
Learning to Identify Drilling Defects in Turbine Blades with Single Stage Detectors

**Andrea Panizza^{1,*}, Szymon Tomasz Stefanek², Stefano Melacci^{2,3},
Giacomo Veneri¹, Marco Gori^{2,3,4}**

¹Department of Artificial Intelligence, Baker Hughes, Florence, Italy

²Siena Artificial Intelligence Lab, Siena, Italy

³DIISM, University of Siena, Siena, Italy

⁴Maasai, Université Côte d'Azur, Nice, France

andrea.panizza@bakerhughes.com, sts@pragma.software,

mela@diism.unisi.it, giacomo.veneri@bakerhughes.com,

marco.gori@unisi.it

Abstract

Nondestructive testing (NDT) is widely applied to defect identification of turbine components during manufacturing and operation. Operational efficiency is key for gas turbine OEM (Original Equipment Manufacturers). Automating the inspection process as much as possible, while minimizing the uncertainties involved, is thus crucial. We propose a model based on RetinaNet to identify drilling defects in X-ray images of turbine blades. The application is challenging due to the large image resolutions in which defects are very small and hardly captured by the commonly used anchor sizes, and also due to the small size of the available dataset. As a matter of fact, all these issues are pretty common in the application of Deep Learning-based object detection models to industrial defect data. We overcome such issues using open source models, splitting the input images into tiles and scaling them up, applying heavy data augmentation, and optimizing the anchor size and aspect ratios with a differential evolution solver. We validate the model with 3-fold cross-validation, showing a very high accuracy in identifying images with defects. We also define a set of best practices which can help other practitioners overcome similar challenges.

1 Introduction

Non-destructive testing (NDT) and structural health monitoring (SHM) techniques are of critical importance for gas turbine OEMS (Original Equipment Manufacturers) [1]. Due to the harsh operating environment, defects in turbine blades can impact the operating life, safety and performance of the gas turbine, so they have to be detected before assembly [2]. Internal defects, such as those which may be generated during the drilling of cooling holes, cannot be detected by visual means, and require the use of NDT techniques [3, 4]. One such technique is to acquire X-ray scans of the blades from different viewing angles. The resulting scans have to be analyzed by specialized workers, and the analysis is tedious and time-consuming, especially if large batches of blades are to be examined. Automating the inspection process is thus key to increased operational efficiency. In the last years, deep learning has dramatically improved the state-of-the-art in speech recognition, visual object recognition, object detection [5] and many other domains, such as drug discovery and genomics [6].

In this work, we apply a single stage detector, RetinaNet [7] to the problem of detecting drilling defects in turbine blades. The application of deep learning to defect detection in turbine blades has

*Corresponding author.

only started recently, and it has been confined to visible light or ultrasonic images of wind turbine blades [8, 9, 10]. To the best of our knowledge, this is the first time when deep learning is applied to the analysis of defects in X-ray images of gas turbine blades. This is a significant challenge because of the following issues:

- The starting images are very large (grayscale images in a proprietary format – resolution 8496×6960).
- The drilling defects are very small (about 14×18 pixels on average), and they are hard to detect for untrained humans. Detecting small objects is known to be a hard problem even for state-of-the-art object detection models [11, 12].
- The internal dataset we have is fairly small by deep learning standards (694 images, of which 134 with at least one defect). Given the necessity for OEMs to protect their Intellectual Property (IP), there are no public datasets of defects in gas turbine blades with which to integrate our internal one.

We overcome these issues by splitting the initial images into smaller tiles, scaling up the tiles so that the annotations become of comparable or larger size than the smallest anchors, and using random data augmentation, as well as anchor scale and aspect ratio optimization, in order to maximize the performance of the classifier. With these modifications, we show that our pipeline is capable of detection with almost human-level accuracy. Also, thanks to the usage of a single stage detector, which is faster than two-stage detectors, we keep inference time low, making the tool usable for industrial usage.

The structure of the paper is as follows: we first describe the dataset (Section 2) and the neural models (Section 3), followed by our experimental results (Section 4). Finally, conclusions and directions for future research are reported (Section 5).

2 Dataset

The blades in the dataset are first stage high pressure turbine blades. The X-ray images are taken at our manufacturing plant, where the drilling process is executed. We have 694 X-ray images, of which 134 with at least one defect. These images correspond to 498 distinct blades, of which 88 with at least one defect. This means that for a few blades multiple X-ray images from different viewing angles were acquired. A X-ray image of a generic turbine blade (not one of ours) is shown in Fig.1.

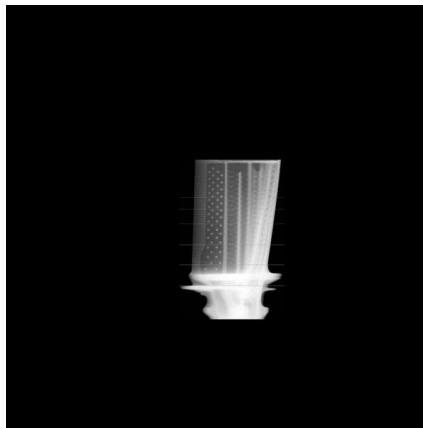


Figure 1: A sample X-ray view of a generic turbine blade. High resolution, large uninformative background region.

The high-quality X-ray scans are very large (8496×6960 pixels), and they include a large uninformative background region that has no relevance in the task of defect detection. The scans highlight holes that are drilled only in the part of the blade which is exposed to the gas flow (aerofoil), but not in the lower part which is secured in the rim of the wheel (foot). We used the `threshold` and `findContours` methods of the OpenCV library [13] to automatically crop out the aerofoil, which

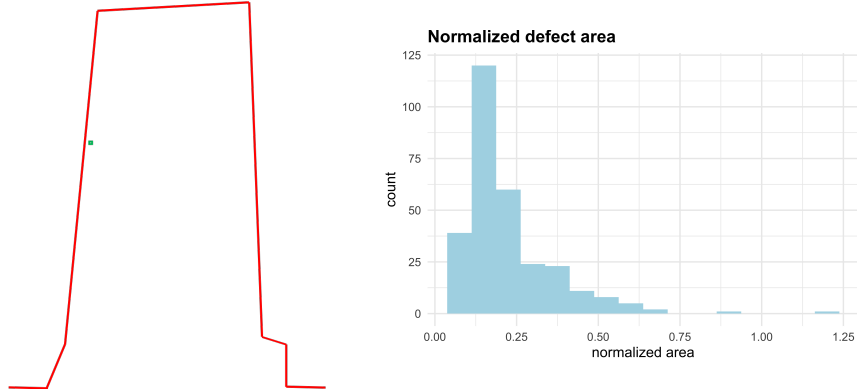


Figure 2: Left: relative size of a defect with respect to the aerofoil. The contour of the airfoil is reported in red, while the defect area is highlighted in green. Right: histogram of defect areas normalized by smallest anchor area.

Table 1: Dataset structure.

Set	Positive	Negative	Total
Train	363	400	763
Validation	153	175	328

results in a 1900×1500 image. The defects are extremely small with respect to the aerofoil, as it can be seen in the left part of Fig. 2. Each defect has been accurately annotated by highly experienced manufacturing engineers, for a total of 294 annotations. The right part of Fig. 2 shows the histogram of the defect areas normalized with respect to the area of the smallest RetinaNet anchor, which is 32×32 pixels. Clearly, most of the defects are smaller than the smallest anchor, which makes their detection extremely complicated.

We postpone to Section 3 the description of the RetinaNet model, and we only describe the dataset we used to validate this intuition. We selected all the 134 positive (with at least one defect) images, while for the negative (no defect) images, we selected only a random subsample of 148 images, for a total dataset size $N = 282$. The reason why we discarded most negative images is that the accuracy of one-stage detectors is affected by the problem of extreme foreground-background class imbalance. RetinaNet implements a focal loss to address this issue [7]. However, even with focal loss, negative (no defect at all) images remain problematic, thus we kept a ratio $\frac{\text{nr. of negative images}}{\text{nr. of positive images}} = 1.1$. We divided the dataset into a training set with 196 images and a validation set with 86 images. After 120 epochs with a minibatch size of 2 (limited by the large size of the images) on a NVIDIA Tesla V100 32Gb GPU, the mAP on the validation set, as computed by the RetinaNet implementation we used, was 0.06, i.e., the model did not learn anything useful.

To overcome this issue, we divided each image into 5×5 overlapping tiles of 500×600 pixels, as shown schematically in Fig. 3 (left), with an horizontal overlap of 250 pixels and a vertical one of 325. The overlap is important because object detection models are often less sensitive towards the boundaries of an image. Thanks to the overlap, if a defect is near the boundary of a tile, it is well inside the interior of at least one other tile, thus the model should be able to detect it. The only exception, of course, is the case in which the defect is near the boundary of the original (unsplit) image. Clearly, since the number of drilling defects in the positive images is limited, most of the tiles end up being empty. Thus, we balanced again the number of positive and negative tiles, ending up with the statistics reported in Table 1. Note that to avoid any data leakage, all the retained tiles from a given image either belong to the training or the validation set. At training time, we scaled the images and annotations up by $2 \times$ in height and width, so that the new histogram of defect areas normalized by the smallest anchor area become as shown in Fig. 3 (right). In Section 4 we will evaluate the effects of each of these operations.

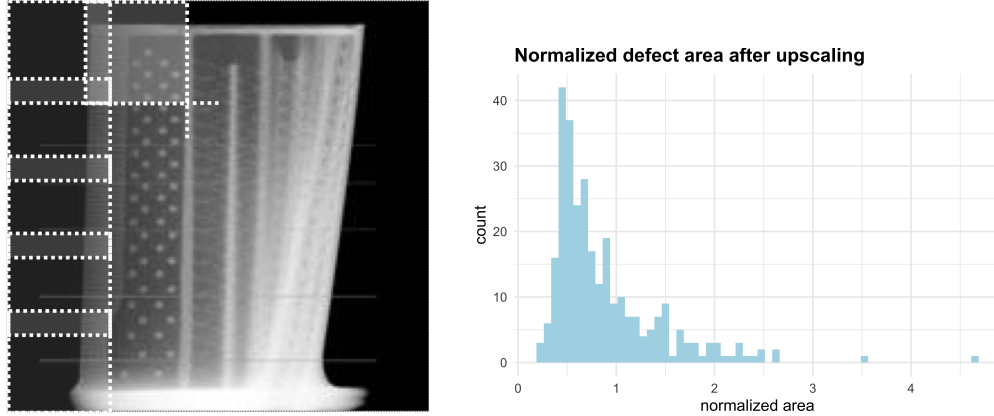


Figure 3: Left: subdivision of an image in overlapping tiles (we show the boundaries of some tiles using dotted lines). Right: histogram of defect areas normalized by the smallest anchor area, after scaling up height and width of each tile by $2\times$.

3 Model

RetinaNet [7] is a single stage model for object detection that uses a focal loss to address the common problem of class imbalance in detection tasks. RetinaNet exploits a feature pyramid network with a top-down architecture to detect objects at different scales [14]. The implementation we used is a fork of the Fizyr² implementation. The RetinaNet model architecture consists of four major components:

- a *bottom-up pathway*, the backbone network which calculates the feature maps at different scales. We used a ResNet-50 model [15].
- *top-down pathway and lateral connections*, the top down pathway upsamples the spatially coarser feature maps from higher pyramid levels, and the lateral connections merge the top-down layers and the bottom-up layers with the same spatial size.
- a *classification subnetwork* predicts the probability of an object being present at each spatial location for each anchor box and object class.
- a *regression subnetwork* regresses the offset for the bounding boxes from the anchor boxes to a nearby ground-truth object, if one exists.

For more details, we point to the original RetinaNet paper. Regarding our specific implementation, the model is pretrained on the MS COCO dataset[16]. We used random data augmentation to improve the model performance. We trained the model for 40 epochs.

3.1 Anchor Optimization

The anchor configuration is a crucial feature in the detection process, and we find the default anchor sizes (32, 64, 128, 256 and 512), aspect ratios (0.5, 1, 2) and scales (1, 1.2, 1.6) to be ineffective for detecting our small defects. We employed the differential evolution search algorithm used by [17] to optimize ratios and scales of anchors on the validation set. The algorithm iteratively improves a population of candidate solutions with regard to an objective function. The implementation is freely available on GitHub³.

4 Results

We present the results of the compared settings in Table 2. The model trained on the original images, without tiling and upscaling, delivers useless results. The model retrained on the splitted images shows a minimal improvement, with a mAP on the validation set of 0.10, since the vast majority of

²<https://github.com/fizyr/keras-retinanet>

³<https://github.com/martinzlocha/anchor-optimization/>

Table 2: Performance metrics.

Setting	mAP	Accuracy
Original images (1500 × 1900)	0.06	0.56
Splitted tiles (500 × 600)	0.10	0.53
Splitted tiles + 2x upscaling (1000 × 1200)	0.73	0.90
Splitted tiles + 2x upscaling (1000 × 1200) + anchor optimization	0.90	0.94

the annotations is smaller than the smallest default anchor. Adding 2× upscaling drives the mAP up to 0.73. This is much better, but still not enough to deploy the model in production. Finally, including anchor optimization brings the mAP to 0.90, which meets our requirements. The results are summarized in Table 2 (mAP).

We also show the results of a different metric, which is more closely related to our business goals (Table 2, Accuracy). We care more about catching blades with one or more defects, rather than localizing the defects exactly. For this reason, we use the following definition of accuracy, for which an image is correctly classified if the following conditions are both true:

- The model predicts a number of defects that is at least half of the number of annotations for the considered image.
- Let S be the union of the predicted bounding boxes, and G the union of the annotations. The Intersection over Union (IoU) of G and S must be more than 0.2.

To estimate the generalization error, we computed this metric in a 3-fold cross-validation setting. The resulting accuracy averaged over the 3 folds is **0.94** (with negligible standard deviation) – very close to human accuracy, and large enough for our goals.

5 Conclusion

We developed a model to identify drilling defects in X-ray images of gas turbine blades. Simply training a very powerful object detection model such as RetinaNet on our dataset did not deliver any usable results. We identified the key issues, which were the large size of the images, preventing the usage of large minibatches, the small average size of the annotations, and the use of default anchor scales and aspect ratios, which were not optimized for our data distribution. The final model delivers very promising results. Before deploying the model in production, we will collect more data to better estimate the generalization error of the model. This work also opens the way to a more general application of deep learning methods to defect detection in gas turbine components.

References

- [1] F. Mevissen and M. Meo, “Detection of cracks in turbine blades with nonlinear ultrasonic frequency modulation,” in *Structural Health Monitoring 2019*, F.-K. Chang, A. Guemes, and F. Kopsaftopoulos, Eds. USA United States: DEStech Publications Inc., Sep. 2019, pp. 824–834.
- [2] J. Aust and D. Pons, “Taxonomy of gas turbine blade defects,” *Aerospace*, vol. 6, no. 5, p. 58, May 2019. [Online]. Available: <http://dx.doi.org/10.3390/aerospace6050058>
- [3] B. Li, L. Chen, Y. Wang, and M. Gao, “3d detection of internal defects for gas turbine blades,” *Insight - Non-Destructive Testing and Condition Monitoring*, vol. 59, no. 7, pp. 364–370, 2017. [Online]. Available: <https://www.ingentaconnect.com/content/bindt/insight/2017/00000059/00000007/art00005>
- [4] L. Chen, B. Li, H. Zhou, Z. Li, and Z. Shang, “Detection of three-dimensional parameter of defects for gas turbine blades based on two-dimensional digital radiographic projective imaging,” *Journal of Nondestructive Evaluation*, vol. 38, no. 4, p. 101, Nov 2019. [Online]. Available: <https://doi.org/10.1007/s10921-019-0640-3>

- [5] L. Liu, W. Ouyang, X. Wang, P. Fieguth, J. Chen, X. Liu, and M. Pietikäinen, “Deep Learning for Generic Object Detection: A Survey,” *International Journal of Computer Vision*, vol. 128, no. 2, pp. 261–318, Feb. 2020, arXiv: 1809.02165 Publisher: Springer.
- [6] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015. [Online]. Available: <http://www.nature.com/articles/nature14539>
- [7] T. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal loss for dense object detection,” in *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 2999–3007.
- [8] A. S. M. Shihavuddin, X. Chen, V. Fedorov, A. Nymark Christensen, N. Andre Brogaard Riis, K. Branner, A. Bjorholm Dahl, and R. Reinhold Paulsen, “Wind Turbine Surface Damage Detection by Deep Learning Aided Drone Inspection Analysis,” *Energies*, vol. 12, no. 4, p. 676, Jan. 2019, number: 4 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: <https://www.mdpi.com/1996-1073/12/4/676>
- [9] Y. Yu, H. Cao, X. Yan, T. Wang, and S. S. Ge, “Defect identification of wind turbine blades based on defect semantic features with transfer feature extractor,” *Neurocomputing*, vol. 376, pp. 1–9, Feb. 2020. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0925231219313396>
- [10] T. Li, Y. Yang, Q. Wan, D. Wu, and K. Song, “Investigation of Wind Turbine Blade Defect Classification Based on Deep Convolutional Neural Network,” in *Proceedings of the 9th International Conference on Computer Engineering and Networks*, ser. Advances in Intelligent Systems and Computing, Q. Liu, X. Liu, L. Li, H. Zhou, and H.-H. Zhao, Eds. Singapore: Springer, 2021, pp. 207–213.
- [11] H. Ji, Z. Gao, T. Mei, and B. Ramesh, “Vehicle detection in remote sensing images leveraging on simultaneous super-resolution,” *IEEE Geoscience and Remote Sensing Letters*, vol. 17, no. 4, pp. 676–680, 2020.
- [12] J. Rabbi, N. Ray, M. Schubert, S. Chowdhury, and D. Chao, “Small-Object Detection in Remote Sensing Images with End-to-End Edge-Enhanced GAN and Object Detector Network,” *arXiv:2003.09085 [cs]*, Apr. 2020, arXiv: 2003.09085. [Online]. Available: <http://arxiv.org/abs/2003.09085>
- [13] G. Bradski, “The OpenCV Library,” *Dr. Dobb’s Journal of Software Tools*, 2000.
- [14] T. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, “Feature pyramid networks for object detection,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 936–944.
- [15] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [16] T. Lin, M. Maire, S. J. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, “Microsoft COCO: common objects in context,” *CoRR*, vol. abs/1405.0312, 2014. [Online]. Available: <http://arxiv.org/abs/1405.0312>
- [17] M. Zlocha, Q. Dou, and B. Glocker, “Improving RetinaNet for CT Lesion Detection with Dense Masks from Weak RECIST Labels,” *arXiv:1906.02283 [cs, eess]*, Jun. 2019, arXiv: 1906.02283. [Online]. Available: <http://arxiv.org/abs/1906.02283>