Automated Multi-Camera Inspection System for Aircraft

Mark David Rice, Gu Ying, Kelvin Wei Lim, Qing Yu Hoo, Liyuan Li,
Lee Jue Ying, Jacky Jie Wei Tan, Lai Xing Ng, Jamie Ng
Institute for Infocomm Research (I^{2}R), A*STAR
{Mark_David_Rice, guy, kelvin_lim, Hoo_Qing_Yu, Li_Liyuan, Lee_Jue_Ying, Jacky_Tan, ng_lai_xing, Jamie_Ng}
@a-star.edu.sg

Abstract

In this paper, we present the development of an automated visual inspection system designed to detect defects on the upper surface of an aircraft airframe. Specifically, the system employs a multi-camera PTZ (Pan-Tilt-Zoom) set-up to capture and process images from designated regions. Custom developed software manages path planning and camera localization, while a hybrid-AI framework is integrated to identify various defect types, including missing and damaged components. The demonstration highlights the system's detection capabilities and prototype functionalities using a large aircraft model, supported by a user interface to monitor progress and visualize results. To help validate this work, performance evaluations were conducted using selected multimodal and object detection models.

Introduction

Visual inspection involves a series of pre- and post-flight checks to determine the airworthiness of an aircraft. This can range from structural damage to physical obstructions, and misaligned aircraft parts. Historically, it is considered a direct and cost-effective means of assessing aircraft safety (FAA 1997). Conventional inspection tools include the use of a flashlight or mirror, relying heavily on human expertise and domain knowledge. In contrast, automated AI inspection offers the potential to reduce or reallocate manpower resources in maintenance, repair, and overhaul (MRO). Not only to modernize workflow processes, but also to provide greater traceability, accountability, and operational safety.

Research studies have included climbing robots that are able to physically attach to an aircraft surface (e.g., Ramalingam et al. 2019; Marquette et al. 2025), ground-based robots (e.g., Leiva et al. 2017), and drone-based systems (e.g., Liu et al. 2022; Ali et al. 2025). Defect detection methods include Mask R-CNN to identify dent damage (Doğru et al. 2020), a modified YOLO-based network to detect irregular shaped objects (Huang et al. 2024), and the use of TensorFlow object detection models to identify features like paint damage and scratches (Plastropoulos et al. 2024).

Yet, despite these advancements, existing solutions lack information on their effectiveness to operate across different airframe structures and materials. To illustrate, drones may be constrained by battery duration, while climbing robots may be limited by their inspection speed (Suvittawat et al. 2025). Moreover, the efficacy of models under varying lighting and surface quality conditions, and diverse defect characteristics, often remains uncertain due to a lack of empirical data. As such, the focus of our work is to consider a more holistic approach to the inspection cycle, with wider real-time scan coverage and reduced inspection times compared to single camera solutions.

System Description

Our overhead multi-camera system automates the inspection process, while keeping the human-in-the-loop to further review any inferred defects. System components include:

Hardware: This consists of a series of ONVIF compliant PTZ (Pan-Tilt-Zoom) cameras, mounted onto a hangar ceiling or walls. These are motorized cameras that allow for controlled lens movement, each connected via a network cable to a PoE+ (Power over Ethernet Plus) switch for data transfer. Camera positions and their quantity are aided through a developed software simulation tool. Here, virtual cameras can be added to a scene to generate the scan coverage of a given aircraft model. Specifically, an estimated scan coverage percentage, which can provide useful reference information prior to any physical installation.

Path planning and localization: To start the inspection process, a one-time manual localization is required to calibrate the PTZ cameras to a parked aircraft. For each camera, a user selects several surface points from a live video stream of the aircraft and matches the same points to a 3D model.

Following 2D-3D matching, camera parameters are recorded and used in a localization algorithm to compute their position and distance. These calibrated cameras serve as input for path planning, which generates a list of pan-tilt-zoom instructions for each camera to capture images at sequential scan points. This process ensures sufficient scan coverage and optimizes image capture routes to minimize mechanical camera movements.

Defect detection: Images captured by the PTZ cameras are encrypted and stored in a scan database before being image processed. The system can currently detect up to 30 defect-types, ranging from missing screws, to corrosion, and tire damage. To address practical, real-world challenges like imbalanced datasets, limited training samples, and high false positive rates, a *Hybrid-AI framework* is proposed. This combines Deep Learning (DL) and Machine Learning (ML) models with Bayesian reasoning.

Interaction: A user can monitor progress through an interface, as scan results highlight the inferred defects and their position on an aircraft. These are visualized on a digital twin (Figure 2). Interactive features include selecting the aircraft model, the live feed from a given camera, and the types of defects to review. In addition, there are options to examine prior image scans.

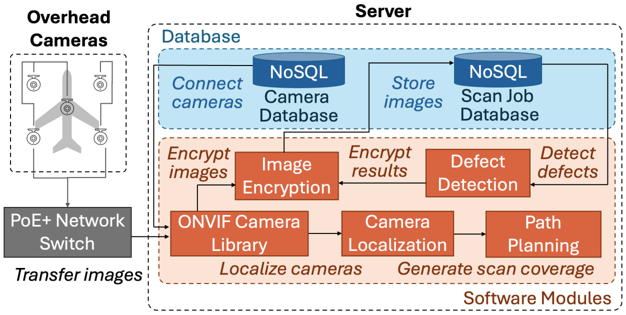

Figure 1: Overview of the camera system architecture.

Evaluation

To illustrate, we developed separate models for detecting *missing screws*, *hairline cracks*, and *open latches*. For missing screws, we implemented a multi-stage detection approach. This includes an initial coarse detection using a Linear SVM, followed by active circle fitting and re-sampling, and fine classification using a Kernel SVM. For crack detection, we developed an algorithm to determine the spatial orientation of associated crack patches. Whilst for open latch detection, we employed a Linear SVM combined with active circle fitting to detect and locate related buttons, enhancing recognition accuracy through Bayesian reasoning.

For benchmarking, we built an image dataset from a commercial aircraft. This contained 332, 157, and 144 images of missing screws, cracks and open latches for training, and 125, 73, and 52 images for testing. We compared our solutions with 1) *Faster R-CNN* (Ren et al. 2015), a DL model that localizes image objects using region proposals, 2) *Grounding DINO* (Liu et al. 2024), an object detection model that uses text prompts, and 3) *GPT-4o* (OpenAI 2024), a frontier Multimodal Large Language Model (MLLM). Training was undertaken with Faster R-CNN, while Grounding DINO, GPT-4o, and our models were evaluated with zero-shot learning.

The results observed in Table 1 show that our models generated the highest F1 scores for all defect types, achieving a good balance between precision and recall. This is crucial for accurate defect classification. Alternatively, object detection models like Grounding DINO performed high in recall, but lower in precision.

Model	Metric	Missing screws	Hairline cracks	Open latches
Faster R-CNN	P	0.627	0.564	0.778
	R	0.879	0.612	0.978
	F1	0.732	0.587	0.867
Grounding DINO	P	0.251	0.489	0.582
	R	**1**	0.752	**1**
	F1	0.401	0.593	0.736
GPT-4o	P	0.610	0.543	**0.987**
	R	0.528	0.432	0.314
	F1	0.566	0.481	0.476
Ours	P	**0.904**	**0.882**	0.911
	R	0.807	**0.900**	0.960
	F1	**0.853**	**0.891**	**0.935**

Table 1: Comparison of results (highest scores are in bold).

Demonstration

This consists of a replica aircraft model, scaled at 1:10. PTZ cameras are mounted on adjustable tripods, with options to vary the number of cameras used for broad or targeted surface scans. For visibility and audience engagement, the user interface is displayed on a large screen, while defects are physically simulated on the surface of the aircraft model. These include loose screws, corroded screws, and bird strike damage. The demonstration showcases the camera calibration process, surface scanning, and defect detection results.

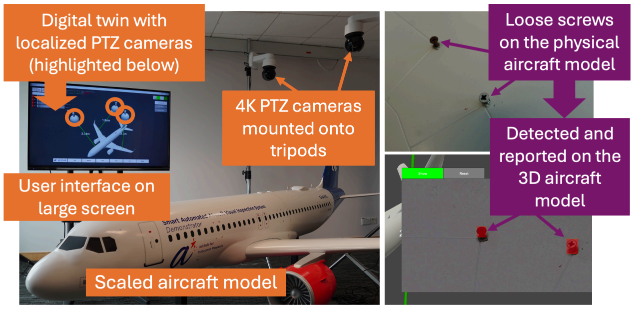

Figure 2: Illustration of the physical set-up.

Acknowledgements

Special thanks to Marcus Wan and Keith Tang for their technical contributions. This research was supported by the Agency for Science, Technology and Research (A*STAR) under the Smart Automated Aircraft Visual Inspection System (SAAVIS) project.

References

Ali, M.A.H., Zulkifle, M.Z.A., Nik Ghazali, N.N., Apsari, R., Meor Zulkifli, M.M.F., and Alkhedher, M. 2025. AI-driven UAV with image processing algorithm for automatic visual inspection of aircraft external surface. *Scientific Reports*, 15:19581.

Doğru, A., Bouarfa, S., Arizar, R., and Aydoğan, R. 2020. Using convolutional neural networks to automate aircraft maintenance visual inspection. *Aerospace*, 7(12):171.

FAA. 1997. *Advisory circular AC 43-204 - Visual inspection for aircraft* (Cancelled). U.S. Department of Transportation.

Huang, B., Ding, Y., Liu, G., Tian, G., and Wang, S. 2024. ASD-YOLO: An aircraft surface defects detection method using deformable convolution and attention mechanism. *Measurement*, 238:115300.

Leiva, J.R., Villemot, T., Dangoumeau, G., Bauda, M.-A., and Larnier, S. 2017. Automatic visual detection and verification of exterior aircraft elements. In *2017 IEEE International Workshop of Electronics, Control, Measurement, Signals and their Application to Mechatronics (ECMSM)*.

Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Jiang, Q., Li, C., Yang, J., Su, H., Zhu, J., and Zhang, L. 2024. Grounding DINO: Marrying DINO with grounded pre-training for open-set object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*: 38-55.

Liu, Y., Dong, J., Li, Y., Gong, X., and Wang, J. 2022. A UAV-based aircraft surface defect inspection system via external constraints and deep learning. *IEEE Transactions on Instrumentation and Measurement*, 71:1-15.

Marquette, W., Schultz, K., Jonnalagadda, V., Wong, B., Garbini, J., and Devasia, S. 2025. Semi-autonomous teleoperation using differential flatness of a crane robot for aircraft in-wing inspection. *IEEE Robotics and Automation Letters*, 10(3): 2742-2749.

OpenAI. 2024. *GPT-4o [Large language model]*. https://openai.com/index/hello-gpt-4o/. Accessed: 12 November 2025.

Plastropoulos, A., Bardis, K., Yazigi, G., Avdelidis, N.P., and Droznika, M. 2024. Aircraft skin machine learning-based defect detection and size estimation in visual inspections. *Technologies*, 12(9):158.

Ramalingam, B., Manuel, V.-H., Elara, M.R., Vengadesh, A., Lakshmanan, A.K., Ilyas, M., and James, T.J.Y. 2019. Visual inspection of the aircraft surface using a teleoperated reconfigurable climbing robot and enhanced deep learning technique. *International Journal of Aerospace Engineering*, 5137139.

Ren, S., He, K., Girshick, R., and Sun, J. 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Proceedings of the 29th International Conference on Neural Information Processing Systems (NIPS)*: 91-99.

Suvittawat, N., Kurniawan, C., Datephanyawat, J., Tay, J., Liu, Z., Soh, D.W., and Ribeiro, N.A. 2025. Advances in aircraft skin defect detection using computer vision: A survey and comparison of YOLOv9 and RT-DETR performance. *Aerospace*, 12(4):356.