Enhanced FOD Detection on Runways Using Deep Learning Based Image Restoration

Amarjeet Kumar
Department of Aerospace Engineering & Autonomous Systems
Defence Institute of Advanced Technology
Pune, India
amarjeetkumarae@gmail.com

Abhishek Dixit
Department of Engineering
National Defence Academy
Pune, India
abhishekdixit84@gmail.com

Pooja Agrawal
Department of Aerospace Engineering & Autonomous Systems
Defence Institute of Advanced Technology
Pune, India
poojaa.diat@gmail.com

doi: https://doi.org/10.37285/bsp.SACAD2025.59

***Abstract*—This work presents a deep learning-based framework for robust Foreign Object Debris (FOD) detection on airport runways under degraded imaging conditions. The framework integrates a Convolutional Neural Network (CNN) based image restoration model with the YOLOv11n detector to enhance visual quality and detection accuracy in runway imagery. Both Gaussian and motion blur, typically introduced due to optical defocus or drone movement, were synthetically simulated and effectively mitigated through supervised image restoration and detection optimization. The framework was trained on an FOD dataset that was synthetically blurred using controlled Gaussian and motion kernels to approximate real-time optical defocus and drone induced vibration. The CNN based restoration model was quantitatively compared with the DeblurGAN baseline to assess restoration performance. The CNN achieved average gains of 7.2% and 11.2% in Peak Signal to Noise Ratio (PSNR) and Structural Similarity Index (SSIM) for Gaussian blur, and 12.1% and 18.2% for motion blur, outperforming DeblurGAN by 5–7%. On real blurred imagery, PSNR improved to 35.12 dB and SSIM to 0.91, demonstrating consistent restoration quality under practical conditions. Runtime evaluation confirmed real-time efficiency, The CNN+YOLO framework achieved real-time performance with an inference time of 0.0099 s per image (≈101 FPS), outperforming DeblurGAN+YOLO at 67.08 FPS. This integration of the image restoration model with YOLOv11n enhances detection reliability and visual consistency under blur distortions, making the framework suitable for real-time runway surveillance and future drone based FOD monitoring applications.**

Keywords*—*FOD, deblurring, YOLOv11, motion blur, gaussian blur, CNN, flight safety

I. Introduction

Ensuring runway safety is one of the most critical aspects of modern aviation management. The presence of Foreign Object Debris (FOD), which includes metallic fragments, tools, stones, and plastic components, poses a serious threat to aircraft operations. Even small debris can cause significant engine or landing gear damage, resulting in costly repairs, operational delays, and major safety risks [1]. Therefore, reliable and timely detection of FOD is essential for maintaining the safety and efficiency of airport operations.

Traditional FOD detection methods, such as manual inspection and fixed surveillance systems, are time consuming, labor intensive, and prone to human error [2]. Automated sensor-based technologies, including radar and infrared systems, were developed to enhance detection range and reduce dependence on manual observation [3], [4]. However, these systems often face challenges under adverse weather conditions, low illumination, or complex runway environments. This leads to false alarms and inconsistent accuracy. Consequently, vision-based detection powered by deep learning has gained significant attention in recent years due to its adaptability and high performance.

Deep learning has transformed computer vision applications by enabling models to learn hierarchical features and spatial relationships directly from data [5]. Early FOD detection studies employed Convolutional Neural Networks (CNNs) such as the Visual Geometry Group (VGG) and ResNet architectures, achieving superior recognition accuracy compared to traditional image processing methods [6]. The introduction of single stage object detectors, particularly the You Only Look Once (YOLO) family, further improved inference speed and precision, making real-time runway monitoring feasible [7]. Research involving YOLOv4, YOLOv5 and YOLOv8 reported notable improvements in mean Average Precision (mAP) and detection reliability over conventional CNN architectures [8], [9], [10]. Additionally, large-scale datasets such as FOD-A [11], comprising more than 30,000 annotated runway images, have enabled robust model training and evaluation under diverse environmental conditions.

Despite these advancements, most existing systems assume that input images are sharp and free from distortions. In real world applications, especially in drone based runway surveillance, image quality often suffers from motion blur, camera vibration, or atmospheric interference [12]. Such degradations obscure object boundaries and fine details, thus severely affecting detection accuracy and increasing false positives. Recent research has explored integrating image restoration and object detection to overcome these challenges. Frameworks such as DeblurGAN [13] and MPRNet [14] have achieved promising results in restoring image clarity before detection, while attention guided deblurring models have shown improved feature recovery under blurred conditions [15]. However, many of these approaches focus on a single type of blur (typically Gaussian or Defocus) and require high computational resources, limiting their real-time applicability [16], [17]. Moreover, limited research has validated such frameworks on datasets simulating drone based runway imagery where both Gaussian and motion blur coexist [18], [19].

To analyse the restoration and detection performance under dual-blur conditions, a CNN based image restoration model was evaluated and compared with a baseline model such as DeblurGAN. Serving as a widely recognized benchmark, DeblurGAN effectively addresses both Gaussian and motion blur, making it suitable for comparative evaluation. Both restoration approaches were evaluated in combination with the YOLOv11n detection framework to assess detection reliability under degraded imaging conditions. This comparison provides a clearer understanding of how restoration quality influences detection accuracy and operational consistency in runway FOD imagery [19], [20]. The dataset primarily uses synthetically blurred images with parameters chosen to replicate real drone motion, while additional validation was performed using real blurred runway samples to ensure realistic evaluation. The computational performance was measured on an RTX 4060 GPU platform to verify real-time feasibility and low latency inference performance. The combined restoration and detection process improves the reliability of debris identification in dynamic and low visibility conditions, contributing to safer and more efficient runway operations.

The main contributions of this research are summarized as follows:

- Development of a deep learning based dual-blur deblurring framework capable of restoring both Gaussian and motion blurred runway imagery using synthetically generated blur to simulate real world degradation, validated on both synthetic and real blurred images.
- Comparative analysis of a CNN framework with a baseline DeblurGAN model, integrated within a YOLOv11n based detection pipeline for accurate and efficient FOD identification.
- Creation of a dedicated FOD dataset comprising 33 categories with synthetically blurred (Motion and Gaussian Blur) samples for comprehensive model training and evaluation.
- Demonstration of improved image clarity and enhanced detection performance through the restoration guided detection approach, supported by statistically validated gains in PSNR and SSIM metrics.
- Validation of the framework's suitability for near real-time deployment through GPU and embedded GPU benchmarks, confirming practical applicability for drone assisted runway inspection.

The following Section II presents the overall methodology, including dataset preparation, training setup, and the detection pipeline with image restoration model integration. Section III discusses the experimental results and analysis, while Section IV concludes the paper with key findings and future scope.

II. METHODOLOGY

This section focuses on the development of a CNN based image restoration and object detection framework designed for FOD identification on airport runways. The workflow integrates image deblurring and detection modules to improve detection accuracy and robustness under degraded visual conditions.

A dedicated dataset was prepared containing 33 categories of runway FOD objects, including metallic fragments, tools, stones, and other debris commonly found on runways, as depicted in Fig. 1. The dataset was divided into training, validation, and testing subsets in a ratio of 7:2:1, ensuring balanced representation across all categories. To simulate real world degradation, Gaussian and motion blur were synthetically applied to sharp images using random kernel sizes and motion vectors. Gaussian blur simulated optical defocus and atmospheric scattering, while motion blur represented camera vibration or drone movement during aerial surveillance. This approach provided a realistic approximation of the visual degradations typically encountered in runway imagery.

Fig. 1. Sample images from the FOD dataset.

To improve model generalization, data augmentation techniques such as random rotation, horizontal flipping, and brightness variation was applied [2]. The inclusion of synthetically blurred data ensured consistent evaluation under varying blur intensities without requiring real drone imagery.

The image restoration stage employs a CNN based framework that follows an encoder–decoder structure with integrated residual blocks for effective feature learning and edge preservation [3]. The encoder extracts multi scale features from the blurred inputs, while the residual connections stabilize training and support texture recovery. The decoder reconstructs deblurred images through transposed convolutions and nonlinear activations, generating sharp yet natural outputs. The model is optimized using a hybrid loss function that combines pixel level, perceptual, and total variation components as

$$L_{\text{total}} = L_1 + \lambda_p L_{\text{perceptual}} + \lambda_{tv} L_{\text{tv}} \tag{1}$$

Where, L_1 represents the mean absolute error, $L_{perceptual}$ measures high level feature similarity using a pretrained VGG16 network [4], and L_{tv} enforces smoothness through total variation regularization [5]. The weight parameters λ_p and λ_{tv} were set to 0.5 and 0.1, respectively. The model was trained for 100 epochs using the Adam optimizer with a learning rate of 1×10^{-4} and a batch size of 8. Training and testing were performed on a workstation equipped with an Intel(R) Core(TM) i9-13900H CPU, 16 GB RAM, and an NVIDIA GeForce RTX 4060 GPU (8 GB VRAM).

Each loss component in the hybrid objective function contributes differently to the restoration quality and overall image fidelity. The L_1 loss minimizes pixelwise intensity differences between the restored and ground truth images, directly improving the PSNR. The perceptual loss, derived from feature maps extracted by a pretrained VGG16 network, preserves high level semantic and structural information [4]. It enhances texture consistency and edge sharpness, leading to improved SSIM. The total variation (TV) loss acts as a regularization term that suppresses minor artifacts and noise while maintaining edge continuity [21]. It enforces smoothness across homogeneous regions and prevents unwanted pixel fluctuations near object boundaries. The combination of these three loss components enables sharper and cleaner restorations with higher PSNR and SSIM compared to using any single term alone.

For a fair performance comparison, the CNN based framework was evaluated alongside a baseline model, DeblurGAN [13]. DeblurGAN employs a conditional generative adversarial network (cGAN) with a ResNet based generator and a PatchGAN discriminator and effectively handles both Gaussian and motion blur through adversarial feature learning. Both restoration models were tested on the same blurred dataset to compare restoration quality in terms of PSNR, SSIM, and visual fidelity. The resulting restored images were subsequently evaluated for their impact on downstream FOD detection accuracy.

For object detection, the YOLOv11n (nano) architecture was trained from scratch using the sharp, high quality dataset to establish a baseline detection model. The YOLOv11n variant, known for its lightweight structure and high computational efficiency, was selected to balance detection accuracy and processing speed, making it suitable for real-time applications. The model was trained at an input resolution of 300 × 300, with a batch size of 8 and a learning rate of 0.001. The trained detector was subsequently used for inference on both blurred and restored images to assess the effect of restoration on detection performance. Evaluation metrics included Precision, Recall, mAP@0.50, and mAP@0.50-0.95.

After successful training, both the restoration and detection networks were combined into a single processing pipeline. During inference, blurred (Gaussian and Motion blur) images were first processed by either the CNN restoration model or the DeblurGAN baseline, and the enhanced outputs were then analyzed by YOLOv11n for FOD detection. Although the experiments were conducted primarily on a workstation GPU to verify real-time deployment capability. The system demonstrated stable inference performance and low latency, confirming its potential applicability for drone assisted runway inspection and aviation safety monitoring.

III. Result

This section presents the evaluation of the CNN and DeblurGAN models integrated with the YOLOv11n detector for FOD identification on runway imagery affected by Gaussian and motion blur. Quantitative and qualitative analyses were performed to assess how image restoration influences detection accuracy under both synthetically generated and real-world blurred conditions.

The restoration performance was evaluated using PSNR and SSIM, derived from the representative Gaussian and motion blur images shown in Fig. 2 and Fig. 3. The corresponding quantitative results are summarized in Table I and Table II. For Gaussian blur, the CNN model achieved an average PSNR of 32.07 dB and SSIM of 0.782, representing respective improvements of 7.2% and 11.2% over the blurred inputs. DeblurGAN achieved 31.66 dB and 0.774, showing gains of 5.7% and 10.1%. For motion blur, the PSNR increased from 29.12 dB (blurred) to 32.73 dB for the CNN model and 32.11 dB for DeblurGAN, corresponding to improvements of 12.1% and 10.1%, respectively. Similarly, SSIM improved from 0.691 to 0.817 for CNN (18.2% gain) and to 0.815 for DeblurGAN (17.9% gain). Both models effectively restored edge sharpness and structural detail, but the CNN based restoration demonstrated slightly higher consistency across different blur intensities and better preservation of fine textures, while DeblurGAN produced marginally sharper but occasionally noisier outputs.

Fig. 2. Comparison of Gaussian blurred and restored images.

Fig. 3. Comparison of motion blurred and restored images.

TABLE I. PSNR COMPARISON FOR GAUSSIAN AND MOTION BLUR IMAGES

Object	Blurred (dB)	DeblurGAN (dB)	CNN (dB)	Increase (%)
Gaussian Blur				
Soda Can	27.62	28.92	29.44	+4.7 / +6.6
Wrench	30.60	32.12	32.86	+5.0 / +7.4
Metal Cap	31.46	32.69	33.14	+3.9 / +5.3
Cutter	26.40	28.64	28.93	+8.5 / +9.6
Screwdriver	33.00	35.26	35.62	+6.8 / +7.9
Pen	30.59	32.24	32.45	+5.4 / +6.1
Average	29.61	31.66	32.07	+5.7 / +7.2
Motion Blur				
Soda Can	26.77	29.44	29.94	+10.0 / +11.9
Wrench	29.30	33.04	33.54	+12.8 / +14.5
Metal Cap	30.69	33.25	33.72	+8.4 / +9.9
Cutter	26.28	29.36	29.48	+11.7 / +12.2
Screwdriver	30.79	35.33	36.31	+14.8 / +18.0
Pen	30.86	33.22	33.41	+7.6 / +8.3
Average	29.12	32.11	32.73	+10.1 / +12.1

TABLE II. SSIM COMPARISON FOR GAUSSIAN AND MOTION BLUR IMAGES

Object	Blurred	DeblurGAN	CNN	Increase (%)
Gaussian Blur				
Soda Can	0.558	0.648	0.660	+16.1 / +18.3
Wrench	0.894	0.914	0.917	+2.2 / +2.6
Metal Cap	0.721	0.756	0.765	+4.9 / +6.1
Cutter	0.547	0.704	0.719	+28.7 / +31.4
Screwdriver	0.820	0.848	0.851	+3.4 / +3.8
Pen	0.680	0.774	0.779	+13.8 / +14.6
Average	0.703	0.774	0.782	+10.1 / +11.2
Motion Blur				
Soda Can	0.549	0.707	0.713	+28.8 / +29.9
Wrench	0.872	0.920	0.920	+5.5 / +5.5
Metal Cap	0.719	0.796	0.796	+10.7 / +10.7
Cutter	0.533	0.762	0.769	+43.0 / +44.3
Screwdriver	0.800	0.873	0.874	+9.1 / +9.3
Pen	0.676	0.831	0.829	+22.9 / +22.6
Average	0.691	0.815	0.817	+17.9 / +18.2

In addition to image quality metrics, the model's convergence was analyzed using the total loss function defined in Equation (1). After training 100 epochs, in (1) the individual losses for the CNN were recorded as $L_1 = 0.010765$, $L_{perceptual} = 0.458133$, and $L_{tv} = 0.021497$. With empirically set weights of $\lambda_p = 0.5$ and $\lambda_{tv} = 0.1$, the total loss was computed as

$$L_{total} = L_1 + \lambda_p L_{perceptual} + \lambda_{tv} L_{tv}$$

$$= 0.010765 + (0.5 \times 0.458133) + (0.1 \times 0.021497)$$

$$= 0.244198$$

Similarly, for DeblurGAN, the individual losses were recorded as $L_1 = 0.013939$, $L_{perceptual} = 0.484339$, and $L_{tv} = 0.019560$

$$L_{total} = 0.013939 + (0.5 \times 0.484339) + (0.1 \times 0.019560)$$

$$= 0.2580645$$

The CNN achieved a 5.4% lower total loss compared to DeblurGAN, demonstrating higher stability and better convergence. This lower total loss directly correlates with improved perceptual fidelity and reduced reconstruction artifacts in restored outputs. Both models enhanced visual clarity, as illustrated in Fig. 2 and Fig. 3. The CNN model produced smoother, more natural restorations, whereas DeblurGAN provided sharper, high-contrast outputs that occasionally exhibited minor artifacts.

For real blurred images, as illustrated in Fig. 4, both restoration methods enhanced perceptual quality but showed different quantitative trends. The PSNR increased from approximately 28.08 dB in blurred inputs to 28.92 dB for the CNN model, indicating a 3.0% improvement, while DeblurGAN slightly reduced PSNR to 26.95 dB, corresponding to a 4.0% decrease. This reduction, particularly observed in clamp part and nut images, suggests that DeblurGAN occasionally amplifies edge contrast at the expense of luminance consistency. In contrast, SSIM improved from 0.502 to 0.581 for the CNN model and to 0.593 for DeblurGAN, showing relative gains of 15.7% and 18.1%, respectively. The CNN model maintained better brightness stability and texture preservation, whereas DeblurGAN produced visually sharper but locally inconsistent restorations.

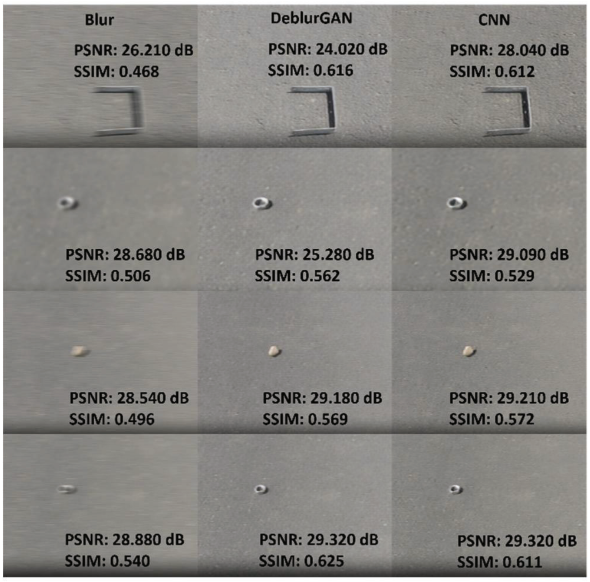

Fig. 4. Comparison of real blurred and restored images.

A statistical analysis was conducted on 330 test images, including both Gaussian and motion blur cases. As shown in Table III, the CNN based framework achieved the highest mean PSNR of 34.94 dB and SSIM of 0.8473. This represents an improvement of 11.6% and 9.1% over the blurred input (31.30 dB, 0.7765), and 4.9% and 0.4% over DeblurGAN

(33.31 dB, 0.8442), respectively. The narrow 95% confidence intervals (±0.13 dB for PSNR and ±0.0022 for SSIM) indicate stable and consistent restoration performance across all test Images.

TABLE III. PSNR AND SSIM (MEAN ±95% CI) FOR BLUR AND RESTORED IMAGES

Image	PSNR (dB) ± 95% CI	SSIM ± 95% CI
CNN Restored	34.94 ± 0.13	0.8473 ± 0.0022
DeblurGAN Restored	33.31 ± 0.13	0.8442 ± 0.0022
Blurred	31.30 ± 0.13	0.7765 ± 0.0022

The computational performance comparison was performed on an NVIDIA GeForce RTX 4060 GPU. Both models were evaluated using 330 test images, and the reported metrics in Table IV represent the average results obtained from three independent trials. The CNN based framework achieved an average inference time of 0.0072 s per image, corresponding to a throughput of approximately 138 FPS, with a GPU utilization of around 31.5%. In contrast, DeblurGAN required 0.0141 s per image (≈71 FPS) and utilized 46.1% of the GPU under identical conditions. Similarly, the integrated CNN+YOLO pipeline achieved an inference time of 0.0099 s (≈101 FPS) with 48.4% GPU utilization, whereas DeblurGAN+YOLO operated at 0.0175 s (≈57 FPS) and 67.1% utilization. This shows that the CNN based restoration framework operates nearly twice as fast as DeblurGAN while consuming fewer computational resources, confirming its suitability for real-time FOD detection and runway monitoring applications.

TABLE IV. COMPUTATIONAL PERFORMANCE COMPARISON

Model	Average Inference Time (s)	FPS	VRAM (GB)	GPU Utilization (%)
CNN+YOLO	0.0099	101.01	0.16	48.37
DeblurGAN+YOLO	0.0175	57.10	0.19	67.08
CNN	0.0072	138.12	0.09	31.48
DeblurGAN	0.0141	70.93	0.11	46.14

The restored images were next evaluated using the YOLOv11n object detection model trained from scratch on 33 FOD categories. The results are presented in Table V. The detector achieved a Precision of 0.973, Recall of 0.936, mAP@0.50 of 0.954, and mAP@0.50–0.95 of 0.596. Since the detection model was trained on sharp images, the same model was used for inference on both blurred and restored datasets to evaluate the effect of image restoration on detection accuracy. Notable improvements were observed in detection confidence and bounding box precision for restored inputs, confirming that enhanced image clarity directly contributes to more reliable detection.

TABLE V. YOLOV11N TRAINING PERFORMANCE METRICS

Metric	Value
Precision	0.973
Recall	0.936
mAP@0.5	0.954
mAP@0.5–0.95	0.596

Detection results were analysed on both synthetic and real blurred datasets to evaluate the influence of restoration on YOLOv11n based object recognition. For synthetic Gaussian blurred scenes, restoration substantially improved detection confidence and localization precision across all categories. As shown in Fig. 5, the average detection confidence increased from 0.68 in blurred images to 0.74 with DeblurGAN and 0.75 with the CNN model, corresponding to relative gains of 8.8% and 10.3%, respectively. The screwdriver and battery classes exhibited the highest improvement, with confidence increases of 26.5% and 7.9% after CNN based restoration.

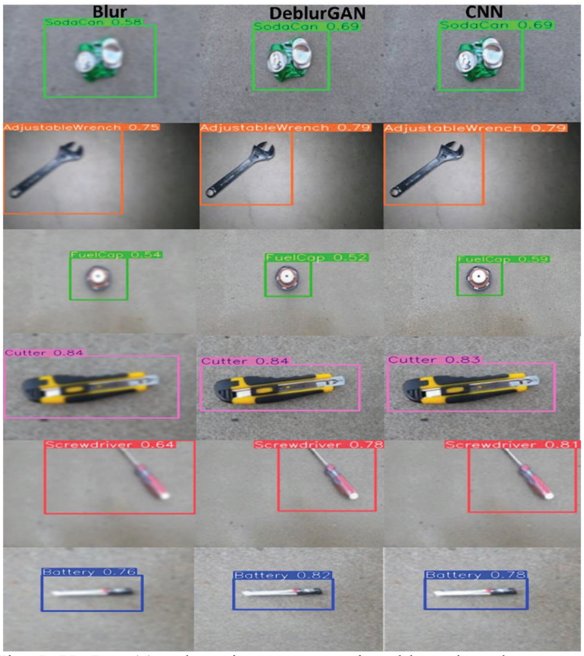

Fig. 5. YOLOv11n detection on gaussian blurred and restored images.

For motion blurred images, visual restoration had an even stronger impact. As shown in Fig. 6, the blurred image of a screwdriver was incorrectly classified as a hammer with a low confidence of 0.26. After restoration, the correct detection confidence improved to 0.79 for both DeblurGAN and the CNN model. Similarly, the soda can and adjustable wrench confidence levels increased from 0.55 and 0.71 (blur) to 0.72 and 0.80 (CNN restored), reflecting 30.9% and 12.7% gains, respectively.

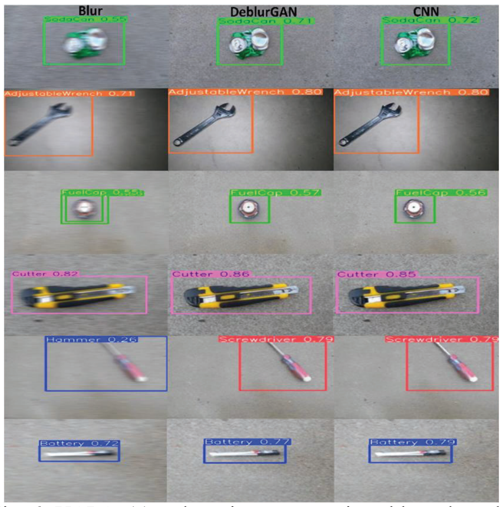

Fig. 6. YOLOv11n detection on motion blurred and restored images.

These results indicate that restoration effectively recovers spatial and textural cues, reducing misclassification and improving object detection confidence, particularly under motion blur.

In real blurred scenarios, as illustrated in Fig. 7, objects such as nuts and ClampPart were missed in the blurred inputs due to the loss of fine texture details and reduced edge contrast. After applying restoration, these objects were successfully identified, with detection confidence improving from 0.33–0.40 in blurred images to 0.41–0.59 in CNN restored outputs and 0.40–0.49 in DeblurGAN results. This corresponds to an average confidence increase of approximately 9.5% for the CNN model and 7.8% for DeblurGAN. The CNN based restoration provided clearer boundaries and higher detection confidence, whereas DeblurGAN produced sharper textures but introduced minor localization variations. These results confirm that both restoration methods enhance object visibility and detection reliability under real-world low contrast conditions.

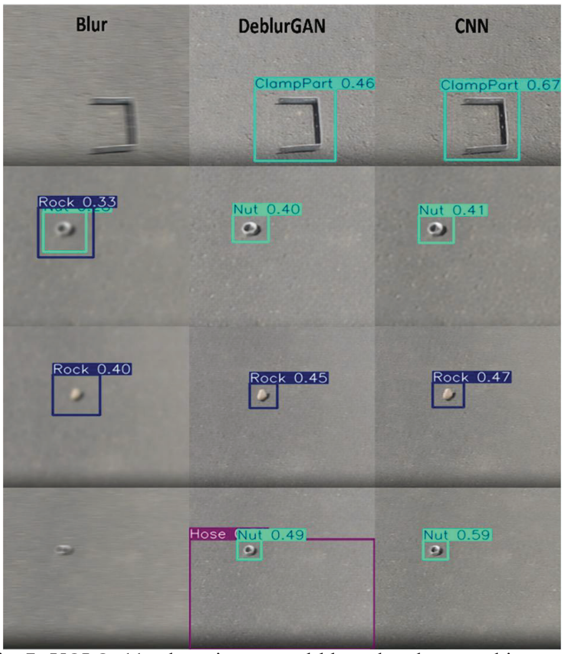

Fig. 7. YOLOv11n detection on real blurred and restored images

The CNN based restoration framework consistently outperformed DeblurGAN across all evaluation parameters, achieving higher PSNR, SSIM, and detection confidence while requiring less computational time. Its inference speed was nearly twice as fast, confirming greater suitability for real-time deployment. Although DeblurGAN generated slightly sharper visual edges in a few high-motion instances, it also introduced minor artifacts and required higher GPU utilization. The CNN model maintained stable reconstruction with smoother textures, balanced sharpness, and improved detection accuracy, particularly in cases of combined Gaussian and motion blur. These characteristics highlight its effectiveness as a lightweight yet reliable restoration component for integration with YOLOv11n, enabling accurate Foreign Object Debris identification under diverse and degraded visual conditions.

IV. CONCLUSION

This study presented deep learning based deblurring frameworks integrated with the YOLOv11n detector for FOD identification under Gaussian and motion blur conditions. Both CNN and DeblurGAN models significantly enhanced image clarity and detection performance. The CNN framework achieved average PSNR and SSIM gains of 7.2% and 11.2% for Gaussian blur, and 12.1% and 18.2% for motion blur, outperforming DeblurGAN by 5–7%. Loss convergence analysis showed that the CNN model achieved a lower total loss by 5.4% compared to DeblurGAN.

For real blurred imagery, PSNR improved from 28.08 dB to 28.92 dB with the CNN model, whereas DeblurGAN reached 26.95 dB. SSIM increased from 0.502 to 0.581 for CNN and 0.593 for DeblurGAN, confirming enhanced structural consistency under practical conditions. Detection confidence improved by 10.3% in Gaussian blur and by 9.5% and 12.6% for real blurred cases using CNN restoration, while DeblurGAN achieved comparable performance. The CNN+YOLO framework achieved an inference speed of 101 FPS compared to 67.08 FPS for DeblurGAN+YOLO, demonstrating an effective balance between restoration accuracy, detection reliability, and computational efficiency for real-time runway safety monitoring. Future work will focus on extending the model to handle multiple image degradations, such as noise, illumination variation, and weather effects, to enable a unified and robust restoration detection system for real-time aviation safety applications.

REFERENCES

[1] M. Kumari, S. Patil, and K. Sharma, "Deep learning-based foreign object debris detection for airport runway safety," *Engineering Research Express*, vol. 6, no. 4, 2024.

[2] A. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," *Journal of Big Data*, vol. 6, no. 60, pp. 1–48, 2019.

[3] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2015.

[4] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," *Proc. Eur. Conf. Comput. Vis. (ECCV)*, pp. 694–711, 2016.

[5] R. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear Phenomena*, vol. 60, no. 1–4, pp. 259–268, 1992.

[6] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, real-time object detection," *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 779–788, 2016.

[7] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "YOLOv4: Optimal speed and accuracy of object detection," *arXiv preprint arXiv:2004.10934*, 2020.

[8] S. Nah, T. H. Kim, and K. M. Lee, "Deep multi-scale convolutional neural network for dynamic scene deblurring," *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 257–265, 2017.

[9] L. Zhang and X. Luo, "Automated foreign object debris detection using runway surveillance systems," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 6, pp. 1822–1834, 2018.

[10] A. Kumar, A. Mittal, A. Dixit, and P. Agrawal, "Deep Learning Based FOD Detection for Runway Health Monitoring," *Proc. 1st Int. Conf. on AIML-Applications for Engineering & Technology (ICAET)*, pp. 1-6, 2025.

[11] A. Gao, P. Zhu, and T. Wang, "Radar-based runway FOD detection system with adaptive thresholding," *Sensors*, vol. 18, no. 11, pp. 3654–3665, 2018.

[12] J. Liu and C. Chen, "Infrared imaging-based debris detection under complex backgrounds," *Infrared Physics & Technology*, vol. 105, 2020.

[13] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 770–778, 2016.

[14] G. Jocher, A. Chaurasia, and J. Qiu, "YOLOv5: Implementation and performance improvements," *Zenodo*, 2022.

[15] Y. Wang, F. Wang, and Z. Liu, "FOD-A: A large-scale dataset for foreign object debris detection on airport runways," *IEEE Access*, vol. 11, pp. 7420–7431, 2023.

[16] T. Kupfer, M. Borth, and D. Klatte, "Image degradation in UAV-based surveillance: Analysis of motion and defocus blur," *Remote Sensing*, vol. 13, no. 19, pp. 3820–3835, 2021.

[17] O. Kupyn, V. Budzan, M. Mykhailych, D. Mishkin, and J. Matas, "DeblurGAN: Blind motion deblurring using conditional adversarial networks," *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 8183–8192, 2018.

[18] S. Zamir, A. Arora, and S. Khan, "MPRNet: Multi-stage progressive image restoration," *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 14821–14831, 2021.

[19] H. Chen, J. Gu, and X. Chen, "Attention-guided image restoration for motion and defocus blur removal," *IEEE Trans. Image Process.*, vol. 32, pp. 1542–1556, 2023.

[20] F. Liu, Z. Yu, and Y. Tang, "A review of UAV-based visual monitoring systems for airport runway safety," *Aerospace*, vol. 11, no. 3, pp. 153–166, 2024.

[21] X. Liu, Y. Li, and J. Wang, "Image restoration using total variation regularized deep convolutional neural networks," *arXiv preprint arXiv:1810.12864*, 2018.