

Synthetic Data for Robust Runway Detection

Estelle Chigot^{*1,2}, Dennis G. Wilson¹, Meriem Ghrib², Fabrice Jimenez², and Thomas Oberlin¹

¹Fédération ENAC ISAE-SUPAERO ONERA, Université de Toulouse, Toulouse, France, {estelle.chigot2, dennis.wilson, thomas.oberlin}@isae.fr

²Airbus, Toulouse, France, {meriem.ghrib, fabrice.jimenez}@airbus.com

Abstract

Deep vision models are now mature enough to be integrated in industrial and possibly critical applications such as autonomous navigation. Yet, data collection and labeling to train such models requires too much efforts and costs for a single company or product. This drawback is more significant in critical applications, where training data must include all possible conditions including rare scenarios. In this perspective, generating synthetic images is an appealing solution, since it allows a cheap yet reliable covering of all the conditions and environments, if the impact of the synthetic-to-real distribution shift is mitigated. In this article, we consider the case of runway detection that is a critical part in autonomous landing systems developed by aircraft manufacturers. We propose an image generation approach based on a commercial flight simulator that complements a few annotated real images. By controlling the image generation and the integration of real and synthetic data, we show that standard object detection models can achieve accurate prediction. We also evaluate their robustness with respect to adverse conditions, in our case nighttime images, that were not represented in the real data, and show the interest of using a customized domain adaptation strategy.

Keywords: Object detection; Synthetic data; Domain adaptation; Runway detection; Vision-based landing

1 Introduction

With the tremendous progress of artificial intelligence, deep vision models are increasingly integrated into critical systems. This paper considers the

^{*}Corresponding author

Accepted in Computer Analysis of Images and Patterns, September 2025 (CAIP 2025).

context of commercial aircraft, where there are major opportunities for both airlines and airports to increase safety and optimize routing or ground operations. In this work, we consider the task of runway detection from onboard cameras.

Deep learning models require large amounts of data to achieve high performance, but in industrial applications, data are often difficult or expensive to obtain. It may be difficult to reliably collect data, impossible to do so safely, or expensive to gather and label data. Aircraft operations make data gathering even more challenging: data collection may be heavily regulated and dangerous to perform outside of nominal conditions. On the other hand, flight simulators seem to provide the ideal solutions. They are cheap, customizable to include adverse conditions such as night, snow or fog, and, most of all, they can generate corresponding precise labels, thanks to the underlying scene modeling process. However, when using simulated data as the only source of data to train a detection model, a performance drop usually occurs when operating the model in real conditions, since the model cannot generalize to real images. This is known as the domain gap, synthetic-to-real gap, or sim-to-real problem. In this context, we study the reliability and performance gains from using synthetic data in a critical runway detection system.

The problem of runway recognition for aircraft has already been studied in the literature, either as a problem of object detection [27, 14, 11] or semantic segmentation [26]. There are now multiple publicly available datasets for runway detection [5, 2, 1]. However, these datasets usually include few to no real images, limiting the range of applicable methods and the evaluation on real use cases. In this study, we have access to a small private dataset of real images, and a large dataset of synthetic images; in both real and synthetic datasets, we have images representing different environmental conditions. We use these data to explore domain adaptation methods and to test models' robustness to unseen situations. In this work, we focus on runway detection at night as an underrepresented case, as we have access to nighttime images in both datasets, enabling us to evaluate the robustness of runway detection models on real images.

We make the following contributions, within the case of runway detection: (1) we demonstrate the benefits of using synthetic data to train a detection model in an industrial use case; (2) we customize a state-of-the-art domain adaptation method to the runway detection problem; (3) we show the advantages of using a domain adaptation approach to enhance model's robustness under adverse conditions, in our case nighttime.

The paper is structured as follows: Section 2 reviews the literature in the field of synthetic data for object detection. Section 3 details our data collection and the methods compared, while results are presented in Section 4. Section 5 concludes the study and draws some perspectives.

2 Related work

2.1 Object detection

Object detection is the computer vision task of determining the location and type of objects in a given image. Recent advances in deep learning have significantly improved object detection, especially in safety-critical industrial applications. One-stage detectors, such as YOLO [18] and SSD [15], prioritize speed over accuracy by directly predicting object location and class probabilities from the image. Two-stage models like Faster R-CNN [19] perform object detection in two steps. First, an image classification model extracts features from the image; this first model is often referred to as the "backbone." Then, a region proposal network (RPN) generates potential object detection proposals from the features. Generally, the backbone is pretrained on an image classification task, such as ImageNet [3], to improve the model's performance on the object detection task. In this work, we focus on Faster R-CNN due to its prevalence in the domain adaptation literature and its good performance.

Runway detection is critical for vision-based aircraft landing systems, in order to properly position the aircraft during the procedure. It has been treated as an object detection task in various studies. Faster R-CNN, in particular, has been used for the detection of runways [14]. A fusion of visual and infrared images has been proposed to improve runway detection at night, also using Faster R-CNN [27]. Federated learning, a distributed means of training models, was shown to improve the robustness of runway detection models [11]. One of the key challenges in the application of deep-learning based models for runway detection is the lack of sufficient real-world data, which synthetic data can help address.

2.2 Synthetic data

In order to train deep learning models and especially computer vision models, it is necessary to gather a large quantity of data. Large public datasets like COCO [13] or ImageNet [3] have accelerated object detection research, but their utility in industrial contexts, such as aviation, is limited. However, the cost, safety or legal concerns and time-consuming labeling process of real-world data collection make datasets difficult to build. Synthetic data are then a valuable alternative to create more data.

Two main paradigms exist for data generation. The first one, domain randomization [24, 25, 7], implies that photorealism doesn't matter as it tries to make data as diverse as possible. Simulators or 3D game engines are used to generate objects with large variations of environmental conditions, lighting or textures, improving model robustness across a wide range of scenarios. The aim is to train the model to perceive real images as just another variation of these conditions, leading to strong detection performance.

The second approach is to generate images as close to the real world as possible. This can be achieved by using computer graphics and simulators to build world replicas, showing realistic physics and real world behaviors. This method is very popular in the automotive industry, where video games like GTA-V [20, 8] or world models developed in 3D engines [21] have been used to generate widely utilized datasets. Additionally, whole driving simulators [4, 12] have been released for this purpose.

Generative models [6, 9, 23] are also a powerful way to create photo-realistic pictures. Generative adversarial networks (GANs) and variational autoencoders learn to approximate the data distribution of their training dataset, while diffusion models are trained to denoise step-by-step random noises into coherent data samples. GANs especially have been leveraged to generate data for industrial applications [16, 10, 29]. However, generative models need lots of real-world input data to produce satisfactory results. Also, they can't generate labels automatically. In this work we will focus on simulators for those two reasons.

In the aviation industry, due to the lack of real open data, datasets have been released to facilitate research in this field. LARD [5] provides a generator for runway detection, based on Google Earth, as well as 1800 real images taken from YouTube videos and manually labeled. On other tasks than runway detection, FS2020 [1] provides a synthetic dataset generated with Microsoft Flight Simulator, meant for runway segmentation and lines detection. Finally, Rareplanes [22] integrates 3D models of aircraft into real images for remote sensing applications. Logically, the works published for runway detection use mainly synthetic data. They either create their own with flight simulators [27], or use available open datasets [11] such as LARD. In a try to assess the benefits of synthetic images, Linden et al. [14] compare the impact on detection performance of several environmental conditions in a synthetic dataset and examines a style transfer approach on automotive data.

2.3 Synthetic to real domain adaptation

Domain adaptation (DA) methods attempt to minimize the synthetic to real gap by bringing closer either the images directly or the feature distributions between domains. Following the first option, SC-UDA [28] uses a neural style transfer method to apply a realistic style on synthetic images, and then uses Faster R-CNN to generate bounding boxes to augment the unlabeled real dataset. On the other hand, CARE [17] tries to align the instances features between synthetic and real images by modifying the traditional Faster R-CNN loss function. Using the available labels, the loss term introduced by CARE computes the difference between the features of a synthetic object and a real one, and tries to minimize it along the other object detection loss terms. CARE also introduces two reweighting terms to adapt to the

difference of object sizes and frequency between the two domains.

In the context of runway detection, DA remains underexplored, likely due to the scarcity of real-world data. In this work we focus on studying DA methods to enhance model generalization from synthetic to real images.

3 Methodology

3.1 Data collection

In our use case, real images are limited in quantity due to their high cost of acquisition. The images are captured using an aircraft equipped with cameras and sensors, navigating across various airports in the U.S. Data is acquired by photographing runways at selected airports during landing sequences. In this study, we have access to a database of real images from 27 airports with their associated metadata. However, the images of the same airports tend to be quite similar, showing the same runways from close points during the flight phases, resulting in limited diversity within the real data. Labeling is performed automatically based on the GPS positions of the aircraft and the runways, leading to potential errors in labels due to factors like miscalibration or cloud occlusions.

In contrast, we generate our own synthetic dataset using the commercial flight simulator XPlane12, which allows us to specify the aircraft’s position, orientation, and various environmental conditions such as weather and time. This customization capacity enables us to collect diverse images from 207 airports worldwide. One of the key benefits of synthetic images is the ability to achieve automatic and accurate labeling of bounding boxes and other parameters that typically require human input in real-world scenarios. We use a custom pipeline rather than relying on existing datasets because it enables the generation of a large volume ($>5,000$) of high-quality images, across a wide variety of airports and under specific environmental conditions.

Taking advantage of this generation capacity, we want to trial synthetic data and assess if using simulated images enables a model to be robust under unseen conditions. To do so, we use the nighttime scenario. At night, runways are very different than during the day. The light conditions differ a lot as they are outlined by lights, in order to enable the pilot to see the runway. The shape is conserved but the texture and colors of the image are not, which is a good way to assess the generalization capabilities of a model. With this idea in mind, we generate nighttime synthetic data from the same airports as the daytime ones.


In Figure 1 we have examples of real and synthetic images for the runway detection task. For copyright reasons, we show synthetic images from [1] which are not the one used to train the model in this study, but closely resemble.



(a) Real runway (Source: Airbus)



(b) Synthetic runway (Source: [1])



(c) Synthetic runway at night (Source: [1])

Figure 1: Examples of real data and synthetic data for runway detection.

3.2 Mixing strategies

In this work, our objective is to assess the use of synthetic data for runway detection. We also want to study the impact of synthetic data in unseen situations. To this end we design several training strategies, applicable for both the nominal and the nighttime studies.

We explore the utility of synthetic data by employing various configurations of real and synthetic datasets. Our study begins with a baseline experiment using only real images of runways (REAL), allowing us to establish a performance benchmark based solely on real data.

Next, we create a synthetic dataset using a flight simulator software (SYNTH). For all datasets, we ensure that each airport is represented by the same number of images, preventing imbalances in both training and validation sets.

To integrate both real and synthetic data during training, we assess different mixing strategies. The standard approach involves randomly sampling from the combined dataset (MIX), resulting in a variable ratio of synthetic to real data at each training step, dependent on the dataset composition.

To mitigate the domain shift between synthetic and real data, we also examine the CARE [17] domain adaptation method (CARE) tailored to our use case. This method introduces an additional loss term during training, encouraging the detection model to represent real and synthetic objects similarly by minimizing the Euclidean distance between their feature maps.

Furthermore, the CARE framework employs a specific sampling strategy that ensures an equal number of synthetic and real images are used at each

training step in the minibatch. While this sampling method is necessary to their loss terms, we also investigate its effectiveness independently, without the accompanying loss terms (SAMPLER).

3.3 Experimental setup

Model overview We consider a standard object detector, Faster R-CNN [19], to follow the recent literature in domain adaptation for object detection [17, 28] and previous works in airport runway detection [27, 14, 11].

The objective function used in this study, which the model tries to minimize, is the standard Faster R-CNN loss function, ℓ_{F-RCNN} .

In the experiments using CARE, we add to the standard loss the alignment term proposed. CARE also incorporates reweighting factors that are unnecessary in our monoclase setting where object size and aspect ratio are consistent. We do not include them in our loss function. The overall training loss then becomes:

$$\ell_{det} = \ell_{F-RCNN} + \lambda \ell_{align}. \quad (1)$$

λ is a balancing factor, which controls the contribution of the alignment loss in the overall loss. ℓ_{align} is the alignment loss proposed by CARE, a cycle loss computing the difference between features of synthetic and real objects.

Implementation details For all the experiments using the real dataset, our data being redundant, we use a small training dataset of 1,000 images from 19 airports. Then, for the simulated dataset we generate 10,000 images to keep a reasonable ratio of synthetic to real images. Those images are from 199 airports, including the ones in the real dataset.

In order to study the robustness to domain shift of each model, we introduce the nighttime condition. For this purpose, we generate another synthetic dataset composed of 5,000 daytime images and 5,000 nighttime images from the same 199 airports.

Our detection model is a Faster R-CNN with FPN and a ResNet50 backbone, pretrained on COCO2017 [13]. The training configuration includes a batch size of 8, a learning rate of 0.002, and is trained during 10,000 iterations. We use $\lambda = 0.1$ for the CARE loss (Equation (1)). We perform the validation at the end of the training.

Evaluation Following previous work [28, 14] we report COCO Average Precision (AP) for each strategy, AP being the average precision over several Intersection over Union values, from 0.5 to 0.95 by 0.05 steps.

To assess the training methods and the impact of training datasets, we build an evaluation dataset of 200 real daytime images from 8 airports. Those

	REAL	SYNTH	MIX	SAMPLER	CARE
DAY AP	58.60	59.04	65.25	64.01	63.66

Table 1: Values of AP for each mixing strategy over daytime validation datasets. Models were trained *without* synthetic nighttime image.

	REAL	SYNTH	MIX	SAMPLER	CARE
DAY AP	58.60	57.55	66.77	64.73	67.50
NIGHT AP	15.21	43.22	43.38	42.00	43.75

Table 2: Values of AP for each mixing strategy over daytime and nighttime validation datasets. Models were trained *with* synthetic nighttime image.

airports are not represented in any training set. In order to evaluate robustness to the nighttime condition, we build an additional validation dataset of 200 real nighttime images from the same 8 airports. Real nighttime images are never seen during training.

4 Results

4.1 Comparison of training strategies

In this section, we evaluate models trained on daytime images only. Table 1 compares the results of the different training strategies and datasets on the daytime validation set. We can observe that all methods outperform the model trained on real data, including the synthetic only experiment (+0,44%). We could explain this behavior with the model’s pretraining; having already seen real images in COCO2017, it acquired features relevant to real images before the synthetic training. In this setting, adding synthetic data into the training dataset seems to bring better performance on the validation dataset, and a simple method such as MIX seems to bring the best results (+6.65% compared to REAL only). This conclusion was expected, as bringing more data with reliable labeling should lead to a better model.

4.2 Results on night domain shift

We analyze the results of models integrating nighttime images in their synthetic training set, compared to the same models trained on daytime images only.

In Table 2, we first compare the models performance on the daytime validation dataset. The REAL model keeps the same training set throughout all experiments and therefore gets the same results. The synthetic only model SYNTH suffers from the addition of nighttime condition. On the contrary, while we expected the same results from the other training strategies,




Figure 2: A real image of an airport runway at night from the validation dataset, and the bounding box determined by inference through each model. The REAL model does not predict any bounding box, while all other models correctly predict the bounding box covering the runway.

but MIX (+8,17%) and especially CARE (+8,90%) seem to benefit from the diversification of conditions. On average SAMPLER attains the same results, showing little use of the nighttime images. Overall all mixing methods still outperform the REAL experiment. The CARE method exhibits the best score showing the benefits of this domain adaptation strategy when diversifying the training dataset.

We also evaluate our models over the nighttime validation dataset. The REAL model experiences a significant drop in performance (-43,39%), and is unable to generalize properly on this unseen condition. All the other methods also experience a drop in performance. However, they are all able to maintain an AP score similar to the SYNTH model (-14,33% compared to daytime validation), way higher than the REAL model on both validation set. In this scenario, adding synthetic data seems to guarantee a correct level of performance, compared to a completely unseen condition. Our customized CARE method still exhibits the best score in this difficult scenario by a small margin.

In Figure 2 we show inference results for all models on one validation image. This picture example is representative of the type of nighttime images we can find in the dataset. The REAL model is unable to see the runway whereas all the models trained on synthetic data could detect it.

Overall, mixing synthetic and real data seems to lead to more robustness to unseen environmental conditions, as we obtain good AP scores in both nominal and adverse conditions. Adding the target data into the synthetic datasets helps getting reasonable performance in difficult conditions. In our

case, domain adaptation seems to improve performance with the customized loss of CARE allowing the model to generalize better. While the sampling method introduced in CARE seems to lower the detection performance, its alignment loss compensates this drop and improves the final results when getting enough diversity in the training dataset.

5 Conclusion and perspectives

In this work, we studied the use of synthetic data for runway detection, specifically how data sampling schemes and domain adaptation can improve model performance. We find that, in general, synthetic data improves model performance and robustness in adapting to rare conditions.

In our first study on the nominal case of regular daytime conditions we find that synthetic images can improve model accuracy, but the presence of both real and synthetic images is necessary.

Synthetic data can be especially valuable to represent rare cases that are underrepresented in real data, in order to ensure model robustness. We found that synthetic images of nighttime conditions can greatly improve model performance on runway detection at night. Even training a model on only synthetic data vastly outperforms models trained only on real data. We consider that this is due to the lack of real nighttime images and the difficulty of adapting runway detection from daytime to nighttime in the absence of nighttime examples.

Domain adaptation seems to be a promising approach to enhance runway detection models. While the benefits of CARE compared to a simple mixing approach seem modest, the customized loss helps overcoming the decline introduced by the sampler, and CARE achieves the best performance overall.

In a future direction we aim to tackle other adverse conditions, such as snowy or foggy weathers that are more complicated to capture in reality or render in simulation. It would also be of interest to explore the impact of the CARE sampler, to further improve the benefits of this domain adaptation approach.

In summary, we show that synthetic data can aid in the task of airport runway recognition. Images from simulation software are especially useful in the representation of rare cases, such as nighttime data. By including synthetic data from diverse cases, we can ensure and improve computer vision model robustness, especially with the help of our customized CARE domain adaptation strategy.

Disclosure of Interests Estelle Chigot, Meriem Ghrib and Fabrice Jimenez are employees of the Airbus company. Dennis G. Wilson and Thomas Oberlin have no competing interests to declare that are relevant to the content of this article.

Acknowledgement This preprint has not undergone peer review or any post-submission improvements or corrections. The Version of Record of this contribution is published in Lecture Notes in Computer Science, and is available online at https://doi.org/10.1007/978-3-032-04968-1_25.

References

- [1] Chen, M., Hu, Y.: An image-based runway detection method for fixed-wing aircraft based on deep neural network. IET Image Processing (2024), <https://www.kaggle.com/datasets/relufrank/fs2020-runway-dataset/data>
- [2] Chen, W., Zhang, Z., Yu, L., Tai, Y.: BARS: a benchmark for airport runway segmentation. Applied Intelligence **53**(17), 20485–20498 (2023)
- [3] Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. pp. 248–255 (2009)
- [4] Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., Koltun, V.: CARLA: An Open Urban Driving Simulator. In: Proceedings of the 1st Annual Conference on Robot Learning. pp. 1–16. PMLR (2017)
- [5] Ducoffe, M., Carrere, M., Féliers, L., Gauffriaux, A., Mussot, V., Pagetti, C., Sammour, T.: LARD–Landing Approach Runway Detection–Dataset for Vision Based Landing. arXiv preprint arXiv:2304.09938 (2023)
- [6] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2. vol. 27 (2014)
- [7] Hinterstoisser, S., Pauly, O., Heibel, H., Martina, M., Bokeloh, M.: An Annotation Saved is an Annotation Earned: Using Fully Synthetic Training for Object Detection. In: 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW) (2019)
- [8] Johnson-Roberson, M., Barto, C., Mehta, R., Sridhar, S.N., Rosaen, K., Vasudevan, R.: Driving in the Matrix: Can virtual worlds replace human-generated annotations for real world tasks? In: 2017 IEEE International Conference on Robotics and Automation (ICRA) (2017)
- [9] Kingma, D.P., Welling, M.: Auto-Encoding Variational Bayes. In: International Conference on Learning Representations (2014)

- [10] Li, B., Zou, Y., Zhu, R., Yao, W., Wang, J., Wan, S.: Fabric Defect Segmentation System Based on a Lightweight GAN for Industrial Internet of Things. *Wireless Communications and Mobile Computing* **2022**(1), 9680519 (2022)
- [11] Li, Y., Angelov, P., Yu, Z., Lopez Pellicer, A., Suri, N.: Federated Adversarial Learning for Robust Autonomous Landing Runway Detection. In: *International Conference on Artificial Neural Networks*. pp. 159–173 (2024)
- [12] Lin, L., Liu, Y., Hu, Y., Yan, X., Xie, K., Huang, H.: Capturing, Reconstructing, and Simulating: The UrbanScene3D Dataset. In: *Computer Vision – ECCV 2022*. pp. 93–109 (2022)
- [13] Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common Objects in Context. In: *Computer Vision – ECCV 2014*. pp. 740–755 (2014)
- [14] Lindén, J., Forsberg, H., Haddad, J., Tagebrand, E., Cedernaes, E., Ek, E.G., Daneshtalab, M.: Curating Datasets for Visual Runway Detection. In: *IEEE/AIAA 40th Digital Avionics Systems Conference (DASC)*. pp. 1–9 (2021)
- [15] Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: SSD: Single Shot MultiBox Detector. In: *Computer Vision – ECCV 2016*. pp. 21–37 (2016)
- [16] Niu, S., Li, B., Wang, X., Lin, H.: Defect Image Sample Generation With GAN for Improving Defect Recognition. *IEEE Transactions on Automation Science and Engineering* **17**(3), 1611–1622 (2020)
- [17] Prabhu, V.U., Acuna, D., Mahmood, R., Law, M.T., Liao, Y.H., Hoffman, J., Fidler, S., Lucas, J.: Bridging the Sim2Real gap with CARE: Supervised Detection Adaptation with Conditional Alignment and Reweighting. *Transactions on Machine Learning Research* (2023)
- [18] Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You Only Look Once: Unified, Real-Time Object Detection. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 779–788 (2016)
- [19] Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **39**(6), 1137–1149 (2017)

- [20] Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: *Computer Vision – ECCV 2016*. pp. 102–118 (2016)
- [21] Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The SYNTHIA Dataset: A Large Collection of Synthetic Images for Semantic Segmentation of Urban Scenes. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 3234–3243 (2016)
- [22] Shermeyer, J., Hossler, T., Etten, A.V., Hogan, D., Lewis, R., Kim, D.: RarePlanes: Synthetic Data Takes Flight. In: *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*. pp. 207–217 (2021)
- [23] Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: *Proceedings of the 32nd International Conference on Machine Learning - Volume 37*. pp. 2256–2265 (2015)
- [24] Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., Abbeel, P.: Domain randomization for transferring deep neural networks from simulation to the real world. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 23–30 (2017)
- [25] Tremblay, J., Prakash, A., Acuna, D., Brophy, M., Jampani, V., Anil, C., To, T., Cameracci, E., Boochoon, S., Birchfield, S.: Training Deep Networks with Synthetic Data: Bridging the Reality Gap by Domain Randomization. In: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. pp. 969–977 (2018)
- [26] Wang, Q., Feng, W., Zhao, H., Liu, B., Lyu, S.: VALNet: Vision-Based Autonomous Landing with Airport Runway Instance Segmentation. *Remote Sensing* **16**(12), 2161 (2024)
- [27] Wang, Z., Zhao, D., Cao, Y.: Visual Navigation Algorithm for Night Landing of Fixed-Wing Unmanned Aerial Vehicle. *Aerospace* **9**(10), 615 (2022)
- [28] Yu, F., Wang, D., Chen, Y., Karianakis, N., Shen, T., Yu, P., Lymberopoulos, D., Lu, S., Shi, W., Chen, X.: SC-UDA: Style and Content Gaps aware Unsupervised Domain Adaptation for Object Detection. In: *2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. pp. 382–391 (2022)
- [29] Zhang, L., Dai, Y., Fan, F., He, C.: Anomaly Detection of GAN Industrial Image Based on Attention Feature Fusion. *Sensors* **23**(1), 355 (2022)