industrial robotic arm (COMAU smart six) with a parallel gripper as end effector and an eye-on-hand camera on board. The camera mounted in this way can be thought as the secondary end-effector and than can be arbitrarily moved with high precision. The image plane of the camera is kept parallel to a table-top scene with randomly arranged objects belonging to the LOOP Dataset (same camera-table configuration seen in subsection III-C). The distance between camera and the table plane is known. Given that the output of the pipeline is a set of oriented prediction $\check{y}_i = \{\check{b}_i = \{x_i, y_i, w_i, h_i\}, \theta_i, c_i\}$ we exploit their position (x_i, y_i) and orientation (θ_i) to build a simple control scheme, for robot guidance, in such a way as to move the camera in order to align a target object with the center of camera viewpoint, oriented as the canonical x axis of the image. A proportional-only control scheme is used, thus, to minimize e_t and e_{θ} , the translational and

Index	Precision					Recall				FScore							OIOU			
Method	Sift	Bold	$LOOP_{10}$	$LOOP_{10}^S$	Sift	Bold	$LOOP_{10}$	$LOOP_{10}^S$	Sift	Bold	$LOOP_{10}$	$LOOP_{10}^S$	Sift	Bold	$LOOP_{10}$	$LOOP_{10}^S$	Sift	Bold	$LOOP_{10}$	$LOOP_{10}^S$
artifact_black	0.50	0.84	0.99	0.97	0.00	0.81	0.96	0.98	0.00	0.82	0.97	0.97	0.75	0.77	0.85	0.84	0.75	0.75	0.85	0.84
artifact_metal	0.00	0.00	0.97	0.95	0.00	0.00	0.98	0.86	0.00	0.00	0.97	0.90	0.00	0.00	0.77	0.77	0.00	0.00	0.77	0.77
artifact_orange	0.54	0.96	0.91	0.78	0.12	0.96	0.91	0.88	0.20	0.96	0.91	0.83	0.71	0.86	0.86	0.86	0.62	0.43	0.78	0.37
artifact_white	0.62	0.00	0.97	0.90	0.85	0.00	1.00	0.98	0.72	0.00	0.98	0.94	0.78	0.00	0.84	0.80	0.78	0.00	0.84	0.79
clip	0.78	0.60	0.95	0.91	0.63	0.60	0.93	0.92	0.70	0.60	0.94	0.91	0.76	0.76	0.81	0.75	0.76	0.74	0.81	0.74
screwdriver	0.76	0.58	0.95	0.97	0.05	0.58	0.85	0.89	0.10	0.58	0.89	0.93	0.71	0.73	0.81	0.79	0.71	0.72	0.81	0.79
Untextured	0.53	0.50	0.95	0.91	0.28	0.49	0.94	0.92	0.29	0.49	0.95	0.91	0.62	0.52	0.82	0.80	0.60	0.44	0.81	0.72
battery_black	0.48	0.82	0.98	0.94	0.81	0.82	1.00	1.00	0.60	0.82	0.99	0.97	0.76	0.71	0.84	0.83	0.75	0.68	0.84	0.83
battery_green	0.65	0.44	0.93	0.90	0.83	0.43	0.93	0.91	0.73	0.43	0.93	0.90	0.69	0.70	0.85	0.85	0.69	0.59	0.82	0.79
box_brown	0.65	0.68	0.97	0.87	0.77	0.68	0.92	0.95	0.71	0.68	0.95	0.91	0.80	0.75	0.74	0.73	0.80	0.65	0.00	0.02
box_yellow	0.49	1.00	0.98	0.98	1.00	1.00	0.98	0.59	0.66	1.00	0.98	0.73	0.82	0.81	0.86	0.80	0.82	0.81	0.86	0.79
glue	0.37	0.99	0.97	0.95	1.00	0.99	0.91	0.97	0.54	0.99	0.94	0.96	0.80	0.81	0.82	0.83	0.80	0.80	0.82	0.83
pendrive	0.78	0.91	0.98	0.92	0.36	0.91	0.94	0.85	0.49	0.91	0.96	0.89	0.75	0.74	0.76	0.73	0.74	0.73	0.76	0.71
Textured	0.58	0.79	0.97	0.92	0.75	0.78	0.94	0.86	0.61	0.78	0.95	0.88	0.77	0.76	0.81	0.79	0.77	0.71	0.67	0.64
global	0.54	0.72	0.96	0.91	0.54	0.65	0.94	0.90	0.54	0.68	0.95	0.91	0.77	0.77	0.82	0.80	0.77	0.70	0.75	0.69

TABLE II

Performances of our approach, compared to SIFT[14] and BOLD[16], in the *Simple Scene* scenario (boldface text highlights the best score in the related pane). The overall Precision/Recall index of LOOP is about 96%/94% showing its flexibility in general purpose real applications.

Index	Precision					Recall FScore						IOU					OIOU				
Method	Sift	Bold	LOOP ₁₀	$LOOP_{10}^S$	Sift	Bold	LOOP ₁₀	$LOOP_{10}^S$	Sift	Bold	$LOOP_{10}$	$LOOP_{10}^S$	Sift	Bold	LOOP ₁₀	$LOOP_{10}^S$	Sift	Bold	$LOOP_{10}$	$LOOP_{10}^S$	
artifact_black	0.00	0.01	0.96	0.90	0.00	0.01	0.94	0.90	0.00	0.01	0.95	0.90	0.00	0.64	0.80	0.75	0.00	0.05	0.80	0.74	
artifact_metal	0.43	0.00	0.97	0.83	0.01	0.00	0.98	0.93	0.02	0.00	0.97	0.87	0.70	0.00	0.82	0.76	0.70	0.00	0.82	0.76	
artifact_orange	0.75	0.00	0.91	0.78	0.41	0.00	0.88	0.92	0.53	0.00	0.89	0.85	0.85	0.00	0.87	0.86	0.84	0.00	0.82	0.68	
artifact_white	0.73	0.00	0.96	0.80	0.10	0.00	0.96	0.98	0.18	0.00	0.96	0.88	0.68	0.00	0.80	0.80	0.68	0.00	0.79	0.76	
clip	0.00	0.40	0.96	0.91	0.00	0.40	0.97	0.87	0.00	0.40	0.97	0.89	0.00	0.71	0.79	0.80	0.00	0.63	0.79	0.78	
screwdriver	0.00	0.11	0.94	0.94	0.00	0.11	0.90	0.48	0.00	0.11	0.92	0.63	0.00	0.75	0.76	0.75	0.00	0.75	0.76	0.74	
Untextured	0.32	0.09	0.95	0.86	0.09	0.09	0.94	0.85	0.12	0.09	0.94	0.84	0.37	0.35	0.81	0.79	0.37	0.24	0.80	0.74	
battery_black	0.59	0.91	0.98	0.91	0.58	0.91	0.99	0.97	0.58	0.91	0.98	0.94	0.74	0.61	0.84	0.81	0.74	0.24	0.84	0.81	
battery_green	0.67	0.04	0.94	0.84	0.90	0.04	0.87	0.50	0.76	0.04	0.90	0.63	0.68	0.62	0.80	0.75	0.68	0.54	0.75	0.71	
box_brown	0.37	0.24	0.78	0.74	0.99	0.24	0.84	0.46	0.54	0.24	0.81	0.57	0.80	0.70	0.74	0.69	0.80	0.69	0.51	0.23	
box_yellow	0.47	1.00	0.96	0.96	0.82	1.00	0.96	0.97	0.60	1.00	0.96	0.97	0.85	0.84	0.84	0.83	0.85	0.84	0.84	0.83	
glue	0.59	0.94	0.99	0.94	0.81	0.94	0.95	0.79	0.68	0.94	0.97	0.86	0.78	0.79	0.84	0.79	0.78	0.79	0.84	0.79	
pendrive	0.88	0.81	0.96	0.90	0.09	0.81	0.85	0.81	0.17	0.81	0.90	0.85	0.64	0.63	0.78	0.73	0.64	0.55	0.78	0.71	
Textured	0.58	0.58	0.93	0.88	0.67	0.57	0.90	0.72	0.53	0.57	0.91	0.78	0.75	0.72	0.80	0.76	0.75	0.67	0.75	0.67	
global	0.53	0.42	0.94	0.87	0.39	0.37	0.92	0.80	0.45	0.40	0.93	0.83	0.77	0.72	0.81	0.78	0.77	0.62	0.78	0.73	

TABLE III

Performances of our approach, compared to SIFT[14] and BOLD[16], in the *Hard Scene* scenario (boldface text highlights the best score in the related pane). Here, the overall Precision/Recall index of LOOP is about 94%/92% showing its robustness against high clutter backgrounds, a typical situation in real industrial environments.

rotational error respectively $\mathbf{e}_t = \{x_i - c_x, y_i - c_y\}$ and $e_{\theta} = -\text{sign}(\sin(\theta_i))(\cos(\theta_i) - 1)$, with (c_x, c_y) the center of the image. The control scheme will produce a linear velocity $\mathbf{v} = K_{P_t} \mathbf{e}_v$ and an angular velocity $\omega = K_{P_{\theta}} e_{\theta}$ to the end effector, i.e. the camera reference frame $(K_{P_t} \text{ and } K_{P_{\theta}} \text{ are}$ tunable proportional gains). The sign(\cdot) function is simply the sign function with sign(0) = 1 to avoid singularities. The above simplified robot guidance scheme is used to lead the robot in a easier condition suitable to estimate the 3D pose of the object because: knowing the extrinsics parameters of the mounted camera and the height of the table w.r.t. the robot base, it is trivial to know the 3D coordinates of an object centered in the camera viewpoint (an alternative solution to estimate 3D coordinates from multiple 2D CNN predictions can be found in [29]). Figure 8 exemplifies the described procedure by presenting a real execution of it. In the supplementary material several runs of experiment are shown from the on-board and off-board cameras. We have accomplished this task on 20 scenes with completely unseen backgrounds. The overall success rate of the final grasp is 100% for each object, except for artifcat_orange and box_brown, which instead scored 60% and 50% respectively. Hance, the success rate for the entire dataset is 92.5%. As expected from the conclusions of subsection IV-D, the grasp for highly symmetrical objects (like artifcat_orange and box brown) becomes very challenging, due to the difficulty of estimating unambiguous orientation for these samples. In these cases, the indecision of the detector combined with the stateless nature of a CNN-based Object Detector (i.e. there is no online tracking method like in [30], where a Recurrent Neural

Network is used in order to achieve a continuous estimate of the pose) leads to a detrimental oscillation of the output in the control module. Then, we corrected the model by modifying the rotational error in such a way as to treat objects as symmetric, with

$$\begin{cases} \mathbf{e}_t = \{x_i - c_x, y_i - c_y\}\\ e_{\theta}^{\text{Sym}} = \operatorname{sign}(\operatorname{tan}(\theta))(|\operatorname{sin}(\theta)|) \end{cases}$$
(4)

Accordingly, the robot guidance scheme leads the camera to move towards the nearest horizontal configuration of the target object (0° or 180° indifferently), reaching 100% of success rate also for each object (the control schemes were tested again over 20 scenes).

V. CONCLUDING REMARKS

We proposed an extension of classical CNN-based Object Detectors able to produce Oriented Bounding Boxes suitable for the 3-DoF pose estimation task. We provided our detector with a simplified procedure to gather a huge amount of training data in the field, with trifling human intervention. With this work we showed how it is possible, with the proposed techniques, to effectively use Deep Learning in a real industrial setting, exploiting a neural network as a module of a more complex control scheme. We are already working to extend our approach in order to deal also with occlusions and symmetries together with a set of more complex discretization functions beyond the simple division into equal bins. An interesting future development could be to remove completely human intervention in



Fig. 8. An exemplification of the pick-and-place system designed using LOOP as a module of the control scheme. The first row depicts frames coming from the on-board camera, with superimposed predictions, showing the robot guidance control phases. The second row shows images coming from an off-board camera framing the grasp sequence downstream of the alignment procedure.



Fig. 9. This picture shows the real output of our software for the simple and hard scene, respectively. The middle column shows Unoriented Predictions (BB) coming from the generic object detector, while the right column shows the derived Oriented Predictions (OBB).

the dataset generation stage in order to deliver a completely automated, and reliable, industrial system.

REFERENCES

- J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. of the IEEE conf. on computer vision and pattern recognition*, 2016, pp. 779–788.
- [2] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *European conf. on computer vision*. Springer, 2016, pp. 21–37.
- [3] 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.
- [4] W. Kehl, F. Manhardt, F. Tombari, S. Ilic, and N. Navab, "Ssd-6d: Making rgb-based 3d detection and 6d pose estimation great again," in *Proc. of the International Conf. on Computer Vision (ICCV), Venice, Italy*, 2017, pp. 22–29.
- [5] B. Tekin, S. N. Sinha, and P. Fua, "Real-time seamless single shot 6d object pose prediction," in *Proc. of the Conf. on Computer Vision and Pattern Recognition*. IEEE, 2018, pp. 292–301.
- [6] M. Sundermeyer, Z. Marton, M. Durner, and R. Triebel, "Implicit 3d orientation learning for 6d object detection from rgb images," in *Proc. of the European Conf. on Computer Vision (ECCV)*, 2018, pp. 699–715.
- [7] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results," http://www.pascalnetwork.org/challenges/VOC/voc2012/workshop/index.html.
- [8] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European Conf. on Computer Vision*. Springer, 2014, pp. 740–755.

- [9] S. Hinterstoisser, S. Holzer, C. Cagniart, S. Ilic, K. Konolige, N. Navab, and V. Lepetit, "Multimodal templates for real-time detection of textureless objects in heavily cluttered scenes," in *International Conf. on Computer Vision (ICCV)*. IEEE, 2011, pp. 858–865.
- [10] T. Hodan, P. Haluza, Š. Obdržálek, J. Matas, M. Lourakis, and X. Zabulis, "T-less: An rgb-d dataset for 6d pose estimation of texture-less objects," in *Winter Conf. on Applications of Computer Vision (WACV)*. IEEE, 2017, pp. 880–888.
- [11] Y. Zhou, Q. Ye, Q. Qiu, and J. Jiao, "Oriented response networks," in Conf. on Computer Vision and Pattern Recognition (CVPR). IEEE, 2017, pp. 4961–4970.
- [12] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *European Conf. on Computer Vision*. Springer, 2014, pp. 818–833.
- [13] M. Jaderberg, K. Simonyan, A. Zisserman, et al., "Spatial transformer networks," in Advances in neural information processing systems, 2015, pp. 2017–2025.
- [14] D. G. Lowe, "Object recognition from local scale-invariant features," in Proc. of the International Conf. on Computer vision, vol. 2. IEEE, 1999, pp. 1150–1157.
- [15] S. A. K. Tareen and Z. Saleem, "A comparative analysis of sift, surf, kaze, akaze, orb, and brisk," in *International Conf. on Computing, Mathematics* and Engineering Technologies (iCoMET). IEEE, 2018, pp. 1–10.
- [16] F. Tombari, A. Franchi, and L. Di Stefano, "Bold features to detect texture-less objects," in *Proc. of the International Conf. on Computer Vision*. IEEE, 2013, pp. 1265–1272.
- [17] J. Chan, J. Addison Lee, and Q. Kemao, "Border: An oriented rectangles approach to texture-less object recognition," in *Proc. of the IEEE Conf.* on Computer Vision and Pattern Recognition, 2016, pp. 2855–2863.
- [18] J. Chan, J. A. Lee, and Q. Kemao, "Bind: Binary integrated net descriptors for texture-less object recognition," in *Proc. CVPR*, 2017.
- [19] B. Liu, X. Shu, and X. Wu, "Fast screening algorithm for rotation invariant template matching," in *International Conf. on Image Processing (ICIP)*. IEEE, 2018, pp. 3708–3712.
 [20] F. Ullah and S. Kaneko, "Using orientation codes for rotation-invariant
- [20] F. Ullah and S. Kaneko, "Using orientation codes for rotation-invariant template matching," *Pattern recognition*, vol. 37, no. 2, pp. 201–209, 2004.
- [21] S. Hinterstoisser, V. Lepetit, S. Ilic, P. Fua, and N. Navab, "Dominant orientation templates for real-time detection of texture-less objects," 2010.
- [22] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," arXiv preprint arXiv:1804.02767, 2018.
- [23] Y. Jiang, S. Moseson, and A. Saxena, "Efficient grasping from rgbd images: Learning using a new rectangle representation," in *International Conf. on Robotics and Automation (ICRA)*. IEEE, 2011, pp. 3304–3311.
- [24] J. Redmon and A. Angelova, "Real-time grasp detection using convolutional neural networks," in *International Conf. on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 1316–1322.
- [25] F.-J. Chu, R. Xu, and P. A. Vela, "Real-world multiobject, multigrasp detection," *Robotics and Automation Letters*, vol. 3, no. 4, pp. 3355– 3362, 2018.
- [26] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "Orb: An efficient alternative to sift or surf," in *International Conf. on Computer Vision* (*ICCV*). IEEE, 2011, pp. 2564–2571.
- [27] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *International Conf. on Intelligent Robots* and Systems (IROS). IEEE, 2017, pp. 23–30.
- [28] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems*, 2012, pp. 1097–1105.
- [29] D. D. Gregorio, R. Zanella, G. Palli, S. Pirozzi, and C. Melchiorri, "Integration of robotic vision and tactile sensing for wire-terminal insertion tasks," *Transactions on Automation Science and Engineering*, pp. 1–14, 2018.
- [30] A. Milan, S. H. Rezatofighi, A. R. Dick, I. D. Reid, and K. Schindler, "Online multi-target tracking using recurrent neural networks." in AAAI, vol. 2, 2017, p. 4.