Proceedings of the 1st international Symposium on AI and Fluid Mechanics
Paper No [leave blank]

Computer Vision-Based Performance Forecasting for Turbomachinery

Anirudh Narayan Rao[1,3,†], Mario Carta[2,3], Tiziano Ghisu[2], Shahrokh Shahpar[1,3], and Francesco Montomoli[1]

[1]Department of Aeronautics, Imperial College London, South Kensington, London, SW7 2AZ, U.K.
[2]Department of Mechanical, Chemical and Materials Engineering, University of Cagliari, Cagliari, 09123, Italy

[3]Fluid Mechanics, Design Engineering, ET&S, Rolls-Royce plc., Derby, DE24 8BJ, U.K.
[†] Corresponding author's email : anirudh.rao@imperial.ac.uk

Abstract. This study investigates the application of a prototype convolutional neural network (CNN) architecture for the image segmentation of aero-engine blades, extracted from borescope videos, to identify deterioration patterns that can be mapped onto existing computational fluid dynamics (CFD) grids. A state-of-the-art CNN architecture, YOLO (You Only Look Once), is employed to segment turbine blades exhibiting in-service degradation, enabling automated damage detection, classification, and quantification of geometric defects. For geometric reconstruction, photogrammetry is utilised to generate a three-dimensional CAD model from a series of 2D images, providing an accurate representation of blade damage. The reconstructed model is then compared with structured light scanning (SLS) data to assess reconstruction accuracy and quantify uncertainties. Additionally, three-dimensional CFD analysis is conducted on the reconstructed geometry using BOXER meshing and the in-house HYDRA solver to evaluate aerodynamic performance. The results indicate that the reconstructed geometries are consistent with the expected performance degradation of end-of-life blades. Lastly, a modified U-Net CNN architecture, designed for real-time computer vision applications, is employed to predict the flow field of the damaged blades. This approach presents a novel and efficient method for predicting blade performance degradation, enabling real-time monitoring and data-driven maintenance strategies in aero-engine applications.

Keywords: image segmentation, CFD, photogrammetry, convolutional neural network (CNN), U-Net

1. Introduction

In recent years, computer vision techniques have advanced significantly, particularly in aerospace engineering, where they are increasingly utilised for the monitoring and maintenance of critical components. A promising approach is the use of borescope videos for condition-based maintenance (CBM) of turbomachinery components [1]. Over time, factors such as erosion, spallation, and de-lamination of thermal barrier coatings (TBC) contribute to the degradation of component performance. Image segmentation is key for identifying damage on turbomachinery blades by segmenting images into regions of interest. This method has been applied in several engineering fields, including crack detection in infrastructure [2, 3]. Convolutional neural networks (CNNs), particularly the You Only Look Once (YOLO) architecture [4, 5], are widely used for real-time object detection and segmentation. YOLO's ability to detect and classify damage in turbine blades from borescope images allows for efficient monitoring, facilitating early detection of damage that could degrade performance [6, 7]. The adoption of YOLO in industrial applications has proven effective in improving detection accuracy and real-time processing, making it ideal for turbine blade inspection [8, 9].

In 3D reconstruction, photogrammetry plays a crucial role in creating accurate models of end-of-life turbine blades from multiple images. These models are essential for assessing degradation, such as erosion, cracks, and material loss, which directly impact aero-thermodynamic performance. Reitenbach et al. [10] explored an end-to-end digital twin approach, integrating 3D scanning, geometric reconstruction, and simulations (CFD and CSM) to evaluate degradation. Their CFD analysis of scanned blades revealed a 1.36% increase in total pressure loss and a 0.3% reduction in power output compared to the design intent, highlighting the significant effect of degradation on turbine performance. High-fidelity CFD studies emphasise the importance of accurate simulations to understand how surface degradation, including erosion and thermal cycling, affects aerodynamic performance [11, 12, 13, 14, 15, 16]. For instance, Li et al. (2020) used CFD to examine the impact of erosion and thermal cycling on turbine blades, shedding light on performance losses [17]. Similarly, Meyer et al. (2022) investigated the effects of TBC degradation on turbine performance [18].

This study integrates image segmentation with high-precision 3D reconstruction to detect, quantify, and model

© Rolls-Royce plc. 2025

turbomachinery degradation. This approach enhances assessment accuracy, correlates damage features with performance changes, and supports extended time-on-wing and component life cycles in critical aerospace applications.

2. Methodology and Insights

This section outlines the methodology employed, beginning with the YOLOv11 analysis of the blades, followed by photogrammetry, and concluding with the real-time CFD prediction. A dataset of high-resolution images extracted from a borescope video of a turbine blade row exhibiting signs of in-service engine degradation was used for the image segmentation analysis. The borescope images typically captured the suction side of the blade, which faces the camera, with the hub positioned on the left-hand side of each image. The suction side, being more exposed to hot gases, is more prone to defects. In each image, one blade was fully visible, while a second blade was only partially visible. The images captured using the borescope would also require correction for barrel distortion inherent in the optical system. The initial task involved manually categorising each image into three distinct categories: blade outline, thermal barrier coating (TBC) loss (also known as chip damage) and shroud damage. This process involved manually identifying the defect boundaries for segmentation and drawing bounding boxes around the damaged areas using LabelImg and Labelme tools for detection [19, 20]. The annotations of each image were stored as a Javascript Object Notation (`.json`) files and converted into the YOLO format, where the file structure sequentially lists the classes and their corresponding coordinates. For the dataset, approximately 83.33% of the available images were used for training, while the remaining 16.66% were set aside for validation and testing. Model training employed a batch size of 5 over 750 epochs, with each epoch lasting approximately 30–35 seconds on an Nvidia K80 GPU. The model weights were saved upon training, and the metrics analysed. To enhance the detection accuracy of even the smallest damages, Slicing-Aided Hyper Inference (SAHI) [21] was applied. The prediction results for an unseen blade (figure 1(a)) for the detection task using the standard YOLOv11 model, the YOLOv11 + SAHI technique (36 slices with an overlap ratio of 0.2 in each dimension), and the segmentation tasks are presented in figures 1(b), 1(c), and 1(d), respectively. SAHI effectively detects even the smallest pitting damage but occasionally misclassified the regions of interest. From the segmented image, the damaged area offers a quantitative assessment, which is essential for accurate and informed decision-making in engine sentencing.

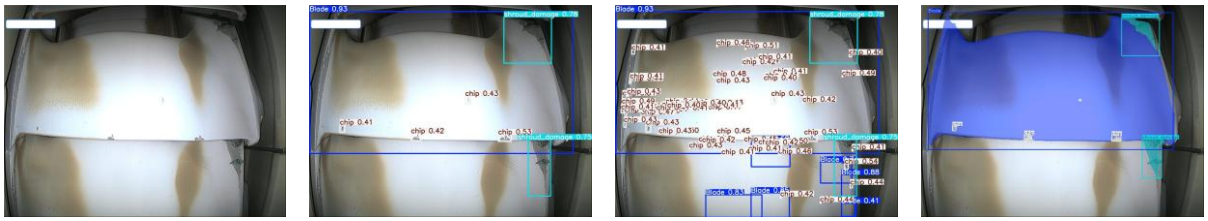

(a) (b) (c) (d)

Figure 1: (a) Original image from the borescope video, (b) Detection of the TBC chips, shroud damage and the blade using the standard YOLOv11 weights and (c) the prediction with SAHI, and (d) the segmented regions of (a). The numbers indicate the confidence of prediction. Images have been resized to maintain confidentiality.

Given that the pressure side of the blade is occluded in the borescope images, three-dimensional damage reconstruction is carried out using photogrammetry, a method that employs photographs captured from multiple angles to generate highly detailed models. Compared to conventional laser scanning, which can take several hours per blade, photogrammetry offers a much faster alternative for blade reconstruction. The blades were mounted on a turntable and photographed at 15° intervals, capturing 24 images per viewing position. To ensure comprehensive coverage, images were taken from three positions: the shroud, midspan, and bottom mounting platform, including the fir tree section, resulting in 72 images per blade. The quality of the reconstruction was influenced by factors such as lighting conditions, camera focus, and the number of images. To validate the accuracy of the reconstructed geometries, they were compared with structured light scanning (SLS), the industry-standard technique for aerospace component assessment. Three photogrammetry tools were used to "stitch" the photographs into a 3D point cloud. Using the open-source software CloudCompare [22], the reconstructed geometries were scaled and aligned with

© Rolls-Royce plc. 2025

the reference SLS point cloud. More than forty corner points along the top and bottom platforms were selected for initial rough alignment, which was then refined using the iterative closest point (ICP) algorithm [23, 24] with a root mean squared error (RMSE) threshold of 10^{-6}. Even small misalignments at this stage could compromise the accuracy of the reconstructed geometry and subsequent downstream analyses. The final alignment and scaling of the reconstructed blade were determined using a transformation matrix, which could be applied to other blades processed through this method. Figure 2 shows the distance field contours of the three reconstructed geometries, measured in Meshlab [25], along with the computed Hausdorff distance between the reference and reconstructed meshes in both front and rear views.

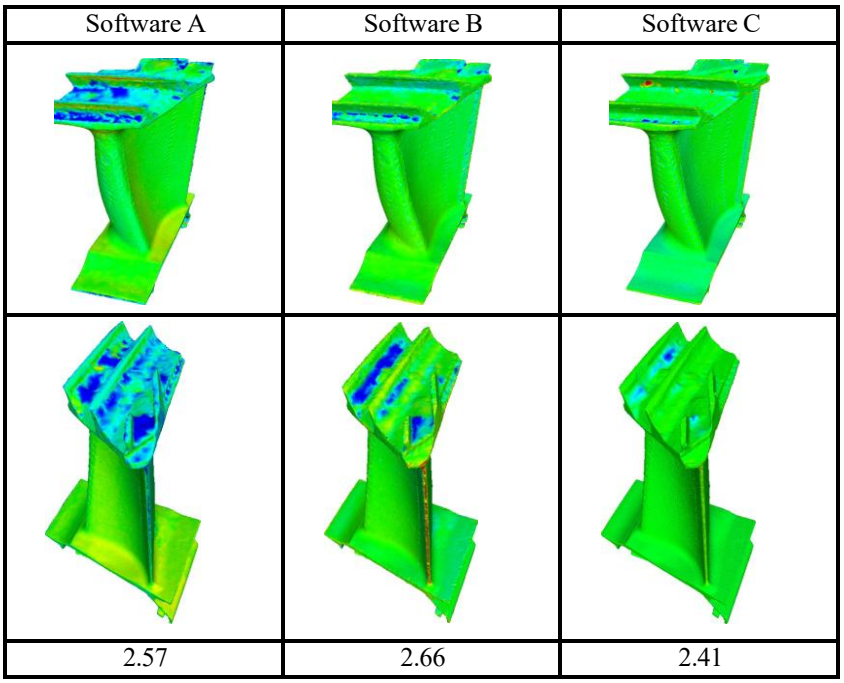

Figure 2: Contours of the distance field between the three reconstructions and the SLS geometry (± 0.25) in perspective views of the front (top row) and rear (bottom row) of the blade. Red regions correspond to negative values (depressions) and blue regions correspond to positive values (protrusions). The Hausdorff distance between the reference and reconstructed geometries is shown at the bottom. Images resized to maintain confidentiality.

The reconstructed blade was exported as an `.stl` file, serving as the geometric input to the Boxer CFD grid generation software [26]. CFD simulations were performed using the Favre-averaged Navier-Stokes solver, HYDRA [27], following the methodology described in [16]. The aerodynamic performance of the reconstructed blade was then compared with the performance trendlines of over 120 SLS-derived geometries (Figure 3(a)). The regression plot of efficiency variation with tip and shroud gap parameters is described in equation 1, where the shroud damage variable, SDV_η, is a linear combination of five critical points on the turbine shroud and the interstitial gap between adjoining turbine blades (Figure 3(b)). The reconstructed blades exhibit close agreement with the regression trendline. It is acknowledged that small variations in the reconstruction process could have influenced the scatter observed in aerodynamic efficiency. Notably, the edges of the shroud between adjacent blades must be aligned precisely to minimize flow leakage, as even small gaps can significantly affect the aerodynamic performance [16, 28].

$$\mathrm{SDV}_\eta = 0.057S_1 + 0.365S_2 + 0.008S_3 + 0.342T_1 + 0.227T_2 \quad (1)$$

Following the photogrammetry process, the objective is to predict the flow field around the reconstructed blade using machine learning techniques. Computer vision methods are leveraged to facilitate rapid prediction of flow features. For this task, the classical U-Net architecture [29], which includes a contracting path, expanding path, and bottleneck layer, is employed. In the contracting path, convolution operations and downsampling of the image occur, reducing the spatial resolution of the feature map while increasing its feature depth. In the expanding

© Rolls-Royce plc. 2025

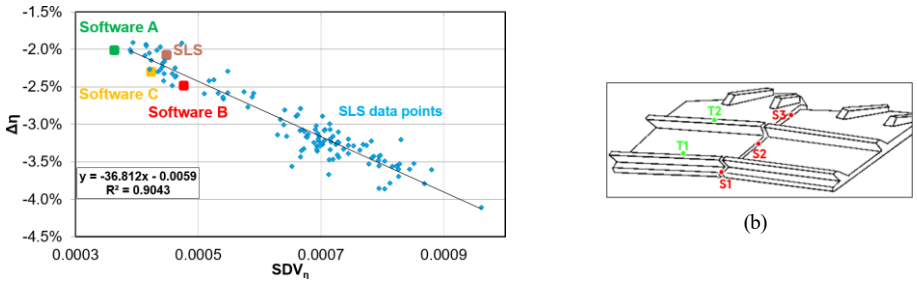

Figure 3: (a) Regression plot illustrating the linear relationship between efficiency and gap widths for the 120 SLS scans (in blue) with an $R^2 = 0.9043$. Annotated in various colours are the three reconstructions and the SLS scanned geometry. (b) Schematic of turbine shroud showing the location the five points of measurement used in Eqn. 1. Images adopted from Carta et al. [16].

path, the depth of the feature maps is reduced while the spatial resolution increases. Skip connections between corresponding encoder and decoder layers transfer fine-grained spatial information, preserving high-resolution features essential for accurate segmentation.The U-Net architecture is predominantly used for image segmentation tasks, processing multi-channel inputs — typically a three-channel RGB image — and generating segmented outputs. These masks, typically binary images, correspond to specific classes within the input image, allowing the network to delineate and identify regions of interest with high accuracy. The design of the network, characterised by its symmetric encoder-decoder structure, facilitates effective localisation and segmentation, even with limited labelled data. Upon completion of training on a dataset of images, the resulting output is typically a segmented image. A novel modification to this architecture is introduced by reversing the typical process: a binary input and a three-channel RGB mask (flow contours from CFD simulations) are used, with the resulting prediction being a three-channel RGB image [30]. To implement this, a dataset of approximately two hundred images from a design of experiments on self-similar damaged blades at a fixed operating condition is employed. The RGB images serve as masks, while the corresponding binary images, forming the blade stencil, are used as inputs to the modified network. Data augmentation is applied due to the small sample size, increasing the dataset by rotating, rolling, and flipping both the RGB images and their corresponding binary images. Training is carried out over several hundred epochs until the metrics of interest, such as validation and training loss stabilise. Images of 256^2 pixels, showing the contours of relative velocity (or any state variable of interest) at the midspan of the turbine cascade, are used, with a 90:10 training-validation split. An unseen stencil of the stacked blades is used as input, and the predicted flow field is compared with the CFD simulation result (Figure 4). Since the values are normalised, the actual flow variable quantities can be recovered and compared with the CFD solution. The structural similarity index [31] of the predicted images was $\gtrsim 0.975$, with a low mean squared error. While obtaining a CFD solution typically requires several hours and significant computational resources, the trained neural network offers instantaneous flow field predictions, enabling engineers to rapidly assess any given cross-section of the blade. Additionally, the model can be trained on images from multiple cross-sections, allowing real-time flow field predictions of blades reconstructed through the photogrammetry process.

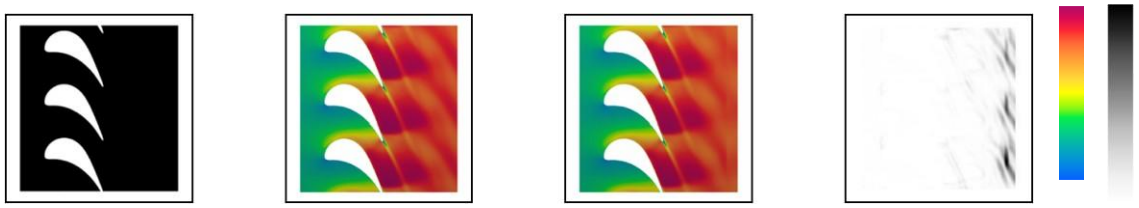

Figure 4: Prediction of the relative velocity contours of the turbine blade at midspan. The prediction is close to the ground truth and the errors are in the far wake. Left to right: The binary stencil, the ground truth, the predicted image, and the difference between the ground truth and the prediction visualised in greyscale. Colourbars run from low to high value. Flow is from left to right in these image, and the images are resized to maintain confidentiality.

© Rolls-Royce plc. 2025

3. Conclusions

This pilot study demonstrates the efficacy of a prototype convolutional neural network (CNN) architecture for segmenting borescope images of turbomachinery blades, enabling the identification of deterioration patterns. The integration of photogrammetry for 3D blade geometry reconstruction, alongside high-fidelity computational fluid dynamics (CFD) simulations, facilitates a thorough assessment of aerodynamic performance. Additionally, a novel modification to the U-Net architecture, involving the reversal of RGB and binary channels, enables real-time predictions of the flow field based on the blade stencil input. This integrated approach offers a robust framework for predicting performance degradation in in-service blades, with the potential to improve residual life assessments and inform data-driven engine maintenance decisions.

Acknowledgements. The authors acknowledge the computing resources provided by Centre for Modelling and Simulation (CFMS), UK and Rolls Royce plc, UK. This study is supported by the EU Horizons Nextair (Nextair) project which has received funding from the European Union's Horizon Europe research and innovation program under grant agreement No. 101056732. **Data confidentiality statement.** Data supporting this study cannot be made available due to commercial confidentiality agreements. **Conflict of interest.** The authors declare no conflicts of interest.

References

[1] Wong, C. Y., Seshadri, P., and Parks, G. T., "Automatic Borescope Damage Assessments for Gas Turbine Blades via Deep Learning," AIAA Scitech 2021 Forum, No. AIAA 2021-1488, 2021, pp. 1–17.

[2] Kheradmandi, N. and Mehranfar, V., "A critical review and comparative study on image segmentation-based techniques for pavement crack detection," *Construction and Building Materials*, Vol. 321, 2022, pp. 126162.

[3] Deng, J., Singh, A., Zhou, Y., Lu, Y., and Lee, V. C.-S., "Review on computer vision-based crack detection and quantification methodologies for civil structures," *Construction and Building Materials*, Vol. 356, 2022, pp. 129238.

[4] Jocher, G. and Qiu, J., "Ultralytics YOLO11," 2024.

[5] Jiang, P., Ergu, D., Liu, F., Cai, Y., and Ma, B., "A Review of Yolo Algorithm Developments," *Procedia Computer Science*, Vol. 199, 2022, pp. 1066–1073.

[6] Li, X., Wang, C., Ju, H., and Li, Z., "Surface Defect Detection Model for Aero-Engine Components Based on Improved YOLOv5," *Applied Sciences*, Vol. 12, No. 14, 2022.

[7] Liu, J., Wang, H., and Jiang, A., "Automated Crack Detection With Image Analysis for the Blades of Steam Turbine," *Journal of Engineering for Gas Turbines and Power*, Vol. 144, No. 8, 06 2022, pp. 081001.

[8] Zubayer, M. H., Zhang, C., Liu, W., Wang, Y., and Haque, M. I., "Automatic Defect Detection of Jet Engine Turbine and Compressor Blade Surface Coatings Using a Deep Learning-Based Algorithm," *Coatings*, Vol. 14, No. 4, 2024, pp. 1–15.

[9] Wu, Z., Zhang, Y., Wang, X., Li, H., Sun, Y., and Wang, G., "Algorithm for detecting surface defects in wind turbines based on a lightweight YOLO model," *Scientific Reports*, Vol. 14, No. 1, 2024, pp. 24558.

[10] Reitenbach, S., Ebel, P.-B., Grunwitz, C., Siggel, M., and Pahs, A., "Digital Transformation in Aviation: An End-to-End Digital Twin Approach for Aircraft Engine Maintenance – Exploration of Challenges and Solutions," Proceedings of Global Power and Propulsion Society (GPPS Chania24), Global Power and Propulsion Society, Chania, Greece, September 4-6, 2024, 2024, pp. 1–13, GPPS-TC-2024-206.

[11] Suttner, A., Heinrich, M., Wuttke, T., and Kroll, M., "Impact of Blade Erosion on Turbine Performance: Simulation and Experimental Validation," *Journal of Turbomachinery*, Vol. 145, No. 7, 2023, pp. 071005.

[12] Khalil, H., Badr, I. K., and Al-Durrah, K., "Aerodynamic Degradation Due to Surface Erosion: A CFD Study of a Gas Turbine Blade," *Energy*, Vol. 255, 2022, pp. 124482.

© Rolls-Royce plc. 2025

[13] Harris, D., Kiani, M. F., and Wang, L., "CFD Simulation of the Effects of Surface Degradation on the Aerodynamic Performance of High-Pressure Turbine Blades," *Computers & Fluids*, Vol. 233, 2021, pp. 104920.

[14] Zhao, X., Wei, Z., and Song, X., "A High-Fidelity CFD Study on the Effect of Surface Erosion on the Aerodynamics of Shrouded High-Pressure Turbine Rotor Blades," *Applied Thermal Engineering*, Vol. 188, 2021, pp. 116596.

[15] Tian, X., Zhang, L., and Zhang, L., "Numerical Simulation of Erosion Effects on Shrouded High-Pressure Turbine Blades Using High-Fidelity CFD," *Journal of Thermal Science and Engineering Applications*, Vol. 15, No. 3, 2023, pp. 031007.

[16] Carta, M., Ghisu, T., and Shahpar, S., "High-Fidelity Computational Fluid Dynamics Analysis of In-Service Shrouded High-Pressure Turbine Rotor Blades," *Journal of Turbomachinery*, Vol. 144, No. 12, 09 2022, pp. 121001.

[17] Li, W., He, J., and Song, P., "High-Fidelity CFD Analysis of Degraded High-Pressure Turbine Blades: Effects of Surface Erosion and Thermal Cycling," *Journal of Turbomachinery*, Vol. 142, No. 7, 2020, pp. 071004.

[18] Meyer, J., Biedenkapp, M., and Mueller, T. R., "Impact of Thermal Barrier Coating Degradation on High-Pressure Turbine Blade Performance: A CFD Study," *Journal of Engineering for Gas Turbines and Power*, Vol. 144, No. 10, 2022, pp. 101004.

[19] Labelme, "Labelme: Image Polygonal Annotation with Python (polygon, rectangle, circle, line, point and image-level flag annotation)," https://github.com/labelmeai/labelme, 2024, Accessed: Jan 2024.

[20] LabelImg, "LabelImg: Graphical Image Annotation Tool (supports YOLO and CreateML formats)," https://pypi.org/project/labelImg/, 2024, Accessed: Jan 2024.

[21] Akyon, F. C., Onur Altinuc, S., and Temizel, A., "Slicing Aided Hyper Inference and Fine-Tuning for Small Object Detection," 2022 IEEE International Conference on Image Processing (ICIP), 2022, pp. 966–970.

[22] Girardeau-Montaut, D., "Cloud Compare v2.1.3 - 3D Point Cloud and Mesh Processing Software, GPL," 2023, Accessed: 2025-01-25.

[23] Besl, P. J. and McKay, N. D., "A method for registration of 3-D shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 14, No. 2, 1992, pp. 239–256.

[24] CloudCompare, "Fine Registration (ICP)," 2024, Accessed: 2025-01-25.

[25] Cignoni, P., Callieri, M., Corsini, M., Dellepiane, M., Ganovelli, F., and Ranzuglia, G., "MeshLab: an Open-Source Mesh Processing Tool," Eurographics Italian Chapter Conference, edited by V. Scarano, R. D. Chiara, and U. Erra, The Eurographics Association, 2008.

[26] Demargne, A. A., Evans, R., Tiller, P., and Dawes, W. N., "Practical and Reliable Mesh Generation for Complex, Real-World Geometries," 52nd Aerospace Sciences Meeting, 2014, pp. 1–9.

[27] Lapworth, L., "Hydra-CFD: a framework for collaborative CFD development," International conference on scientific and engineering computation (IC-SEC), National Harbor, Maryland, Vol. 30, 2004, pp. 1–6.

[28] Carta, M., Ghisu, T., and Shahpar, S., "Heat Transfer Analysis of Damaged Shrouded High-Pressure Turbine Rotor Blades," *International Journal of Turbomachinery, Propulsion and Power*, Vol. 8, No. 3, 2023, pp. 1–12.

[29] Ronneberger, O., Fischer, P., and Brox, T., "U-Net: Convolutional Networks for Biomedical Image Segmentation," *CoRR*, Vol. abs/1505.04597, 2015.

[30] Yu, Y., Chen, S., and Wei, H., "Modified UNet with attention gate and dense skip connection for flow field information prediction with porous media," *Flow Measurement and Instrumentation*, Vol. 89, 2023, pp. 102300.

[31] Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P., "Image Quality Assessment: From Error Visibility to Structural Similarity," *IEEE Transactions on Image Processing*, Vol. 13, No. 4, 2004, pp. 600–612.

© Rolls-Royce plc. 2025