The Impact of Domain Randomization on Object Detection: A Case Study on Parametric Shapes and Synthetic Textures*

Atabak Dehban^{1,2}, João Borrego¹, Rui Figueiredo^{1,3}, Plinio Moreno¹, Alexandre Bernardino¹, and José Santos-Victor¹

Abstract—Recent advances in deep learning-based object detection techniques have revolutionized their applicability in several fields. However, since these methods rely on unwieldy and large amounts of data, a common practice is to download models pre-trained on standard datasets and fine-tune them for specific application domains with a small set of domain-relevant images. In this work, we show that using synthetic datasets that are not necessarily photo-realistic can be a better alternative to simply fine-tune pre-trained networks. Specifically, our results show an impressive 25% improvement in the mAP metric over a fine-tuning baseline when only about 200 labelled images are available to train. Finally, an ablation study of our results is presented to delineate the individual contribution of different components in the randomization pipeline.

I. INTRODUCTION

With the availability of advanced object detectors [1], [2], [3], [4], these systems and their variations have found many applications ranging from face detection [5], to medical applications [6], and to robotics [7].

However, training these systems from scratch is still a challenge as these methods rely on the availability of large, annotated, and high quality datasets. One common approach to circumvent this issue is to re-use detectors that were pre-trained on large and available datasets such as Common Objects in COntext (COCO) [8] and ImageNet [9] and later, apply some form of domain adaptation technique [10] for the particular task at hand using a smaller, domain specific dataset [11], [12], [13]. This approach results in varying degrees of success (refer to [14] for a study on how knowledge can be transferred across different tasks). This line of research has been accelerated, thanks to the availability of high quality open source implementations of *state-of-the-art* object detectors [15], [16].

These domain adaptation techniques usually assume, implicitly, that the underlying data–structure is shared across the domains. However, this premise does not strongly apply to many applications, especially when the target domain does not significantly overlap with the outdoor images that make up a large portion of both ImageNet and COCO.

To overcome these challenges, in this work, we are investigating the usage of *domain randomization* [17] to facilitate

*This work is partially supported by the Portuguese Foundation for Science and Technology (FCT) project [UID/EEA/50009/2019].

¹Institute for Systems and Robotics, Instituto Superior Tećnico, Universidade de Lisboa, Lisbon, Portugal

{adehban,jborrego,ruifigueiredo,plinio, alex,jasv}@isr.tecnico.ulisboa.pt

³ Institute for Systems and Robotics, Universidade de Coimbra, Coimbra, Portugal



Fig. 1: An example image from the test set, annotated by the object detector. Annotations are red for boxes, blue for cylinders and green for sphere.

the adaptation of an object detector, namely Single–Shot Detector (SSD) [2], to detect three categories of objects: cylinders, spheres, and boxes.

This task is accomplished based on our previous work on an open source plugin [18] developed for Gazebo simulator [19]. This plugin was selected as it streamlines the generation and rendering of different objects as long as their mesh description is available. In addition, adding parametric classes of objects using this plugin is quite straightforward. Finally, Gazebo is the current *de facto* standard for robotics which covers several physics engines, families of robots, and different types of actuators¹. These features may encourage other researchers to develop *domain randomization* experiments where multiple robots can interact with several objects while learning new skills.

According to our experiments, domain randomization can substantially increase the accuracy of object detectors at least in situations where only a relatively small domain–specific dataset of annotated images is available. More importantly, the fact that the synthetic dataset does not necessarily need to be photo-realistic helps to significantly lower the barrier in applying this technique in different applications.

The aforementioned benefits notwithstanding, to the best of our knowledge domain randomization has not been subject to meticulous studies characterizing its impact on multicategory object detection. Considering that domain randomization entails programming a desired task in a simulation environment, this issue can stymie its adoption in practice. Thus, the focal question we are trying to address here is the

² Champalimaud Centre for the Unknown, Lisbon, Portugal

¹https://en.wikipedia.org/wiki/Robotics_ simulator as of August 6, 2019

benefit one can expect from applying domain randomization in real-world robotic tasks.

The main contributions of this paper can be summarized as following:

- Domain randomization is often used to apply the models trained in simulation, directly to real tasks. Contrary to this popular approach, we have shown **substantial improvements** in the accuracy of SSD can be expected when fine-tuning starts from models trained on synthetic data, provided that at least a small, domain–specific dataset is available;
- We conducted a comprehensive study in order to determine the contribution of individual components of the domain randomization pipeline and discuss the importance of viewpoint variations, different types of textures and number of available images for training;
- We have made significant contributions to our previously introduced open source Gazebo plugin, which has resulted in *doubling* the speed of scene generation pipeline, by effectively removing redundant object load times. These modifications have greatly facilitated the study of domain randomization in object category detection.

The rest of this paper is organized as follows: in section II we examine the related work on different domain adaptation techniques relevant to object detection, that have been studied in the literature. In section III we explain the setup of the experiments, as well as our contributions to the Gazebo plugin which has made this work possible. For the sake of completeness, a brief overview of SSD is also provided. Section IV discusses the results of using domain randomization on object detection and the significance of different components in the domain randomization pipeline. It also benchmarks the importance of our contributions to the Gazebo plugin for scene generation. Finally, we draw our conclusions and discuss promising future research directions in section V.

II. RELATED WORK

Recent advances on deep learning and parallel computing have boosted research and contributed to many breakthroughs in machine learning and computer vision. Being capable of learning the underlying non-linear structure in high dimensional data, they have achieved *state-of-theart* performance in image classification [20], detection [4] and segmentation tasks [21]. However, supervised training of deep neural networks relies on the availability of large datasets, hand-labeled in a laborious and time consuming manner.

In this section we overview the main concepts and related work on automated, computer driven data augmentation techniques for computer vision applications.

A. Reality Gap

The discrepancy between the real world and simulated, computer generated environments is often referred to as the reality gap. There are two common approaches to bridge this disparity: either reducing the gap by attempting to increase the resemblance between the two domains or explore methods that are trained in a more generic domain, representative of both simulation and reality domains simultaneously. To achieve the former, one may increase the accuracy of the simulators in an attempt to obtain high-fidelity results [22], [23]; or use Generative Adversarial Networks (GANs) to produce photo-realistic images from simulated renders [24]. Both methods require great effort in the creation of systems which model complex physical phenomena to attain realistic simulations. Our work focuses mainly on the latter approach. Instead of diminishing the reality gap in order to use traditional machine-learning methods, we analyze methods that are aware of this disparity.

B. Data Augmentation

An alternative method to obtain large amounts of annotated training data is to enrich a small dataset with new labelled elements. In [25], the authors generate synthetic composite images for training neural networks for object detection. They propose methods in which 2D cropped object images are superimposed into a real-world RGB-D scene. Moreover, their proposal integrates scene contextual information in the data generation process.

This work demonstrates that the performance of state-ofthe-art object detectors significantly improves when trained with both synthetic and real data compared to training with real data alone. The data generation method is tested with two publicly available datasets, GMU-Kitchens [26] and Washington RGB-D Scenes V2 [27].

Compositional Language and Elementary Visual Reasoning diagnostics (CLEVR) [28] can be considered an alternative domain randomization dataset with the same type of categories to ours. It was initially proposed for the task of visual question answering and Blender [29] is used as the rendering engine. However, the high fidelity renders and low texture variations in this dataset are orthogonal to domain randomization spirit, and any obtained results could be either attributed to the similarity of the objects to real examples or to limited texture variations. Thus, we treat this dataset as another test set and report the results of applying different variations of our detectors on this dataset.

C. Domain Randomization

Rather than attempting to perfectly emulate reality, one may create models that strive to achieve robustness to high variability in the environment. Domain randomization [17] is a simple yet powerful technique for generating training data for machine-learning algorithms. The goal is to synthetically generate or enhance the data, in order to introduce random variances in the environment properties *that are not essential to the learning task*. This idea dates back to at least 1997 [30], with Jakobi's observation that evolved controllers exploit the unrealistic details of flawed simulators. While working on evolutionary robotics, he concluded that controllers can evolve to become more robust by introducing random noise in all the aspects of simulation which do not have a basis in reality, and only slightly randomizing the remaining which do. It is expected that given enough variability in the simulation, the transition between simulated and real domains is perceived by the model as a mere disturbance, to which it has became robust.

Different to our study, [17] has applied domain randomization for *single object localization*. Here, we have focused on a different application, namely *multi-category detection*. Additionally, we believe [17] has introduced a very powerful idea, however, there are some extra add-ons and limitations which should be peeled by the sort of ablation studies presented here:

Relying on object's mesh. The localizers of [17] require the mesh of the object for simulation. Our experiments find it an unnecessary complication for detection.

No ablation study on texture variations. [17] did not study the influence of different synthetic textures. More importantly, the most useful synthetic texture, Perlin noise (c.f. Fig. 10) was absent from that work.

Training one localizer per object. In [17], the authors have trained one object detector per target object. Thus, no generalization across objects is achieved.

Concurrent to our work, [31] reports the effect of overlaying real textures on the accuracy of multiple state-of-theart object detectors in a single-class outdoor car detection scenario. In contrast, we report the impact of overlaying synthetically generated patterns with different characteristics and increasing complexity on the accuracy metrics in a **multiple category** indoor detection of parametric shape primitives scenario.

III. METHODS

In order to apply an object detector in a new domain, it is often required to collect some training samples from the domain at hand. Labelling data for object detection is harder than labelling it for object classification, as bounding box coordinates are needed in addition to target object's identity, which increases to the importance of optimally benefiting from the available data.

After data collection, a detector is selected, commonly based on a trade-off between speed and accuracy, and is finetuned using the available "target domain" data. Our proposal is to use a synthetic dataset, with algorithmic variations in irrelevant aspects of simulation, instead of relying on pretrained networks on datasets which share little resemblance to the task at hand. This approach is further detailed in this section.

A. Single–Shot Detector

In all of our experiments, SSD was used as the base detector as it is one of the few detectors that can be applied in real-time while showing a decent accuracy. However, there is nothing in SSD which makes it more or less susceptible to benefits of domain randomization, thus, we expect our results to directly generalize to other deep learning based detectors.

The inner workings of SSD is briefly described here, however, readers should refer to the original publication [2] for a comprehensive study of the detector. At the root of all deep learning based object detectors, there exists a base Convolutional Neural Network (CNN) which is used as feature extractors for further down-stream tasks, *i.e.* bounding box generation and foreground/background recognition. Similar to YOLO architecture [3], SSD takes advantage from the concept of *priors* or *default boxes*² where each cell identifies itself as including an object or not, and where this object exists, relative to a default location. However, unlike YOLO, SSD does this at different layers of the base CNN. Since neurons in different layers of CNNs have different receptive fields in terms of size and aspect ratios, effectively, objects of various shapes can be detected.

During training, if a ground truth bounding box matches a default box, *i.e.* they have an Intersection over Union (IoU) of more than 0.5, the parameters of how to move this box to perfectly match the ground truth are learned by minimizing an smooth L1 metric. Hard negative mining is used to create a more balanced dataset between foreground and background boxes. Finally, Non-Maximum Suppression (NMS) is used to determine the final location of the objects.

Unlike the original SSD architecture, we used MobileNet [32] as the base CNN for feature extraction in all experiments. MobileNet changes the connections in a conventional CNN to drastically reduce its number of parameters, without having a notable toll on performance, relative to a comparable conventional CNN.

B. Contributions to Gazebo Plugin

Our contribution to the open-source Gazebo plugin for domain randomization [18] consists of an optimization in the scene composition, which almost doubled performance. Originally, each scene required parametric objects to be generated from a Simulation Description Files (SDF)³ formatted string, which was altered during run-time in order to allow for different object dimensions and visuals. Furthermore, objects were created and destroyed in between scenes.

Instead, we first spawn the maximum number of objects on scene of each type. Then, in each scene we alter their visual properties from within the Gazebo engine, by for instance changing their scale and pose, which results in a substantial performance boost.

In addition, we improved the existing auxiliary pattern generation module in order to exploit parallelism in batch generation of Perlin noise textures, using OpenMP⁴ framework.

C. Experiment Design and Setup

We have conducted various experiments and tests to quantify the results of different scenarios. Initially, two sets of 30k synthetic images are generated. The modifications we mentioned in the previous subsection have greatly facilitated this process. These two sets differed from one another by the degree in which the virtual camera in the scene has

²Called anchor box in YOLO.

³http://sdformat.org/ as of August 6, 2019

⁴https://www.openmp.org/ as of August 6, 2019

changed its location. In the first set, the viewpoint was fixed, whereas in the other set, its location varied largely across the scene (Fig. 5). More details will be provided in section IV.

Four types of textures were used in the generation of synthetic images, which have been employed in recent research applying domain randomization [17], [33]. Specifically, these include flat colors, gradients of colors, chess patterns, and Perlin noise [34], which can be seen in Fig. 2. The textures are randomly attached to either objects or to the background.



Fig. 2: Example synthetic scene employing all 4 texture patterns, labelled by the plug-in. The ground has a flat color, box has gradient, cylinder has chess, and sphere has Perlin noise. Any texture type can be applied to any object.

In addition, we have collected 250 real images in the lab, out of which 49 contain objects unseen in training, for the sole purpose of reporting final performance (Fig. 1). The train, validation and test partitions of our real image dataset is specified in table I.

TABLE I: Count of real images in train, validation and test partitions.

| Training | Validation | Test | Total |
|----------|------------|------|-------|
| 175 | 26 | 49 | 250 |

In this dataset, there was no consideration to explicitly keep the percentage of different classes balanced (table II), as such, we have also reported precision-recall curves for each class. Finally, all our reported metrics are calculated with an IoU of 0.5.

TABLE II: Amount of different classes in the real dataset.

| Partition | # Box | # Cylinder | # Sphere | Total |
|-----------|-----------|------------|----------|-------|
| Train set | 502 (63%) | 209 (26%) | 86 (11%) | 797 |
| Test set | 106 (40%) | 104 (40%) | 53 (20%) | 263 |

For baseline calculations, we have used SSD, trained on COCO, and fine-tuned it on the train set until the performance on validation set stopped improving.

In other experiments, we have used MobileNet which was trained on ImageNet as the CNN classifier of SSD and first fine-tuned the detector on synthetic datasets with bigger learning rates and later, in some experiments, finetuned again with smaller learning rates on the real dataset.

Finally, smaller synthetic datasets of 6k images were generated, each with a type of texture missing, and an additional baseline for comparison, which includes every pattern

type. These datasets allowed us to study the contribution of each individual texture in the final performance, as well as performance comparison of the smaller synthetic datasets.

IV. EXPERIMENTS AND RESULTS

All synthetic images have Full-HD (1920 \times 1080) resolution and are encoded in JPEG lossy format, to match the training images taken by Kinect v2.0 that were used in our experiments. For training and testing, images are down-scaled to half these dimensions (960 \times 540) which is the resolution employed for all test scenarios in our pipeline. Examples of the real datasets can be seen in Fig. 3.



(a) Training Set Examples;



(b) Test Set Examples;

Fig. 3: Example images from real (a) training and (b) test sets, annotated with ground truth and detector outputs, respectively.

Networks were trained with mini-batches of size 8, on a machine with two Nvidia Titan Xp GPUs, for a duration depending on the performance in a real image validation set. We have only used horizontal flips and random crops, with parameters reported in original SSD paper, as the preprocessing step, since we are interested in studying the effects of synthetic data and not different pre-processings. Finally, in compliance with the findings in [31], all the weights of the network are being updated in our experiments.

Our code and dataset are currently hosted on GitHub⁵ and our Laboratory's webpage⁶. Accordingly, some training hyper parameters are omitted from the current report.

A. Benchmarking Contributions to Gazebo Plugin

In [18], we reported that a dataset of 9.000 Full-HD (1920 \times 1080) images took roughly 3 hours to generate.

```
<sup>5</sup>https://github.com/jsbruglie/
```

- tf-shape-detection, as of August 6, 2019
- ⁶http://vislab.isr.ist.utl.pt/datasets/
- #shapes2018, as of August 6, 2019



Fig. 4: Precision-recall curves of different variants of the detectors after fine-tuning on the real dataset.

In a similar computer, we tested the updated plugin with our modifications and obtained almost double of the speed performance, generating 9.000 synthetic images in little over 1h30min, while benefiting from a larger set of available random textures (a total of 60.000 textures, compared to the reported 20.000), which expectedly should have increased run-time.

Our novel approach allows us to alter the properties of the objects directly through the rendering engine API, which is much more efficient than spawning and removing objects with different features. Specifically, objects are spawned below the ground plane and moved to desired location in the new scene. By changing its scale vector we can effectively morph the object shape. Finally, we load the random textures as Gazebo resources on launch, and can apply them directly, although they are only loaded into memory once they are required by the rendering engine.

B. Effects of Domain Randomization on Object Detection

In this subsection, we wish to quantify how much an object detector performance would improve due to the usage of synthetic data. To this purpose, initially, we fine-tuned a SSD, pre-trained on COCO dataset with our real image dataset. In the subsequent sections we refer to this network as baseline.

Afterwards, we trained a SSD with only its classifier pretrained on ImageNet, using our two synthetic datasets of 30k images each, as described in section III-C.

Both of these datasets contain simulated tabletop scenarios with a random number of objects $N \in [2,7]$, each in one of three classes: box, cylinder or sphere. These objects are placed randomly on the ground plane in a 3×3 grid, to avoid overlap.

In the first dataset, the camera pose is randomly generated for each scene, such that it points to the center of the object grid. This generally results in high variability in the output, which may improve generalization capabilities of the network at the expense of added difficulty to the learning task, as, for instance, it exhibits higher levels of occlusion. In the second dataset, the camera is fixed overlooking the scene at a downward angle, which is closer to the scenario we considered in the real dataset. Example scenes with viewpoint candidates for each dataset are shown in Fig. 5.



(a) Moving Viewpoint;



(b) Fixed Viewpoint;

Fig. 5: Viewpoint candidates in synthetic scene generation. (a) Viewpoint changes both position and rotation in between scenes. Subfigure represents four possible camera poses. (b) Viewpoint is static.

In addition to the camera, the scene light source is always allowed to move in a manner akin to the camera, in the first dataset.

The networks were trained on these datasets for over 90 epochs, based on their performance on the validation set, similar to the baseline. These networks were then directly applied to the test set (which has real images) without any fine-tuning on our dataset of real object data, in order to quantify how much knowledge can be directly transferred from synthetic to real domain.

Finally, these two detectors were fine-tuned on the real dataset. The result of this analysis is depicted in Fig. 6.



Fig. 6: Per class AP and mAP of different detectors. MV: Moving Viewpoint; FV: Fixed Viewpoint; Real: finetuned on the real dataset

The network trained on the dataset with no camera pose

variation and fine-tuned on real data exhibits the best performance at 0.83 mAP, which corresponds to an improvement of 26% over baseline.

Perhaps unintuitively, variations in camera position have slightly decreased the mAP score. We re-emphasize that the goal of this work was to generalize detections across unseen instances of the categories, given only a small real dataset and inevitably, data artifacts always leak in such small datasets. However, as noted by Jakobi (refer to Sec. II-C), for domain randomization to optimally work, one should only randomize aspects of the environment that are expected to vary in the real scenario *e.g.* moving viewpoint is useful for a robot with moving eyes but detrimental for fix-installed cameras. This result holds, even though we did not put any effort in aligning the simulated camera to the real one and the fixed pose of the camera is different in train and test sets.

However, it is expected that the network trained on the moving viewpoint dataset is more robust and would perform better if it is tested against a test set with varying camera/light positions.

Fig. 7 shows the precision-recall curves of different networks for each class. Consistently, the networks trained on the fixed viewpoint dataset and fine-tuned on the real dataset out-perform other variations. This trend is only less prominent in the case of *sphere* class, where, seemingly, due to the smaller examples of this class in the real dataset (table II) the training benefits less from fine-tuning on the real dataset for some values of iso-f1 surfaces. We hypothesize that more real sphere examples could help the detector in improving the Average Precision (AP) score for spheres.

C. Performance Evaluation on CLEVR

Our real dataset which is used throughout this paper was collected in house and thus, it has a limited size. Even though our number of available real samples is representative of many robotic scenarios, it still may cause *sample bias* and limit the generalization of our results. To investigate further, we have used CLEVR as another real dataset to test the effect of domain randomization. Fig. 8 shows some examples from this dataset.

Despite the high fidelity of these samples, there are not many variations in the dataset in terms of size and texture of objects. We have tested the variations of the detector described in the previous section against CLEVR validation set and report the results in Fig. 9.

According to our experiments, it seems the rendering artefact even in CLEVR is very prominent and only using the fixed view-point version of our dataset is enough to achieve the highest performance, *i.e.* any fine-tuning on real images hinders performance on CLEVR detection task.

In other experiments, due to low variations of textures in CLEVR, we managed to achieve almost perfect detection (mAP>0.98%) on the entire validation set by training a detector only on a small sample of 200 images.

D. Individual Contribution of Texture Patterns

A valid question in domain randomization research is the contribution of including various textures as well as the

importance of sample sizes. To study this question, we have created smaller synthetic datasets with only 6k images, where in each of them one specific texture is missing. Similar to previous subsection, MobileNet pre-trained on ImageNet was selected as the classifier CNN, but the detectors were instead trained on these smaller synthetic datasets and then, in some cases, fine-tuned on the real dataset.

The training of all networks on synthetic datasets lasted for 130 epochs, where the mAP stopped improving over the validation set. Finally, these networks were fine-tuned with the real-image dataset.

The results of these experiments are reported in Fig. 10. By comparing figures 10a and 10b it is clear that all variations have benefited from fine-tuning with the real dataset.

About the individual performances, generally speaking, by removing more and more complex textures before finetuning (flat to be the least complex and Perlin to be the most complex), the performance hurts, and we found Perlin noise to be a vital texture for object detection, while the flat texture has the least significance. Even after fine-tuning, flat texture seems to be the most detrimental and removing it improves performance on the real test set. Consistent with this observation, according to Fig. 10b, the small dataset with all the textures cannot always compete with some of the datasets where a texture is missing.

According to figures 10a and 11, the detector trained on the small dataset with all texture classes, performed comparably to other variations on the **validation set** during training, however, presumably due to smaller number of samples and simultaneously, so many texture classes, overfitted to the objects in the train set and failed to generalize as well as others to the objects in the test set.

Regarding the number of samples, our in house study with 200k synthetic images (unreported) in line with the findings in [31], suggests that more is not always better. The network trained on our smaller dataset of only 6k images without the "flat" texture has even slightly out-performed the network that was trained on 30k synthetic images. This result seems to be consistent for detectors with classifiers trained on real images, trained on synthetic data and then again fine-tuned with real samples. After a certain limit on number of synthetic images, the mAP performance oscillates for one or two percent.

V. CONCLUSIONS AND FUTURE DIRECTIONS

In this work we have shown that multi-category object detection pipelines can significantly benefit from pre-training on synthetic non-photo-realistic datasets. Our modifications to an open-source plugin have enabled us to rapidly test different variations in the synthetic data and assess the importance of various components such as texture complexity and sample size.

According to our experiments, increasing texture complexity in the synthetic data should be compensated by larger number of samples, however, big gains in detector accuracy can be obtained with synthetic datasets that are orders of



Fig. 7: Precision-recall curves of different variants of the detectors. For abbreviations refer to Fig. 6.



Fig. 8: samples from CLEVR dataset



Fig. 9: Per class AP and mAP of different detectors on CLEVR dataset. For abbreviations refer to Fig. 6.



Fig. 10: Performance of SSD on test set during training on smaller datasets of 6k images, each missing a type of texture, with the exception of the baseline, (a) prior and (b) after finetuning on real image dataset.



(b) After fine-tuning;

Fig. 11: Performance of SSD on validation set during training on smaller datasets of 6k images, each missing a type of texture, with the exception of the baseline, prior and after fine-tuning on real image dataset ((a), (b) respectively).

magnitude smaller than COCO or ImageNet as long as a classifier trained on real datasets is being used.

Our modifications to the plugin for the synthetic data generation will facilitate the creation of scenes for other types of studies in domain randomization, such as the impact of clutter and the increasing number of new object classes. The choice of Gazebo will stimulate the creation of scenes for deep learning experiments in robotics, such as object tracking and mobile manipulation.

In real scenarios where final performance metric is usually the most pertinent consideration, various data augmentation techniques such as color intensity distortions, random crops, etc. should be added to the training pipeline of domain randomization to improve the generalization capabilities of the detector at test time.

We believe enriching the plugin with more texture categories and combinations of categories can significantly improve the synthetic data quality for domain randomization studies. More specifically, currently no synthetic object can have more than one texture, where as in reality, e.g. a box can have different textures on each side.

Finally, in this work we have studied the simplest form of domain adaptation, namely fine-tuning, as it is by far the most popular approach due to its simplicity, and we didn't want to dilute the effect of domain randomization with advanced adaptation techniques. However, there are other adaptation methods, e.g. adversarial adaptation [35], that are suitable when the cost of collecting real data is even more significant, such as instance segmentation tasks, because they allow the transfer knowledge to "target domain" without supervised signal.

ACKNOWLEDGMENT

The Titan Xp GPUs used for this research were donated by the NVIDIA Corporation.

References

- J. Dai, Y. Li, K. He, and J. Sun, "R-fcn: Object detection via region-based fully convolutional networks," in Advances in neural information processing systems, 2016, pp. 379–387.
- [2] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "SSD: Single shot multibox detector," in *ECCV*. Springer, 2016, pp. 21–37.
- [3] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," 2018.
- [4] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards realtime object detection with region proposal networks," in Advances in neural information processing systems, 2015, pp. 91–99.
- [5] S. Zhang, X. Zhu, Z. Lei, H. Shi, X. Wang, and S. Z. Li, "S³3fd: Single shot scale-invariant face detector," in *ICCV*, 2017, pp. 192–201.
- [6] W. Zhu, C. Liu, W. Fan, and X. Xie, "Deeplung: Deep 3d dual path nets for automated pulmonary nodule detection and classification," in WACV, 2018, pp. 673–681.
- [7] E. Maiettini, G. Pasquale, L. Rosasco, and L. Natale, "Interactive data collection for deep learning object detectors on humanoid robots," in *Humanoids*, 2017, pp. 862–868.
- [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 *ECCV*. Springer, 2014, pp. 740–755.
- [9] 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.
- [10] V. M. Patel, R. Gopalan, R. Li, and R. Chellappa, "Visual domain adaptation: A survey of recent advances," *IEEE signal processing magazine*, vol. 32, no. 3, pp. 53–69, 2015.
- [11] M. Ferguson, R. Ak, Y.-T. T. Lee, and K. H. Law, "Automatic localization of casting defects with convolutional neural networks," in *IEEE International Conference on Big Data*, 2017, pp. 1726–1735.
- [12] H. Maeda, Y. Sekimoto, T. Seto, T. Kashiyama, and H. Omata, "Road damage detection using deep neural networks with images captured through a smartphone," *arXiv preprint arXiv:1801.09454*, 2018.
- [13] G. Van Horn, O. Mac Aodha, Y. Song, Y. Cui, C. Sun, A. Shepard, H. Adam, P. Perona, and S. Belongie, "The inaturalist species classification and detection dataset," 2018.
- [14] A. R. Zamir, A. Sax, W. Shen, L. Guibas, J. Malik, and S. Savarese, "Taskonomy: Disentangling task transfer learning," in *CVPR*, 2018, pp. 3712–3722.
- [15] R. Girshick, I. Radosavovic, G. Gkioxari, P. Dollár, and K. He, "Detectron," https://github.com/facebookresearch/detectron/, 2018.
- [16] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama *et al.*, "Speed/accuracy trade-offs for modern convolutional object detectors," in *CVPR*, 2017, pp. 3296–3297.
- [21] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *ICCV*. IEEE, 2017, pp. 2980–2988.

- [17] 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 *IROS*, 2017, pp. 23–30.
- [18] J. Borrego, R. Figueiredo, A. Dehban, P. Moreno, A. Bernardino, and J. Santos-Victor, "A generic visual perception domain randomisation framework for gazebo," in *ICARSC*. IEEE, 2018, pp. 237–242.
- [19] N. Koenig and A. Howard, "Design and use paradigms for gazebo, an open-source multi-robot simulator," in *IROS*, vol. 3, 2004, pp. 2149– 2154.
- [20] W. Rawat and Z. Wang, "Deep convolutional neural networks for image classification: A comprehensive review," *Neural computation*, vol. 29, no. 9, pp. 2352–2449, 2017.
- [22] M. Johnson-Roberson, C. Barto, R. Mehta, S. N. Sridhar, K. Rosaen, and R. Vasudevan, "Driving in the matrix: Can virtual worlds replace human-generated annotations for real world tasks?" in *ICRA*. IEEE, 2017, pp. 746–753.
- [23] Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, and A. Farhadi, "Target-driven visual navigation in indoor scenes using deep reinforcement learning," in *ICRA*. IEEE, 2017, pp. 3357–3364.
- [24] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb, "Learning from simulated and unsupervised images through adversarial training," in *CVPR*, vol. 3, no. 4, 2017, p. 6.
- [25] G. Georgakis, M. A. Reza, A. Mousavian, P.-H. Le, and J. Košecká, "Multiview RGB-D dataset for object instance detection," in *International Conference on 3D Vision (3DV)*. IEEE, 2016, pp. 426–434.
- [26] G. Georgakis, M. A. Reza, A. Mousavian, P. Le, and J. Kosecka, "Multiview RGB-D dataset for object instance detection," *CoRR*, 2016. [Online]. Available: http://arxiv.org/abs/1609.07826
- [27] K. Lai, L. Bo, and D. Fox, "Unsupervised feature learning for 3d scene labeling," in 2014 IEEE International Conference on Robotics and Automation (ICRA), May 2014, pp. 3050–3057.
- [28] J. Johnson, B. Hariharan, L. van der Maaten, L. Fei-Fei, C. L. Zitnick, and R. Girshick, "Clevr: A diagnostic dataset for compositional language and elementary visual reasoning," in *CVPR*. IEEE, 2017, pp. 1988–1997.
- [29] Blender Online Community, Blender a 3D modelling and rendering package, Blender Foundation, Blender Institute, Amsterdam. [Online]. Available: http://www.blender.org
- [30] N. Jakobi, "Evolutionary robotics and the radical envelope-of-noise hypothesis," *Adaptive Behavior*, vol. 6, no. 2, pp. 325–368, 1997. [Online]. Available: https://doi.org/10.1177/105971239700600205
- [31] J. Tremblay, A. Prakash, D. Acuna, M. Brophy, V. Jampani, C. Anil, T. To, E. Cameracci, S. Boochoon, and S. Birchfield, "Training deep networks with synthetic data: Bridging the reality gap by domain randomization," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018, pp. 969–977.
- [32] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," *arXiv* preprint arXiv:1704.04861, 2017.
- [33] S. James, A. J. Davison, and E. Johns, "Transferring end-to-end visuomotor control from simulation to real world for a multi-stage task," in *CoRL*, 2017, pp. 334–343.
- [34] K. Perlin, "Improving noise," ACM Trans. Graph., vol. 21, no. 3, pp. 681–682, Jul. 2002. [Online]. Available: http://doi.acm.org/10.1145/ 566654.566636
- [35] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, "Adversarial discriminative domain adaptation," in CVPR, vol. 1, no. 2, 2017, p. 4.