

A Novel Method for Object Detection using Deep Learning and CAD Models

Igor Garcia Ballhausen Sampaio¹, Luigy Machaca¹, José Viterbo¹ ^a and Joris Guérin² ^b

¹Computing Institute, Universidade Federal Fluminense, Brazil

²LAAS-CNRS, ONERA, Université de Toulouse, France

Keywords: Object Detection, CAD Models, Synthetic Image Generation, Deep Learning, Convolutional Neural Network.

Abstract: Object Detection (OD) is an important computer vision problem for industry, which can be used for quality control in the production lines, among other applications. Recently, Deep Learning (DL) methods have enabled practitioners to train OD models performing well on complex real world images. However, the adoption of these models in industry is still limited by the difficulty and the significant cost of collecting high quality training datasets. On the other hand, when applying OD to the context of production lines, CAD models of the objects to be detected are often available. In this paper, we introduce a fully automated method that uses a CAD model of an object and returns a fully trained OD model for detecting this object. To do this, we created a Blender script that generates realistic labeled datasets of images containing the object, which are then used for training the OD model. The method is validated experimentally on two practical examples, showing that this approach can generate OD models performing well on real images, while being trained only on synthetic images. The proposed method has potential to facilitate the adoption of object detection models in industry as it is easy to adapt for new objects and highly flexible. Hence, it can result in significant costs reduction, gains in productivity and improved products quality.


1 INTRODUCTION


Recently, Deep Learning (DL) has produced excellent results for Object Detection (OD) (Liu et al., 2020). On the one hand, a typical limitation with DL is the requirement of large labeled datasets for training. Indeed, although there are various large databases available online for OD, for specific industrial applications it is always necessary to create custom datasets containing the objects of interest. While the scenarios present in public datasets are useful from both a research and application standpoint, it was found that industrial applications, such as bin picking or defect inspection, have quite different characteristics that are not modeled by the existing datasets (Drost et al., 2017). As a result, methods that perform well on existing datasets sometimes show different results when applied to industrial scenarios without retraining. The process of generating a specific dataset for retraining is tedious, and can be error-prone when conducted by non-professional technicians. Moreover, generating a

new dataset and labeling it manually can be very time consuming and expansive (Jabbar et al., 2017).

On the other hand, OD for industrial production lines presents the specificity that the manufacturers often have access to the CAD models of the objects to detect. Thanks to advances in computer graphics techniques, such as ray tracing (Shirley and Morley, 2003), the generation of photo-realistic images is now possible. In such artificially generated images, the computer can be employed to obtain bounding box labeling for free. The use of synthetic images rendered from CAD models to train OD models has already been proposed in (Peng et al., 2015), (Rajpura et al., 2017) and (Hinterstoisser et al., 2018). However, their approaches are not automated as they require manual scene creation by Blender artists. In addition, the objects used in these works are usually generic, such as buses, airplanes, cars or animals.

The main contribution of this paper is to present a new method for training OD models in synthetic images generated from CAD models that is fully automatic and thus well suited for industrial use. The proposed method consists of the automatic generation of realistic labeled images containing the objects to be

^a  <https://orcid.org/0000-0002-0339-6624>

^b  <https://orcid.org/0000-0002-8048-8960>

detected, followed by the fine-tuning of a pretrained OD model on the artificial dataset. An extensive study is conducted to properly select the user-defined parameters so that it maximizes the performance on real world images. Our method is evaluated using the CAD models of two industrial objects for training, as well as real labeled images containing the objects for evaluation. The results obtained are very promising as we manage to get F1-scores above 90% on real images while training only on synthetic images.

This paper is organized as follows. Section 2 presents the related work in the field of object detection and deep learning for industry. Section 3 provides detailed explanations about the proposed method created. Section 4 describes our experiments, presents the results obtained and discusses them. Finally, conclusions and directions for future work are presented in Section 5.

2 RELATED WORK

This section presents related work about OD, industrial applications of DL-based computer vision, as well as computer vision methods using CAD models.

2.1 Object Detection

OD is a challenging computer vision problem that consists in locating instances of objects from predefined categories in natural images (Prasad, 2012). It has many applications in various domains such as autonomous driving, security and medical diagnosis (Xiao et al., 2020). Deep learning techniques have emerged as a powerful strategy for learning characteristic representations directly from data and have led to significant advances in the field of generic object detection (Liu et al., 2020). In the last decade, many competitions for object detection have been held to provide large annotated datasets to the community, and to unify the benchmarks and metrics for fair comparison between proposed methods (Everingham et al., 2010), (Lin et al., 2014), (Zhou et al., 2017), (Kuznetsova et al., 2018).

Some examples of OD methods proposed within the last few years include (He et al., 2015), where the author proposes a new network structure, called SPP-net, which can generate a fixed-length representation, regardless of the size/scale of the image. Other works such as (Jana et al., 2018) aim to improve processing speed and at the same time efficiently identify objects in the image. Finally, deeper CNNs have led to record-breaking improvements in the detection of more general object categories, a shift which came

about when DCNNs began to be successfully applied to image classification (Liu et al., 2020).

2.2 OD for Industrial Applications

Although general purpose OD methods have greatly improved thanks to the availability of large public datasets, the detection of instances in the industrial context must be approached differently, since annotated images are generally not available or rare. Indeed, to train a deep learning model, hundreds of annotated images for each object category are needed. Specific datasets need to be collected and annotated for different target applications. This process is time-consuming and laborious, and increases the burden on operators, which goes against the goal of industrial automation (Cohen et al., 2020), (Ge et al., 2020).

A public datasets adapted to the industrial context was developed in (Drost et al., 2017). Unlike other 3D object detection datasets, this work models industrial waste collection and object inspection tasks that often face different challenges. In addition, the evaluation criteria are focused on practical aspects, such as execution times, memory consumption, useful measures of correction and precision. Other examples of datasets adapted to the industrial context include (Gu erin et al., 2018b) and (Gu erin et al., 2018a).

Finally, in (Yang et al., 2019), a method to detect defects of tiny parts in real time was developed, based on object detection and deep learning. To improve their results, the authors consider the specificities of the industrial application in their method such as the properties of the parts, the environmental parameters and the speed of movement of the conveyor. This is a good example to adapt OD training methods to the specific constraints of the industrial context.

2.3 CAD Models and OD

The first commercial CAD programs came up in the 1970s, providing functions for 2D-drawing and data archival, and evolved into the main engineering design tool (Lindsay et al., 2018), (Hirz et al., 2017). These models can provide a scalable solution for intelligent and automatic object recognition, tracking and augmentation based on generic object models (Ben-Himane et al., 2010). For example, CAD models have been used to support multi-view detection (Zhang et al., 2013). In (Peng et al., 2015), 3D models were used as the primary source of information to build object models. In other works, 3D CAD models were used as the only source of labeled data (Lin et al., 2014), (Everingham et al., 2010), but they are limited to generic categories, such as cars and

motorcycles .

3 PROPOSED METHOD

An overview of this proposed method can be seen in Figure 1. First, a custom Blender code is used to generate labeled training images containing the rendered CAD model in context. Then, a pretrained object detection model is fine-tuned on the generated dataset. Finally, the model can be used for inference on real images (Figure 1b).

3.1 Image Generation

For the automatic generation of the training images, the software Blender (Blender Online Community, 2018) is used. Blender is a powerful software for 3D design, which includes features such as modeling, rigging, simulation and rendering. Blender has a good Python API, is open-source and has good GPU support.

In order to generate a synthetic training image sample, our code requires several elements. First, a CAD model of the object of interest as well as several other industrial CAD models need to be available. In the experiments of this paper, we use the two objects shown in Figure 2, for which we also have real world test images. The other objects serve as distractors to help the model focusing on the right object. The CAD models for the distractors are gathered from the Grabcad website¹. Different textures for the different distractors as well as for the background are gathered from the Poliigon website². Finally, the color and texture of the object of interest are reproduced manually.

Once we have access to all the elements above, the generation code goes as follows. A floor and a table are created and some distractors are sampled. Using physics simulation, the distractors are dropped from a random height on the table. The position of the object of interest is also randomly sampled. Once the 3D scene is created, textures and colors are sampled for the backgrounds and the distractors and the entire scene is textured. Light sources and cameras are also sampled and placed randomly. Constraints on the camera pose are applied, in order to ensure that the object appears in the camera view. Once the scene has been created, the rendering occurs and generates an image. By removing the light sources and making the object of interest a light source itself, we can generate another image which can be used for bounding box labeling. This procedure is necessary because

¹<https://grabcad.com/>

²<https://www.poliigon.com/>

even if we know the location of the object, it can be partly hidden by distractors and thus distort the labeling. Example images generated using our Blender code can be seen in Figure 3.

3.2 Model Training

In this work, we did not train a new CNN architecture from the scratch. Instead, we used one of the pre-trained models provided by TensorFlow Object Detection API (Huang et al., 2017). This approach is called transfer learning and consists in starting training from a model that already knows basic feature extraction skills and is less likely to overfit the synthetic datasets. Indeed, the diversity that we can create with Blender is limited as we cannot get an infinite amount of textures and distractors, and the diversity already encountered by the network during pre-training can help reduce overfitting. In addition, using a network pre-trained on real images can prevent the network from learning detection features that depend too much on the generation procedure.

There exists several models in the TensorFlow OD model zoo. More information on the performance of the detection, as well as the reference execution times, for each of the available pre-trained models, can be found on the Github page of the API³. In practice, the model used in this paper is the *faster_rcnn_inception_v2_coco* model, which provides a good trade-off between performance and speed.

Faster R-CNN, the model used in this work, takes as input an entire image and a set of object proposals. The network first processes the whole image with several convolutional and max pooling layers to produce a convolutional feature map. Then, for each object proposal, a region of interest (RoI) pooling layer extracts a fixed-length feature vector from the feature map. Each feature vector is fed into a sequence of fully connected layers that finally branch into two sibling output layers: one that produces softmax probability estimates over K object classes plus a catch-all “background” class and another layer that outputs four real-valued numbers for each of the K object classes. Each set of 4 values encodes refined bounding-box positions for one of the K classes. For a more detailed view about Faster-RCNN, we refer the reader to the original paper (Ren et al., 2015), or to the following tutorial (Ananth, 2019).

To train the final OD model, the TensorFlow OD API requires a specific file structure of the training images and labels. This step is carried out automatically by our script.

³<https://github.com/tensorflow/models>

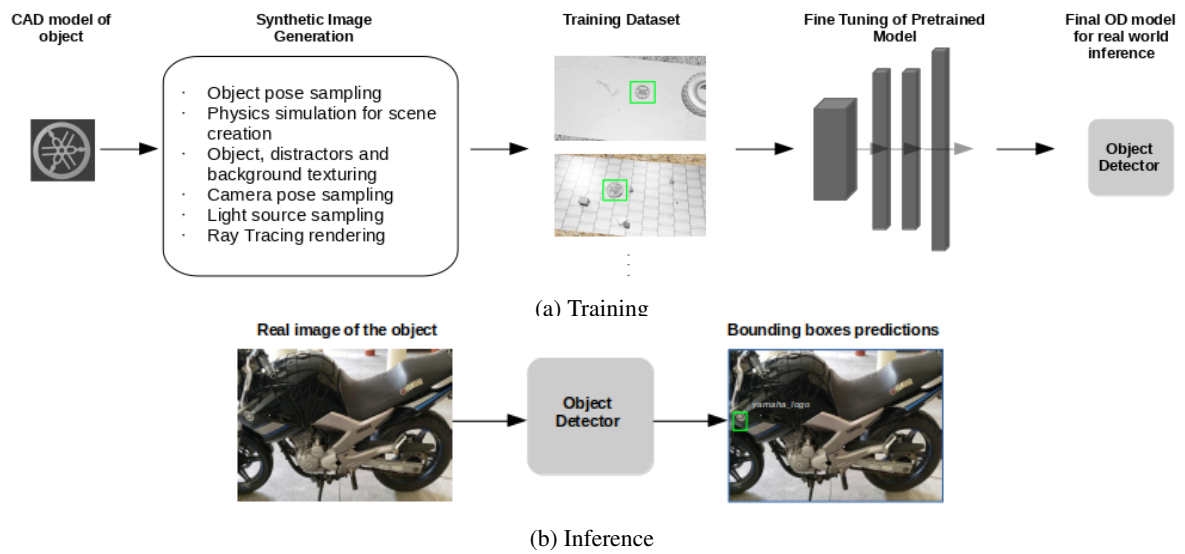


Figure 1: Overview of the proposed method for training an object detection network using a CAD model.




Figure 2: CAD models used in our experiments.

3.3 Parameter Selection Procedure

The Blender script used for image generation has many hyperparameters that must be chosen before using it, such as the number of distractors, the number of scenes generated or the resolution of synthetic images. Hence, we conduct a set of experiments to properly select these parameters in order to optimize the OD results for inference on real images. In this section, we explain the parameter selection procedure. In other words, we present the dataset on which the different sets of parameters were evaluated, as well as the metrics used to assess the quality of the results obtained with a given set of parameters. The results obtained for this parameter selection procedure are presented in Section 4.

3.3.1 Test Dataset

The objective of this work is to validate that an object detector trained on synthetic images can generalize to real world industrial cases. Hence, we use a test dataset composed of 380 real images containing the

objects corresponding to the CAD models used for training. The bounding box annotation files for the test images are generated manually using a software called LabelImg⁴. This application allows us to draw and save the annotations of each image as xml files in the PASCAL VOC format (Everingham et al., 2010).

3.3.2 OD Metrics

In order to evaluate the quality of the trained model on real images, and thus to be able to select the best hyperparameters for image generation and training, we used standard OD metrics that are presented here.

Intersection over Union (IoU): is an evaluation metric used to measure how much a predicted bounding box matches with a ground truth bounding box. For a pair of bounding boxes, IoU is defined as the area of the intersection divided by the area of the union (Figure 4). If A corresponds to the ground-truth box and B to the predicted box, then, IoU is computed as :

$$IoU = \frac{|A \cap B|}{|A \cup B|}, \quad (1)$$

where $|\cdot|$ denotes the area of a given shape. The numerator is called the overlap area and the denominator is called the combined area. IoU ranges between 0 and 1, where 1 means that the bounding boxes are the same and 0 that there is no overlap.

Precision, Recall, F1-Measure. We call confidence score, the probability that an anchor box con-

⁴<https://github.com/tzutalin/labelImg>



(a) Adblue

(b) Yamaha logo

Figure 3: Example of images generated using our custom Blender script.

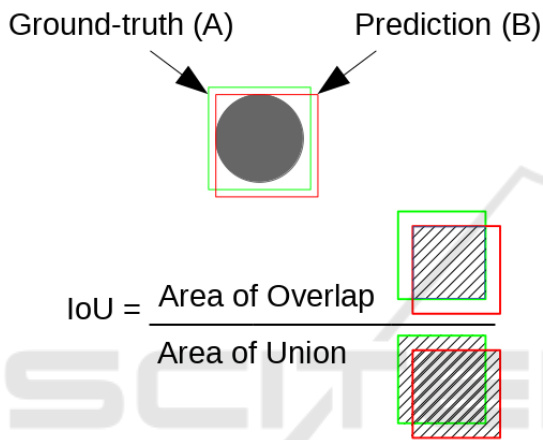


Figure 4: Intersection over Union (IoU) computation.

tains an object from a certain class. It is usually predicted by the classifier part of the object detector. The confidence score and IoU are used as the criteria to determine whether a detection is a true positive or a false positive. Given a minimal threshold on the confidence score for bounding box acceptance, and another threshold on IoU to identify matching boxes, a detection is considered a true positive (TP) if there exists a ground truth such that: confidence score $>$ threshold; the predicted class matches the class of the ground truth; and $\text{IoU} > \text{threshold}_{\text{IoU}}$. The violation of any of the last two conditions generates a false positive (FP). In case multiple predictions correspond to the same ground-truth, only the one with the highest confidence score counts as a true positive, while the others are considered false positives. When a ground truth bounding box is left without any matching predicted detection, it counts as a false negative (FN).

If we note TP, FP and TN respectively the number of True Positives, False Positives and False Negatives in a dataset, we can define the following metrics:

$$\text{Precision} = \frac{TP}{TP + FP}, \quad (2)$$

$$\text{Recall} = \frac{TP}{TP + FN}, \quad (3)$$

$$F_1 = \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}. \quad (4)$$

A high precision means that most of the predicted boxes had a corresponding ground truth, i.e., the object detector is not producing bad predictions. A high recall means that most of the ground truth boxes had a corresponding prediction, i.e., the object detector finds most objects in the images. The F_1 -Score is the harmonic mean of the precision and recall, it is needed when a balance between precision and recall is sought.

In the case of object detection on production lines, a low precision means that sometimes a part might be absent and the model would not see it, whereas a low recall means that sometimes the part is present and the model raises an alert anyways. For this reason, both a good recall and a precision are required and the choice of using the F_1 -Score metric seems appropriate.

Average Precision. After an OD model has been trained, the computation of Precision, Recall and F_1 -score depends on the value of the two thresholds defined above (for the confidence score and IoU). In order to properly choose the values of these thresholds, it is interesting to analyze the Precision x Recall curves. For each class, and for a given value of the IoU threshold, the confidence threshold is set as a variable and sampled between 0 and 1 to plot a parametric curve with precision and recall as the x and y-axis.

A class-specific object detector is considered good if the precision remains high as the recall increases, meaning that if you vary the confidence limit, the precision and recall will still be high. Hence, to compare between curves we generally rely on a numerical metric called Average Precision (AP). Since 2010, the

standard computation method for AP consists in calculating the area under the curve (AUC) of the Precision x Recall curve (Everingham et al., 2010).

4 EXPERIMENTS AND RESULTS

The results obtained for the parameter selection procedure as well as our final evaluations are presented here. These experiments were conducted on a Nvidia Quadro P5000 GPU and a 2.90GHz Intel Xeon E3-154M v5 processor (16 GB of RAM).

4.1 Hyperparameter Tuning

The parameter selection procedure is conducted exclusively on the Yamaha logo object, the best set of parameters is then tested on the Adblue object to ensure that it also performs well. The influence of four tunable parameters on the final results is studied here. For each parameter, three values were selected for the tests. These parameters and their studied values are:

- **Resolution:** 640x480, 960x540, 1080x720
- **Camera poses:** 2, 5, 20
- **Number of scenes:** 20, 50, 200
- **Number of distractors:** 0, 5, 20

From simple preliminary experiments that are not presented here, we concluded that the number of textures used for the floor, the distractors and the support should be set to the maximum number of textures available (in our case 7 for the floor and 6 for the two others). The parameter values used in this work were chosen empirically, that is, after several test scenarios, these values were the ones that generated the best performance regarding the metrics.

In total, from the values selected for the four parameters, we sampled more than 30 combinations and compared the OD results on the testing set of real images. For each combination tested, we trained the Faster-RCNN CNN on the synthetic images that were generated. We note that, for each hyperparameters combination, the experiments were repeated 10 times in order to attenuate the influence of the random components in the generation and training process. For reasons of space in this article, it was not possible to present all results. However, in order to demonstrate the importance of this parameter selection step, Table 1 shows the best and the worst configuration that were tested.

In Table 1, we can see that the distractors are an essential element in our proposed pipeline for image generation. Indeed, when removing them, we can see

Table 1: Best and worst hyperparameters configurations obtained and their corresponding results.

Parameters	Best Case	Worst Case
Resolution	960x540	960x540
cam_poses	5	5
n_scenes	20	20
n_images	100	100
n_distractors	20	0
Generation Time	1257.30	749.92
n_samples	10	10
Precision %: Avg (Std.Dev)	78.06 (15.41)	57.34 (16.27)
Recall %: Avg (Std.Dev)	96.223 (4.07)	90.71 (6.45)
F1-Score %: Avg (Std.Dev)	85.19 (10.57)	66.80 (11.49)

a drop of around 23% in F1-score, on average across the 10 experiment samples. We also tried combinations with few distractors, but the F1-score results dropped significantly. This makes sense as the real images evaluated had several distractors as well.

Another important point is that the resolution of the images generated should be greater than the inference images. In all of our tests, this scenario always produced the best results. It also makes sense as it is easier to learn from a more detailed/complex model and then evaluate in a less detailed/complex scenario.

Finally, we also tried to increase the number of generated training images to see if this would lead to an increase in performance. Surprisingly, we acknowledged that the performance dropped for the case with 20 distractors, 20 camera poses and 50 scenes (1000 images). This might mean that when presented too many synthetic images, the model starts overfitting to the biases involved by our generation process, and it also indicates that we do not need a large number of images to train our model. In addition to this performance drop, generating ten times more images also makes the proposed pipeline almost 25 times slower (31037.16 seconds).

4.2 Results

In this section, we evaluate the results of the best combination of parameters (Best Case from Table 1) in more details. These results are presented in Table 2, they correspond to using a confidence threshold of 0.9 and an IoU threshold of 0.5. From Table 2, we can see that the best parameters identified using only the Yamaha logo produce similar results when applied to another object (Adblue). This suggests that the proposed parameters for our method seem to be well suited for different objects and thus could generalize well to various industrial use cases without additional parameter tuning.

Table 2: Results obtained with the best set of hyperparameters.

Object	Precision %	Recall %	F1-Score %
Adblue	85.11	80.00	81.93
Yamaha logo	78.06	96.22	85.19

4.3 Discussion

It is difficult to compare our results with other works in the literature. Indeed, as far as we know, the approach presented in this work is the first proposal to build a fully automated pipeline that takes as input the CAD model of an object and outputs a trained object detection model for this object without any real image. For fair comparison we would need to compare our work with other end-to-end systematic approaches to build OD models from CAD models, which is impossible as it does not exist. Else, we hope that the results presented in this work can serve as a good baseline for comparison of future works in this research direction.

However, we give a rough comparison with other relevant works to give an idea of how well our approach is performing. In (Mazzetto et al., 2019) the detection of objects in an automobile production line was implemented, using only real images of the objects. In this work, the estimated detection accuracy was around 90 %, which is only about 5 to 10 % better than the results obtained in our work using only synthetic images. In (Jabbar et al., 2017), the authors also train an OD model using synthetic images generated with Blender and evaluate the results in real images. However, this approach is not entirely automated since the scenes are created manually by Blender artists to ensure photo-realism. The object used in this work for evaluation is a glass of wine and the maximum AP obtained is 71.14 %. We can see that our systematic approach seems to work better than this approach, however, we cannot reproduce the method on our objects, as we cannot create the scenes manually in the same way that they would. The potential better performance of our approach can be explained by the fact that the loss of photo-realism can be compensated by the higher number of images in our synthetic datasets. Indeed, with our fully automated approach, it is faster and requires no effort to generate more data, unlike in (Jabbar et al., 2017).

5 CONCLUSIONS

This work presents a systematic approach to train object detection models to address industrial scenarios, using only a CAD model of the object of interest as

input. The method first generates realistic synthetic images using a custom Blender script, and then trains a faster-RCNN OD model using the TensorFlow OD API. To understand and optimize the different parameters in the proposed pipeline, a systematic parameter selection study is conducted using a Yamaha logo CAD model for training and real images containing the same object in context for evaluation. The selected hyperparameters are then tested on an other object, showing that they can generalize to different scenarios.

Over the last decade, successful deep learning methods have been developed to tackle the challenging problem of generic object detection. However, when it comes to the problem of OD in an industrial environment, the availability of good quality data becomes a bottleneck. To address this issue we proposed to use synthetic images for training, which is challenging as it might not reflect the high variability found in in real industrial environment (objects, pieces and scenery, etc.). In addition, there is also a difficulty in finding CAD models of specific industrial objects so that they can be trained and other approaches can be tested and compared. Thus, as a consequence of this work, a set of data was produced and made publicly available for future research⁵.

Therefore, the main conclusion from this work is that it is possible to train an object detection model on a set of synthetic images generated from CAD models with excellent performance. In addition, it was shown that a large set of images is not needed to obtain a significant result. Our experiments indicate that the proposed rendering process is sufficient to obtain good performances and that the way of building and rendering the scenes is crucial for the final result.

ACKNOWLEDGEMENTS

Our work has benefited from the AI Interdisciplinary Institute ANITI. ANITI is funded by the French “Investing for the Future – PIA3” program under the Grant agreement n°ANR-19-PI3A-0004.

REFERENCES

- Ananth, S. (2019). Faster R-CNN for object detection, a technical paper summary.
- Ben-Himane, S., Hintestroisser, S., and Navab, N. (2010). Computer vision CAD models. US Patent App. 12/682,199.

⁵https://github.com/igorgbs/systematic_approach_cad_models

- Blender Online Community (2018). *Blender - a 3D modelling and rendering package*. Blender Foundation.
- Cohen, J., Crispim-Junior, C., Grange-Faivre, C., and Tougne, L. (2020). CAD-based learning for ego-centric object detection in industrial context. In *15th International Conference on Computer Vision Theory and Applications*, volume 5, pages 644–651. SCITEPRESS.
- Drost, B., Ulrich, M., Bergmann, P., Hartinger, P., and Steger, C. (2017). Introducing MVTEC ITODD—a dataset for 3d object recognition in industry. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 2200–2208.
- Everingham, M., Van Gool, L., Williams, C. K., Winn, J., and Zisserman, A. (2010). The pascal visual object classes (VOC) challenge. *International journal of computer vision*, 88(2):303–338.
- Ge, C., Wang, J., Wang, J., Qi, Q., Sun, H., and Liao, J. (2020). Towards automatic visual inspection: A weakly supervised learning method for industrial applicable object detection. *Computers in Industry*, 121:103232.
- Guérin, J., Gibaru, O., Nyiri, E., Thiery, S., and Palos, J. (2018a). Automatic construction of real-world datasets for 3D object localization using two cameras. In *IECON 2018-44th Annual Conference of the IEEE Industrial Electronics Society*, pages 3655–3658. IEEE.
- Guérin, J., Gibaru, O., Nyiri, E., Thiery, S., and Boots, B. (2018b). Semantically meaningful view selection. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1061–1066. IEEE.
- He, K., Zhang, X., Ren, S., and Sun, J. (2015). Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 37(9):1904–1916.
- Hinterstoisser, S., Lepetit, V., Wohlhart, P., and Konolige, K. (2018). On pre-trained image features and synthetic images for deep learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Hirz, M., Rossbacher, P., and Gulanová, J. (2017). Future trends in CAD—from the perspective of automotive industry. *Computer-aided design and applications*, 14(6):734–741.
- Huang, J., Rathod, V., Sun, C., Zhu, M., Korattikara, A., Fathi, A., Fischer, I., Wojna, Z., Song, Y., Guadarra, S., et al. (2017). Speed/accuracy trade-offs for modern convolutional object detectors. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7310–7311.
- Jabbar, A., Farrowell, L., Fountain, J., and Chalup, S. K. (2017). Training deep neural networks for detecting drinking glasses using synthetic images. In *International Conference on Neural Information Processing*, pages 354–363. Springer.
- Jana, A. P., Biswas, A., et al. (2018). YOLO based detection and classification of objects in video records. In *2018 3rd IEEE International Conference on Recent Trends in Electronics, Information & Communication Technology (RTEICT)*, pages 2448–2452. IEEE.
- Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Mallocci, M., Duerig, T., et al. (2018). The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *arXiv preprint arXiv:1811.00982*.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. (2014). Microsoft COCO: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer.
- Lindsay, A., Paterson, A., and Graham, I. (2018). Identifying and quantifying inefficiencies within industrial parametric CAD models. In *Advances in Manufacturing Technology XXXII: Proceedings of the 16th International Conference on Manufacturing Research*, volume 8, page 227. IOS Press.
- Liu, L., Ouyang, W., Wang, X., Fieguth, P., Chen, J., Liu, X., and Pietikäinen, M. (2020). Deep learning for generic object detection: A survey. *International journal of computer vision*, 128(2):261–318.
- Mazzetto, M., Southier, L. F., Teixeira, M., and Casanova, D. (2019). Automatic classification of multiple objects in automotive assembly line. In *2019 24th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*, pages 363–369. IEEE.
- Peng, X., Sun, B., Ali, K., and Saenko, K. (2015). Learning deep object detectors from 3D models. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1278–1286.
- Prasad, D. K. (2012). Survey of the problem of object detection in real images. *International Journal of Image Processing (IJIP)*, 6(6):441.
- Rajpura, P. S., Bojinov, H., and Hegde, R. S. (2017). Object detection using deep CNNs trained on synthetic images. *arXiv preprint arXiv:1706.06782*.
- Ren, S., He, K., Girshick, R., and Sun, J. (2015). Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99.
- Shirley, P. and Morley, R. K. (2003). *Realistic ray tracing*. AK Peters/CRC Press.
- Xiao, Y., Tian, Z., Yu, J., Zhang, Y., Liu, S., Du, S., and Lan, X. (2020). A review of object detection based on deep learning. *Multimedia Tools and Applications*, pages 1–63.
- Yang, J., Li, S., Wang, Z., and Yang, G. (2019). Real-time tiny part defect detection system in manufacturing using deep learning. *IEEE Access*, 7:89278–89291.
- Zhang, X., Yang, Y.-H., Han, Z., Wang, H., and Gao, C. (2013). Object class detection: A survey. *ACM Computing Surveys (CSUR)*, 46(1):1–53.
- Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., and Torralba, A. (2017). Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1452–1464.