

Generative Adversarial Networks for Data Augmentation in Structural Adhesive Inspection

Ricardo Silva Peres ^{1,2,*} , Miguel Azevedo ¹, Sara Oleiro Araújo ^{1,2} , Magno Guedes ³  and Fábio Miranda ³ and José Barata ^{1,2} 

¹ Electrical Engineering Department, School of Science and Technology, NOVA University of Lisbon, 2829-516 Caparica, Portugal; mn.azevedo@campus.fct.unl.pt (M.A.); s.araujo@uninova.pt (S.O.A.); jab@uninova.pt (J.B.)

² UNINOVA-Centre of Technology and Systems (CTS), FCT Campus, 2829-516 Caparica, Portugal

³ Introsys S.A., Estrada dos 4 Castelos 67, 2950-805 Quinta do Anjo, Portugal; magno.guedes@introsys.eu (M.G.); fabio.miranda@introsys.eu (F.M.)

* Correspondence: ricardo.peres@uninova.pt

Abstract: The technological advances brought forth by the Industry 4.0 paradigm have renewed the disruptive potential of artificial intelligence in the manufacturing sector, building the data-driven era on top of concepts such as Cyber-Physical Systems and the Internet of Things. However, data availability remains a major challenge for the success of these solutions, particularly concerning those based on deep learning approaches. Specifically in the quality inspection of structural adhesive applications, found commonly in the automotive domain, defect data with sufficient variety, volume and quality is generally costly, time-consuming and inefficient to obtain, jeopardizing the viability of such approaches due to data scarcity. To mitigate this, we propose a novel approach to generate synthetic training data for this application, leveraging recent breakthroughs in training generative adversarial networks with limited data to improve the performance of automated inspection methods based on deep learning, especially for imbalanced datasets. Preliminary results in a real automotive pilot cell show promise in this direction, with the approach being able to generate realistic adhesive bead images and consequently object detection models showing improved mean average precision at different thresholds when trained on the augmented dataset. For reproducibility purposes, the model weights, configurations and data encompassed in this study are made publicly available.

Keywords: deep learning; data augmentation; quality control; industry 4.0; structural adhesive



Citation: Peres, R.S.; Azevedo, M.; Araújo, S.O.; Guedes, M.; Miranda, F.; Barata, J. Generative Adversarial Networks for Data Augmentation in Structural Adhesive Inspection. *Appl. Sci.* **2021**, *11*, 3086. <https://doi.org/10.3390/app11073086>

Academic Editor: Radu Godina

Received: 9 March 2021

Accepted: 29 March 2021

Published: 30 March 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The timely inspection of parts and components along the production line is critical to ensure that each part adheres to strict quality criteria necessary to guarantee the safety of the product's end-user, particularly in sectors such as aerospace, naval and automotive. Regarding the latter, the adoption of lighter and robust materials has made it so that some parts cannot be welded, with structural adhesives playing an important role as an alternative that contributes to the reduction of noise, vibrations and infiltrations.

However, the inspection of parts bonded with this method typically involves destructive tests that require the separation of bonded parts as a way to enable the analysis of the continuity, spread and consistency of the adhesive. This makes such inspections difficult, being both time-consuming and costly in terms of resources, materials and waste. Moreover, common defects like discontinuities and blobs are impossible to correct if not detect prior to the bonding process, since parts undergo considerable mechanical and structural changes.

While the recent advances in Artificial Intelligence (AI), particularly concerning deep learning, and the Industry 4.0 paradigm have made significant progress towards the automation of in-line quality inspections across varied domains [1–4], the destructive nature

and relatively low frequency of these tests make it so that data availability remains a tremendous challenge for the industrial viability and adoption of such solutions. Even in cases for which defect data can be made available, it is difficult to ensure that a balanced number of samples of each defect type is included.

Consequently, one of the main barriers in modern deep learning-based approaches concerning quality control is the vast amount of training data necessary to develop these solutions. Large datasets require considerable human effort to generate from scratch, whilst generally being costly, time-consuming and error-prone. To mitigate this issue, data augmentation with synthetic data is emerging as a possible solution to decrease the burden of data collection and annotation [5].

Recently, a systematic review [6] of the current Industrial AI landscape has shown that the usage of synthetic data with the purpose of artificially augmenting datasets in industrial applications was present in around 20% of publications included in the study. With the important role that realistic and high-resolution synthetic images can play in industrial computer vision tasks, Generative Adversarial Networks (GANs) [7] are becoming more and more attractive as a way to reliably generate additional samples in the manufacturing domain.

While generative methods have been prevalent in recent literature as a way to generate novel samples from high-dimensional data distributions (for instance in the case of images), previous methods such as those based on variational autoencoders [8] tended to produce blurry results due to restrictions in the model. GANs are generally capable of producing sharper images [9] with recent advances making it possible to generate synthetic samples with increasingly large levels of quality [10,11].

Some examples of GANs in the manufacturing domain can be found in the literature, addressing applications in which the data are scarce and difficult to obtain [12–14]. Other relevant examples include efforts addressing cases in which labeling the entire dataset may be unfeasible, with semi-supervised learning methods showing promise in this direction [15–17].

While a large amount of real training data is still necessary in order to train GAN models to generate samples with sufficient quality, researchers at NVIDIA have recently proposed StyleGAN2-ADA [18], which makes considerable progress towards the successful training of GANs with limited data. This was made possible by employing an Adaptive Discriminator Augmentation (ADA) mechanism that stabilizes training in limited data regimes, thus avoiding typical issues of overfitting the discriminator.

Based on this, we propose a novel approach leveraging these recent advances in GAN training to generate synthetic data as a way to augment scarce training sets in manufacturing quality control tasks, specifically regarding structural adhesive applications. Additionally, we demonstrate that such an approach can improve the viability and performance of automated inspections based on deep learning. We carried out preliminary testing using a pilot application cell to generate real images of defective beads, creating a dataset to train this generative model for data augmentation. Using the augmented training set, we trained a state-of-the-art object detection model to automate the quality inspection task, then compared its performance with the same model trained only on a real training set of limited size.

This directly addresses the future research direction delineated in previous work [19], which performed data augmentation using simulation. Shifting to a GAN-based method facilitates the employment of synthetic data generation without requiring the specific modeling of a simulation that closely resembles the real application. This is especially important in cases for which modeling may be too complex or unfeasible.

The remainder of this article is organized as follows: Section 2 describes the materials and methods adopted in this work, providing sufficient detail and references for its replication. Section 3 summarizes the preliminary results along with a brief interpretation, followed by a more thorough discussion of their implication and directions for future work

on this topic in Section 4. Finally, Section 5 summarizes the conclusions and provides some closing remarks.

2. Materials and Methods

This section briefly addresses the methods necessary for the replication of the proposed approach. Firstly, the generation of the real dataset and its characterization are discussed. Then, the generation of the synthetic adhesive images using StyleGAN2-ADA is presented. Finally, the implementation, training and validation steps for the YOLOv4-Tiny object detection model are discussed. The model weights, configurations and data encompassed in this study are made publicly available at https://github.com/RicardoSPeres/GAN_Synth_Adhesive (accessed on 30 March 2021).

2.1. Dataset Characterization

The dataset considered in this study consists of images collected from a structural adhesive application cell located at Introsys S.A. facilities in Castelo Branco, Portugal, a company specializing in industrial automation (particularly in the automotive sector) that operates in the international market since 2004. It encompasses two stations, one for the adhesive application carried out by an ABB IRB 2400 robotic arm, and the other for the visual quality inspection with two Teledine cameras, as depicted in Figure 1. Additional descriptions of the cell are available in [19,20]. For this study two types of defects are considered based on the requirements of the use case at hand, namely discontinuities and blobs (excess of material). Still, the approach can easily accommodate additional defect types, so long as they are identified during the labeling process. To generate the original dataset, the controller was manually reconfigured by an operator for each product to produce different defects of these types at varying positions during the adhesive application.

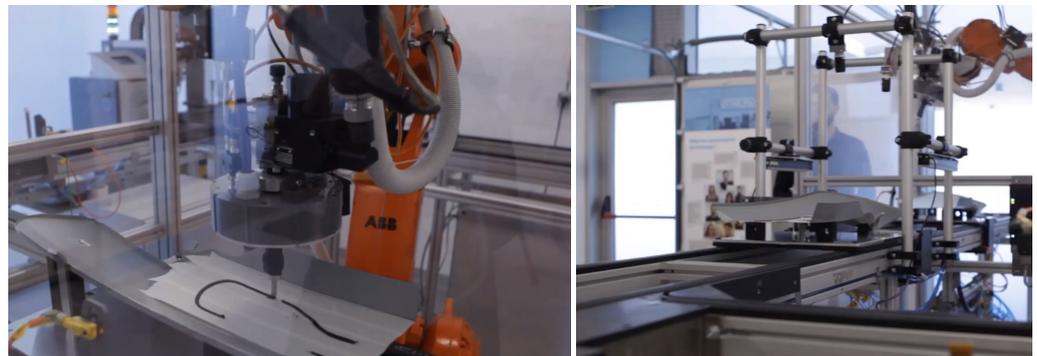


Figure 1. (left) The structural adhesive application station showing a discontinuity defect. (right) The quality inspection station downstream.

For the training set of 143 real images used in the GAN model, the original images were cropped to a resolution of 1024×1024 to match the input format of StyleGAN2-ADA, namely square-shaped resolution and the same power-of-two dimensions, and centered on the adhesive bead. The resulting model was then used to generate 536 synthetic images of defective beads to augment the original dataset.

Regarding the object detection model to automate quality inspection, 116 real images from the original dataset were selected (removing those that had no defects of the contemplated types), with this set being split 50/50 into training and validation sets. The initial training set of 58 images was then augmented with the synthetic set, resulting in a training set of 594 1024×1024 images, which were manually annotated using LabelImg (<https://github.com/tzutalin/labelimg>, accessed on 9 March 2021). The real training set is imbalanced, consisting of 75 instances of discontinuity defects and 34 of excess defects, since the former occur more frequently in the present use case. The augmentation addresses the class imbalance, resulting in 582 instances of discontinuities and 580 of excess defects.

2.2. Generating Synthetic Adhesive Images

GANs are an emerging trend in synthetic image generation that is founded on the idea of training a pair of neural networks in competition with one another [21]. A general depiction of a basic GAN structure is provided in Figure 2.

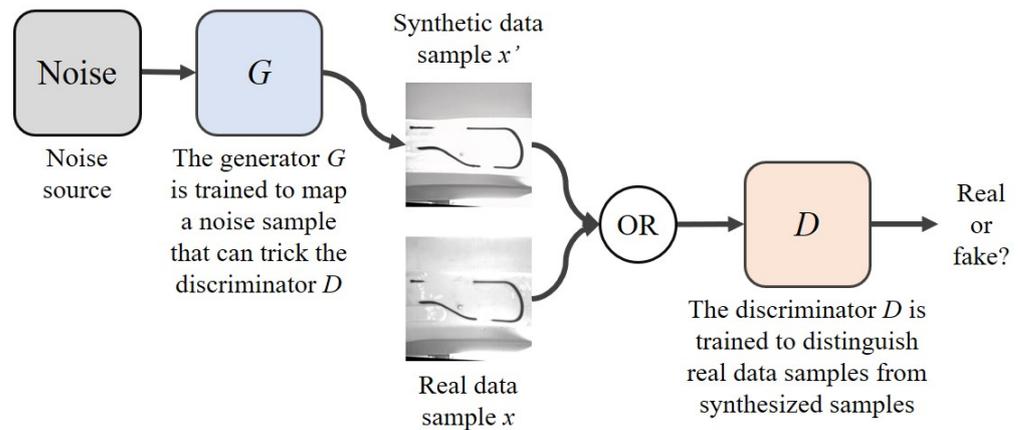


Figure 2. Basic structure of a Generative Adversarial Network (GAN). The generator G and discriminator D are generally implemented by multi-layer networks consisting of convolutional and/or fully-connected layers. Adapted from [21].

During this simultaneous training process, the error signal to the discriminator is provided from the ground truth of whether or not a sample was real or synthetic. Since the generator has no direct access to real images, the error signal can be used to train the generator (via the discriminator), enabling it to generate synthetic images with improved quality.

However, it remains challenging to collect datasets that are large enough to meet the requirements to train modern high-resolution GANs. This is particularly true in domains such as manufacturing quality control, where manufacturers generally optimize against the occurrence of defects. In this light, the recent introduction of StyleGAN2-ADA has made it possible to use much smaller datasets than before for this purpose.

Hence, for the generation of synthetic images, the base implementation of StyleGAN2-ADA available at <https://github.com/NVLabs/stylegan2-ada> (accessed on 9 March 2021) was adopted. To speed up convergence and reduce data requirements, transfer learning was used to start from a model pre-trained on the FFHQ 1024×1024 dataset [10], as opposed to random initialization. As stated in [18], transfer learning provides significantly better results than from-scratch training, with its success apparently depending mainly on the diversity of the source dataset, instead of the similarity between subjects.

Since datasets are stored as multi-resolution TFRecords, the first step was to convert the structural adhesive dataset to the correct format. Then, training was carried out on a single NVIDIA Tesla V100 GPU, taking approximately eight hours. For inference on the same setup, synthetic images could be generated at an average rate of 105 images per minute, using different seeds and varying truncation values.

2.3. Object Detection with Synthetic Data

Concerning object detection models, Scaled-YOLOv4 [22] has recently achieved state-of-the-art performance on the Microsoft Common Objects in COntext (MS COCO) dataset [23]. However, to enable a quicker iteration of experiments, the YOLOv4-Tiny model [24] was chosen for the validation of this approach due to shorter training times. Training was also carried out in a single NVIDIA Tesla V100 GPU using Darknet [25], an open source neural network framework written in C and CUDA supporting CPU and GPU computation. The base implementation is available at <https://github.com/AlexeyAB/darknet> (accessed on 9 March 2021).

To assess the impact of the proposed augmentation approach, two distinct models were trained over 6000 iterations each, one using only the real training set of 58 structural adhesive bead images, and the other using the full augmented set of 594 images. Afterwards, the mean Average Precision (mAP) of each model was calculated at different Intersection over Union (IoU) thresholds, first for the validation set of 58 real images, then for a second holdout set of 19 images generated in a separate experiment. This metric was chosen as it common in most modern object detection tasks [23,26]. The holdout set included images with defects occurring in portions of the bead that were not common (or present at all) in the original dataset, representing a more difficult validation scenario.

3. Preliminary Results

This section provides a brief description of the preliminary experimental results, both for the synthetic data generation and the object detection, along with some interpretations of their implication in the context of this work.

3.1. Synthetic Image Generation

Regarding the generation of synthetic structural adhesive defects, different seed and truncation values were used to generate a variety of images, with some examples shown in Figure 3 where the first column presents real images, while all others are generated by the trained GAN.

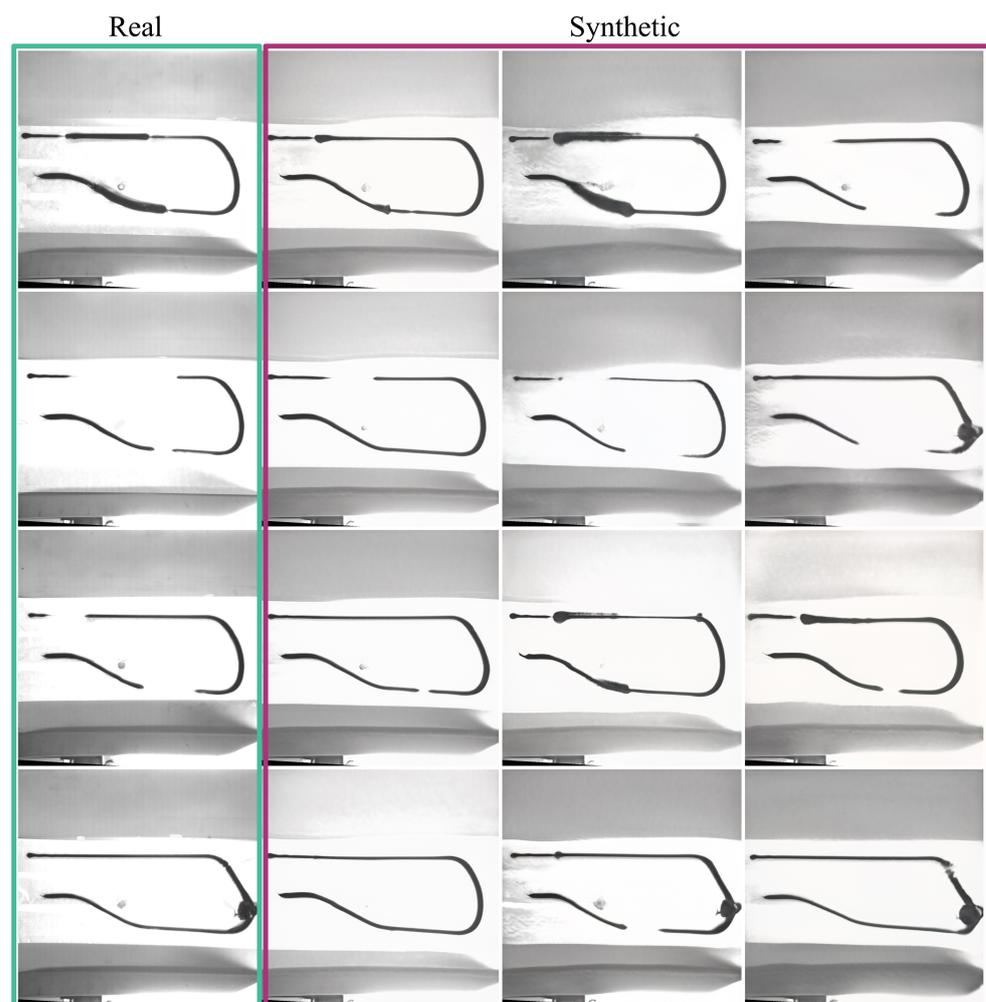


Figure 3. Real and synthetic images of defective structural adhesive beads. First column depicts real images from the original dataset. Second, third and fourth columns correspond to synthetic images generated with truncation values of 0.5, 1.5 and 2.0, respectively.

It can be observed that not only is the model capable of generating realistic images of different beads, but also that varied defects can be generated along different segments of the bead, both for discontinuity and blob defects.

3.2. Automated Defect Detection Using Deep Learning

Following the method specified in Section 2.3, one of the YOLOv4-Tiny models was trained on the real training set and the other on the augmented one, as described in Section 2.1. To further facilitate the comparison, a third model was trained only on the synthetic dataset, with a fourth model being trained on a dataset augmented with a previous alternative approach based on simulation [19]. A summary of the results from this process is provided in Table 1.

Table 1. Evaluation metrics for the models trained with synthetic set (536 images), the scarce real set (58 images) and the augmented sets (594 images for both). Models were first validated on a set of 58 real images, then on a holdout set of 19 real images. The number of instances of each defect is listed in the table. The mean Average Precision (mAP) is computed for different Intersection over Union (IoU) thresholds, namely 0.15, 0.30 and 0.50. Best results are highlighted in bold.

Dataset	mAP@0.15	mAP@0.30	mAP@0.50
Validation Set (60 discontinuity, 60 excess)			
Synthetic	0.8141	0.7724	0.6755
Real	0.8662	0.8239	0.7221
Augmented (Simulation)	0.8870	0.8265	0.7339
Augmented (GAN)	0.9131	0.8569	0.7708
Holdout Set (40 discontinuity, 17 excess)			
Synthetic	0.5969	0.5154	0.3344
Real	0.5986	0.5249	0.4125
Augmented (Simulation)	0.6631	0.6189	0.4406
Augmented (GAN)	0.7675	0.7256	0.5788

In both cases, the model trained on the training set augmented using the proposed approach attained superior results. Furthermore, the impact of the augmentation becomes even more evident when testing the model with the holdout set. While the validation set contains images with defects that are similar to those present in the real training set, the ones held out are considerably different, for instance with excess defects occurring in parts of the bead that were previously unseen (in both GAN and object detector training sets). Several examples are showcased in Figure 4, in which the first column (left) depicts results from the model trained only on real data, while the second column (right) shows the detections of the augmented model on the same image.

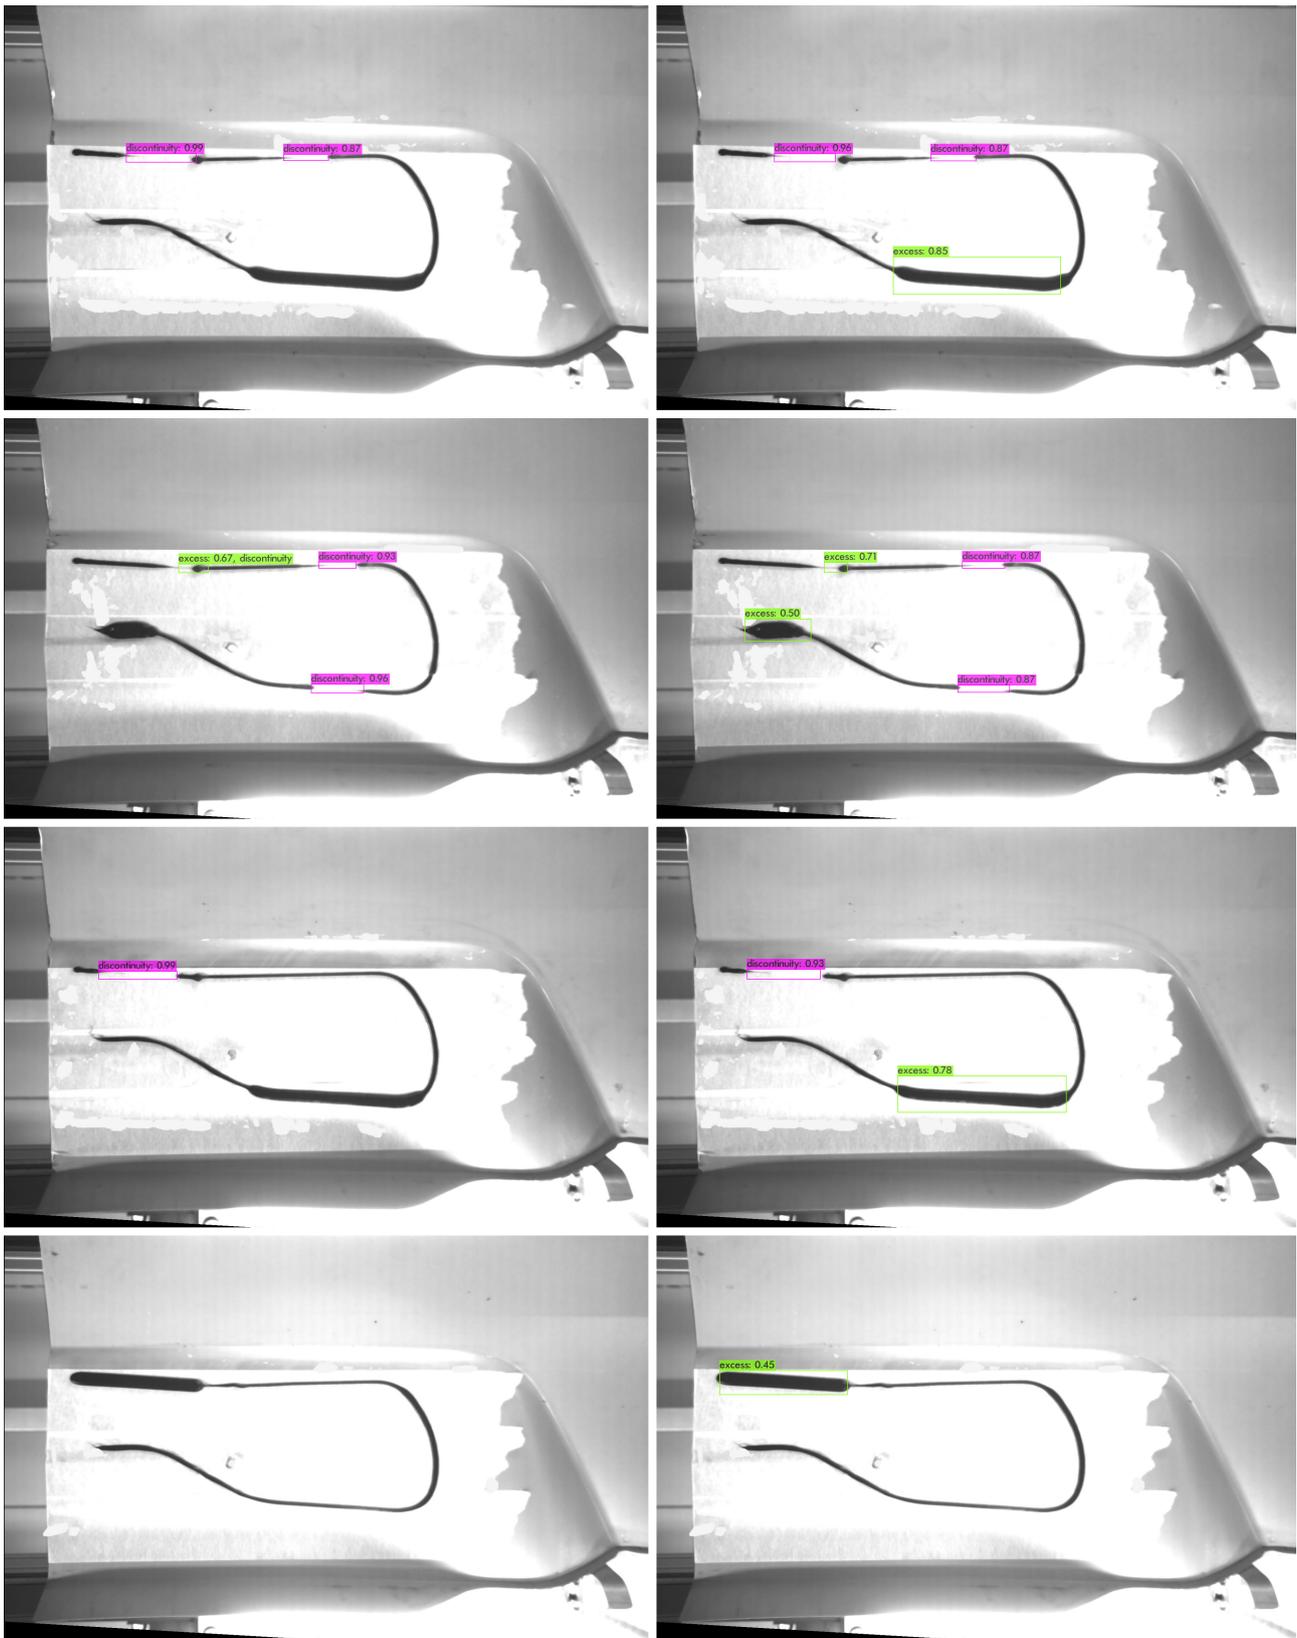


Figure 4. Examples from the output of the YOLOv4-Tiny models running inference on the holdout set. Left column depicts the output from the model trained only on real data, while the right column represents the corresponding result from the model trained on the augmented set.

In addition to this, given the imbalanced nature of the original (real) dataset, it is interesting to take a closer look at the Average Precision (AP) for each class (discontinuity and excess) between the different training sets. These results are presented in Table 2.

Table 2. Average Precision (AP) per defect class with an IoU threshold of 50%, validated first on a set of 58 real images, then on a holdout set of 19 real images. The number of instances of each defect is listed in the table.

Dataset	AP Discontinuity	AP Excess
Validation Set (60 discontinuity, 60 excess)		
Synthetic	0.7577	0.5934
Real	0.7735	0.6707
Augmented (Simulation)	0.7720	0.6958
Augmented (GAN)	0.7933	0.7483
Holdout Set (40 discontinuity, 17 excess)		
Synthetic	0.4404	0.2284
Real	0.5421	0.2829
Augmented (Simulation)	0.5812	0.3001
Augmented (GAN)	0.5711	0.5864

From the analysis of Table 2 it becomes evident that there is a clear difference in AP between the classes for the model trained only on the real imbalanced dataset. More specifically, the difference amounts to ~10% in the validation set, and ~26% in the holdout set. The considerably larger gap in the latter can be explained once more due to the more pronounced difference between the defects observed in the holdout set and those of the initial dataset. Clear examples of this are provided in Figure 4, with excess defects occurring at varied parts of the bead.

In contrast, the performance of the model trained on the balanced augmented set is not only superior but also much more harmonized across classes, with a difference of only ~4.5% in the validation set and ~1.5% in the holdout set, clearly showcasing the impact of the augmentation in terms of model performance for this task. While the performance is marginally better for discontinuity defects using a previous approach based on simulated data (0.01 AP), the performance gain enabled by the novel approach for the minority class is much more significant (from 0.3001 to 0.5864 AP).

Lastly, the model trained only on synthetic data resulted in worst performance across all tests, which reveals how important of a role the inclusion of real data still plays in this process, even considering a limited amount of samples.

4. Discussion

Recent advances in the field of GANs have made it possible to train these models with limited data while often achieving results close to previous state-of-the-art approaches. While generally a few thousand training images are still required for this purpose, we show that the usage of a much smaller dataset (under 200 images) still yielded results capable of greatly improving the performance of object detection models in the specific task of structural adhesive inspection, particularly in the case of imbalanced training data. Based on the original recommendation it can be hypothesized that results could be further improved by either increasing the size of the original dataset, or by employing for instance smaller partially overlapping crops.

These results suggest that the proposed approach has the potential to significantly mitigate the issues of data availability in such tasks, as it not only reduces the costs associated with the generation of images (energy, personnel and resources), but also the time required for this process. Naturally, it is generally unfeasible to dedicate a production line solely for the purpose of generating specific defects to create sufficiently large datasets.

Even in cases where a pilot line can be used for this purpose, such as the case presented herein, considerable effort is still required to manually trigger and supervise the process, whilst adjusting the parameters for defects with sufficient variety to be generated. Hence, these results reveal the potential of the proposed approach to improve the viability and likelihood of adoption of such a solution at an industrial level.

Particularly regarding Table 1, the results from the holdout set suggest that the key contribution of the augmentation is not only the increased volume of data, but also the variety of the defects generated by the GAN. This increased variety enables the object detector to perform better in previously unseen settings. This conclusion is further corroborated by Table 2, which highlights the performance differences between training with the small imbalanced set of real images, and the the larger balanced set of both real and synthetic images.

Nevertheless, some limitations remain to be tackled. On the one hand, the possibility to further increase the variety of the generated images with a finer degree of control on the type of defects generated would be tremendously useful. Besides increasing the size of the dataset, recent methods to explore the model's latent space such as GANSpace [27] could be experimented with to achieve this, on top of fine-tuning or extending the base architecture. On the other hand, the effort required to annotate the synthetic dataset is still considerable, as in this case this process was performed manually. A possible direction to explore here could be the usage of semi-supervised approaches to automatically label newly generated images and include them in the training loop, hence reducing the annotation effort.

In addition to this, future work will further explore the comparison of this approach with those based on simulated data (e.g., images generated in simulations such as those created in CoppeliaSim [28] or the Unity engine). Interesting points of focus can be the time to generate images and the degree of control. In this case, such approaches could prove to be a viable alternative if there are not sufficient real images to train the GAN at an initial stage, so long as the modeling effort required to generate synthetic images with sufficient quality is not too great. Thus, deriving a methodology or guidelines addressing which approach to adopt for different scenarios could prove useful.

5. Conclusions

In this article we propose a novel approach to address the challenge of data availability in the quality inspection of structural adhesive applications, leveraging recent advances regarding GAN training in limited data regimes.

We show that not only can realistic images of a variety of defects be generated quickly with this method, but also that the synthetic dataset can be used to augment scarce training sets for automated inspection solutions based on deep learning, greatly improving their performance on this task. We validate this in a real structural adhesive application line for automotive parts, with preliminary results showing considerable improvements in the mAP of state-of-the-art object detection models at different IoU thresholds when performing the automated defect detection.

Moreover, the proposed approach greatly reduces the costs of generating additional training data with sufficient samples of each defect. The process of generating datasets with sufficient quality and variety for training deep learning models is generally time consuming and costly when using traditional methods, particularly in terms of energy consumption, material and personnel costs. By enabling this at a fraction of the cost/time, this approach greatly contributes to improve the viability of deep learning for such quality inspection tasks.

Author Contributions: Conceptualization, R.S.P. and J.B.; methodology, R.S.P. and M.A.; software, R.S.P. and M.A.; validation, M.G., F.M.; formal analysis, R.S.P., M.A., M.G. and F.M.; data curation, R.S.P., S.O.A., M.G. and F.M.; writing—original draft preparation, R.S.P.; writing—review and editing, R.S.P., M.A., S.O.A., M.G., F.M. and J.B.; visualization, R.S.P., M.G. and F.M.; supervision, R.S.P. and J.B.; project administration, R.S.P. and J.B.; funding acquisition, R.S.P., M.G., F.M. and J.B.; All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded in part by the FCT/MCTES (UNINOVA-CTS funding UIDB/-00066/2020 and by the Compete 2020 program of the European Union within the scope of the SEE-Q project with the reference number POCI-01-0247-FEDER-034072.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are openly available in Github at https://github.com/RicardoSPeres/GAN_Synth_Adhesive (accessed on 30 March 2021).

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
ADA	Adaptive Discriminator Augmentation
AP	Average Precision
GAN	Generative Adversarial Network
IoU	Intersection Over Union
mAP	Mean Average Precision
MS COCO	Microsoft Common Objects in COntext
PLC	Programmable Logic Controller
YOLO	You Only Look Once

References

1. Tabernik, D.; Šela, S.; Skvarč, J.; Skočaj, D. Segmentation-based deep-learning approach for surface-defect detection. *J. Intell. Manuf.* **2020**, *31*, 759–776. [CrossRef]
2. Lee, K.J.; Kwon, J.W.; Min, S.; Yoon, J. Embedding Convolution Neural Network-Based Defect Finder for Deployed Vision Inspector in Manufacturing Company Frontec. In Proceedings of the AAAI Conference on Artificial Intelligence, New York, NY, USA, 7–12 February 2020; Volume 34, pp. 13164–13171.
3. Yang, Y.; Pan, L.; Ma, J.; Yang, R.; Zhu, Y.; Yang, Y.; Zhang, L. A high-performance deep learning algorithm for the automated optical inspection of laser welding. *Appl. Sci.* **2020**, *10*, 933. [CrossRef]
4. Cui, W.; Zhang, Y.; Zhang, X.; Li, L.; Liou, F. Metal Additive Manufacturing Parts Inspection Using Convolutional Neural Network. *Appl. Sci.* **2020**, *10*, 545. [CrossRef]
5. 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 Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops, Seoul, Korea, 27–28 October 2019.
6. Peres, R.S.; Jia, X.; Lee, J.; Sun, K.; Colombo, A.W.; Barata, J. Industrial Artificial Intelligence in Industry 4.0-Systematic Review, Challenges and Outlook. *IEEE Access* **2020**, *8*, 220121–220139. [CrossRef]
7. Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative adversarial nets. In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2014; pp. 2672–2680.
8. Kingma, D.P.; Welling, M. Auto-Encoding Variational Bayes. In Proceedings of the 2nd International Conference on Learning Representations, ICLR, Banff, AB, Canada, 14–16 April 2014.
9. Karras, T.; Aila, T.; Laine, S.; Lehtinen, J. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In Proceedings of the International Conference on Learning Representations, Vancouver, BC, Canada, 30 April–3 May 2018.
10. Karras, T.; Laine, S.; Aila, T. A style-based generator architecture for generative adversarial networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 16–20 June 2019; pp. 4401–4410.
11. Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; Aila, T. Analyzing and improving the image quality of stylegan. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 13–19 June 2020; pp. 8110–8119.
12. Cabrera, D.; Sancho, F.; Long, J.; Sánchez, R.V.; Zhang, S.; Cerrada, M.; Li, C. Generative adversarial networks selection approach for extremely imbalanced fault diagnosis of reciprocating machinery. *IEEE Access* **2019**, *7*, 70643–70653. [CrossRef]
13. Mao, W.; Liu, Y.; Ding, L.; Li, Y. Imbalanced fault diagnosis of rolling bearing based on generative adversarial network: A comparative study. *IEEE Access* **2019**, *7*, 9515–9530. [CrossRef]
14. Zhou, F.; Yang, S.; Fujita, H.; Chen, D.; Wen, C. Deep learning fault diagnosis method based on global optimization GAN for unbalanced data. *Knowl. Based Syst.* **2020**, *187*, 104837. [CrossRef]

15. Dai, Z.; Yang, Z.; Yang, F.; Cohen, W.W.; Salakhutdinov, R.R. Good semi-supervised learning that requires a bad gan. In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2017; pp. 6510–6520.
16. Kumar, A.; Sattigeri, P.; Fletcher, T. Semi-supervised learning with gans: Manifold invariance with improved inference. In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2017; pp. 5534–5544.
17. Di, H.; Ke, X.; Peng, Z.; Dongdong, Z. Surface defect classification of steels with a new semi-supervised learning method. *Opt. Lasers Eng.* **2019**, *117*, 40–48. [[CrossRef](#)]
18. Karras, T.; Aittala, M.; Hellsten, J.; Laine, S.; Lehtinen, J.; Aila, T. Training Generative Adversarial Networks with Limited Data. *arXiv* **2020**, arXiv:2006.06676.
19. Peres, R.; Guedes, M.; Miranda, F.; Barata, J. Simulation-based Data Augmentation for the Quality Inspection of Structural Adhesive with Deep Learning. *Techrxiv* **2021**. [[CrossRef](#)]
20. Cabrita, J.; Miranda, F.; Guedes, M. Smart control system for zero-defect adhesive application using industrial robots. In Proceedings of the Workshop on Disruptive Information and Communication Technologies for Innovation and Digital Transformation, Braganca, Portugal, 20 December 2019; p. 39.
21. Creswell, A.; White, T.; Dumoulin, V.; Arulkumaran, K.; Sengupta, B.; Bharath, A.A. Generative adversarial networks: An overview. *IEEE Signal Process. Mag.* **2018**, *35*, 53–65. [[CrossRef](#)]
22. Wang, C.Y.; Bochkovskiy, A.; Liao, H.Y.M. Scaled-YOLOv4: Scaling Cross Stage Partial Network. *arXiv* **2020**, arXiv:2011.08036.
23. 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 Proceedings of the European Conference on Computer Vision, Zurich, Switzerland, 6–12 September 2014; Springer: Berlin/Heidelberg, Germany, 2014, pp. 740–755.
24. Bochkovskiy, A.; Wang, C.Y.; Liao, H.Y.M. Yolov4: Optimal speed and accuracy of object detection. *arXiv* **2020**, arXiv:2004.10934.
25. Redmon, J. Darknet: Open Source Neural Networks in C. 2013–2016. Available online: <http://pjreddie.com/darknet/> (accessed on 9 March 2021).
26. Everingham, M.; Van Gool, L.; Williams, C.K.; Winn, J.; Zisserman, A. The pascal visual object classes (voc) challenge. *Int. J. Comput. Vis.* **2010**, *88*, 303–338. [[CrossRef](#)]
27. Härkönen, E.; Hertzmann, A.; Lehtinen, J.; Paris, S. GANSpace: Discovering Interpretable GAN Controls. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 9841–9850.
28. Rohmer, E.; Singh, S.P.N.; Freese, M. CoppeliaSim (formerly V-REP): A Versatile and Scalable Robot Simulation Framework. In Proceedings of the International Conference on Intelligent Robots and Systems (IROS), Tokyo, Japan, 3–7 November 2013. Available online: www.coppeliarobotics.com (accessed on 9 March 2021).