

Deep-learning-based Character Recognition of Aerospace Alloy Components for Automated Quality Inspection

Ming-Chan Lee,* Guan-Liang Lin, Guan-Chen Jian,
Ting-Rui Lin, and Chia-Chuan Tai

Department of Electrical Engineering, National Kaohsiung University of Science and Technology,
Kaohsiung 807, Taiwan

(Received November 18, 2025; accepted December 4, 2025)

Keywords: aerospace component, automated recognition, nickel-cobalt alloy, deep learning

As aerospace engine manufacturing advances toward intelligent automation, traditional manual component identification methods suffer from low efficiency and high error rates. We present a vision-based automated recognition system that integrates image processing and deep learning algorithms for the identification of aerospace engine components. Nickel–cobalt alloy components (Inconel 718) are used as test samples. The system employs optical character recognition (OCR) and You Only Look Once version 8 (YOLOv8) algorithms for character detection and classification. Image preprocessing includes grayscale conversion, Gaussian blur, edge detection, and morphological operations to enhance image quality. To overcome the limitations of single-frame recognition under varying lighting conditions and complex backgrounds, a multiframe fusion mechanism is implemented to enhance stability and accuracy. Comparative experiments were conducted to evaluate algorithm performance in terms of recognition accuracy, processing speed, and computational efficiency. The system integrates a programmable automation controller (PAC) and positioning sensors for automated operation. Experimental results demonstrate an approximately 96% recognition accuracy with a processing time of 0.35 seconds per component, confirming the effectiveness of this approach for automated quality inspection in aerospace manufacturing environments.

1. Introduction

Aerospace manufacturing prioritizes traceability, stability, and safety, and consequently, in-line component identification and visual inspection increasingly rely on real-time computer vision (CV) systems. The You Only Look Once (YOLO) series, particularly YOLOv8, offers an optimal balance between detection speed and accuracy, making it suitable for edge deployment.⁽¹⁾ Meanwhile, aerospace manufacturing is increasingly adopting data-driven approaches to reduce human variability and improve manufacturing yield.

*Corresponding author: e-mail: mclee@nkust.edu.tw
<https://doi.org/10.18494/SAM6048>

Recent studies have led to significant advances in real-time object detection algorithms and integrated detection–recognition systems. Hussain systematically reviewed the real-time capability and extensibility of YOLO variants,⁽²⁾ leading to subsequent improvements such as DWRB module-SADown module-LASPPF structure (DSL)-YOLO for enhanced small-object and occlusion detection⁽³⁾ and YOLO- Average SPP-Fast Block, C2f_SimDCNv2, and Space-to-Depth (ADS) for balanced defect detection efficiency.⁽⁴⁾

In automatic optical inspection (AOI) applications, deep learning approaches have proven effective even with limited training data.⁽⁵⁾ YOLOv8-based systems have been successfully applied to visual anomaly quantification,⁽⁶⁾ while multiscale residual attention networks address challenges in tiny character and oblique text recognition.⁽⁷⁾ Two-stage detection–optical character recognition (OCR) systems have been developed for various industrial scenarios, including expiry-date reading under varying illumination,⁽⁸⁾ embossed tire text recognition, geometric alignment pipelines,⁽⁹⁾ and small-character recognition with contrast enhancement and template matching.⁽¹⁰⁾

In automatic number plate recognition (ANPR), YOLOv8–OCR integration has demonstrated high-speed processing capabilities.^(11,12) In comparative studies, cloud-based OCR and handwritten text recognition services have been evaluated for optimal tool selection.⁽¹³⁾ Methodologically, dual-branch architectures and global feature extraction have been explored to improve surface defect inspection performance.^(14,15)

Aerospace components manufactured from advanced alloys—including titanium, nickel-cobalt-based superalloys, and high-strength steels—pose significant optical inspection challenges. Nickel–cobalt alloy components (Inconel 718) are particularly challenging to inspect owing to intense specular reflection, low contrast, and surface oxidation or coatings, which reduce the signal-to-noise ratio and edge definition.⁽¹⁶⁾ Traceability markings are often small and irregularly oriented, requiring multiscale, pose-invariant recognition.⁽¹⁷⁾ Metal surface character recognition is sensitive to pose and illumination variations, necessitating precise fixturing and process control. While cloud-based OCR offers convenience, industrial deployment demands real-time latency and data governance compliance. Edge computing on programmable logic/automation controllers (PLC/PAC) reduce latency but require lightweight models owing to computational constraints, and the industrial Internet of Things (IIoT) integration adds communication overhead.

Recent research in industrial character recognition has predominantly employed general-purpose deep learning models or cloud-based OCR services to enhance inspection accuracy. However, these approaches often struggle with the specific optical and sensing challenges of aerospace materials—such as the intense specular reflection and low contrast exhibited by Inconel 718 surfaces—and frequently lack the necessary integration with industrial control systems for real-time production monitoring. Furthermore, existing vision-based inspection systems typically require substantial network bandwidth and computational resources, limiting their deployment feasibility in edge computing environments where immediate feedback and deterministic response times are critical for production line operations.^(14–17) In this study, we present an industrial vision system for automated character recognition on Inconel 718 aerospace components using image sensing and deep learning techniques. The system integrates the

YOLOv8-nano (YOLOv8-n) object detection architecture for efficient character region localization, utilizing targeted image preprocessing on regions of interest (ROIs). Character recognition incorporates Microsoft Azure OCR with rule-based disambiguation logic ($O \leftrightarrow 0$, $I \leftrightarrow 1$) and spatial layout analysis for correct sequence reconstruction. The complete sensing pipeline operates on PLC/PAC with event-driven triggering and IIoT connectivity for real-time production monitoring. Experimental validation reveals reliable performance on high-gloss nickel alloy surfaces, achieving 96.2% detection accuracy, 98% recognition accuracy under controlled conditions, and 85–90% accuracy in actual production environments, with an average processing time of 0.35 s per component, confirming the practical feasibility for industrial deployment in aerospace manufacturing quality control.

2. Materials and Methods

We develop an integrated detection–OCR system for automated character recognition on aerospace component identity markings. The investigation focuses on Inconel 718 engine hanger components as representative specimens of nickel-based superalloy components. The system (Fig. 1) employs YOLOv8 for character region detection, applies targeted preprocessing within detected bounding boxes, and performs OCR using the Microsoft Azure Computer Vision Read Application Programming Interface (API). The system operates with event-driven triggering at 3–5 s intervals to accommodate network latency while maintaining production throughput. Hardware and control architecture details are provided in Fig. 2.

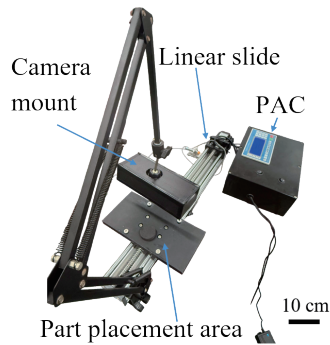


Fig. 1. (Color online) Hardware and control architecture of the PAC-controlled inspection system.

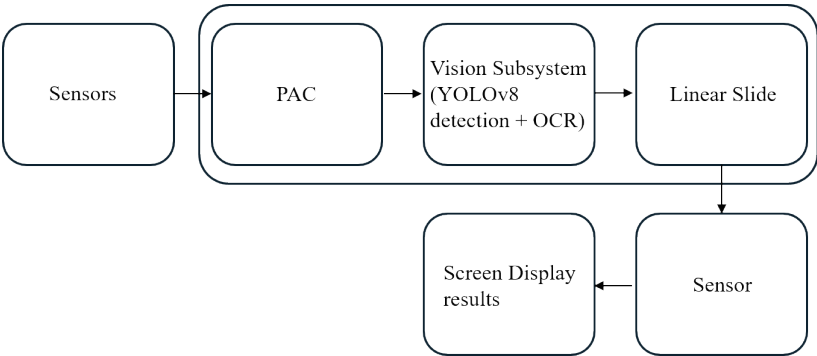


Fig. 2. Block diagram of the PAC-controlled vision system platform.

2.1 Character region detection using YOLOv8

2.1.1 Image acquisition and labeling

A custom dataset was constructed from images captured on industrial production lines. Alphanumeric characters and symbols on Inconel 718 component surfaces were manually annotated with bounding boxes and class labels. The Roboflow platform was employed for duplicate image detection, annotation quality verification, and dataset preparation in YOLO-compatible format. During dataset partitioning, samples with diverse lighting conditions and surface glare were intentionally retained in the test set to evaluate model robustness under varying illumination while maintaining fixed camera positioning. The final dataset comprises 540 images. Images were acquired using a camera [PCM-550, Nakay Co., Ltd. (KINYO), New Taipei City, Taiwan]. The dataset was partitioned into training (70%), validation (20%), and testing (10%) subsets, with the test set specifically including challenging illumination scenarios to assess real-world deployment performance.

2.1.2 Local preprocessing

Inconel 718 surfaces exhibit characteristic high specular reflectivity and low contrast, which significantly degrade detection accuracy and character recognition performance. To mitigate these effects, targeted preprocessing is applied exclusively within each detected bounding box rather than to the entire image, reducing computational overhead while improving local image quality. The preprocessing pipeline consists of four sequential operations defined by Eqs. (1)–(4).

$$I_{gray} = 0.299R + 0.587G + 0.114B \quad (1)$$

Here, I_{gray} is the grayscale intensity, R , G , and B are the red, green, and blue channel intensities (8-bit values ranging from 0 to 255), and the coefficients 0.299, 0.587, and 0.114 are perceptual luminance weighting factors based on human visual sensitivity.

$$I_{smooth} = I * G\sigma \quad (2)$$

$G\sigma$ is a two-dimensional Gaussian kernel with standard deviation σ , and $*$ denotes the convolution operation. This step suppresses high-frequency noise while preserving edge information.

$$t^* = \arg \max_t \sigma_b^2(t) \quad (3)$$

Here, t^* is the optimal threshold that maximizes the interclass variance $\sigma_b^2(t)$, defined as

$$\sigma_b^2(t) = \omega_0(t) \omega_1(t) (\mu_0(t) - \mu_1(t))^2. \quad (4)$$

Here, ω_0 and ω_1 are the probabilities of background and foreground classes, respectively, and μ_0 and μ_1 are their corresponding mean intensities.

2.1.3 Detector

YOLOv8-n, the most lightweight variant in the YOLOv8 series, is employed for real-time character region detection. The network architecture comprises three main components: a CSPDarknet backbone with CSP bottleneck with 2 convolutions (C2F) and spatial pyramid pooling fast (SPPF) modules for hierarchical feature extraction, a path aggregation network with a feature pyramid network (PAN-FPN) neck for multiscale feature fusion, and a decoupled detection head that separately predicts bounding box coordinates and class confidence scores. This lightweight configuration enables real-time inference on resource-constrained industrial controllers while maintaining sufficient accuracy for character localization.

2.1.4 Training setup and evaluation

Models are trained at 640×640 pixel resolution with a batch size of 16 on compute unified devices architecture (CUDA)-enabled graphics processing units (GPUs). Data augmentation includes mosaic augmentation (probability 0.25), random rotation ($\pm 7^\circ$), and hue–saturation–value (HSV) saturation variation (factor 0.4) to enhance robustness under varying illumination and orientation. Mosaic augmentation is turned off during the final 20 epochs to stabilize convergence. Training durations of 50, 100, 200, and 300 epochs are compared to determine the optimal balance between performance and cost. Early stopping with 80-epoch patience prevents overfitting; training terminates when the validation mean average precision ($mAP@0.5$) improvement falls below 0.001 for 80 consecutive epochs, retaining the best checkpoint.

Model performance is evaluated on both validation and test sets using standard object detection metrics. Intersection over union (IoU), precision, and recall are computed according to

$$IoU = \frac{|B \cap G|}{|B \cup G|}, \quad (5)$$

$$Precision = \frac{TP}{TP + FP}, \quad (6)$$

$$Recall = \frac{TP}{TP + FN}, \quad (7)$$

where B is the predicted box and G the ground-truth box; TP , FP , and FN denote true positives, false positives, and false negatives, respectively. mAP at IoU threshold 0.5 ($mAP@0.5$) and

across IoU thresholds from 0.5 to 0.95 with 0.05 increments ($mAP@0.5:0.95$) are reported as primary performance indicators.

2.1.5 Model selection

YOLOv8-nano is selected for three primary reasons: (i) its decoupled detection head and optimized loss functions provide superior localization accuracy for small-scale, low-contrast character markings typical of aerospace alloy surfaces; (ii) its lightweight architecture and native Open Neural Network Exchange (ONNX)/TensorRT export capability enable deployment on resource-constrained PLC/PAC platforms; and (iii) it achieves the optimal balance between detection accuracy and inference speed required for real-time industrial inspection. Comparative studies have demonstrated that YOLOv8 achieves higher mAP and faster inference than YOLOv5 under similar configurations.

2.2 Character recognition system

Following character region detection, a two-stage recognition pipeline extracts alphanumeric content from detected bounding boxes. First, individual regions are cropped, preprocessed, and spatially ordered. Second, OCR is performed using Microsoft Azure Computer Vision Read API, followed by rule-based postprocessing to correct common misclassifications.

2.2.1 Region extraction and spatial ordering

Each detected bounding box is cropped from the source image and preprocessed using the operations described in Sect. 2.1.2. Recognized character tokens are spatially ordered on the basis of their centroid coordinates—left to right for horizontal text or top to bottom for vertical arrangements—and concatenated into continuous strings.

2.2.2 Cloud OCR and postprocessing

Character recognition is performed using Microsoft Azure Computer Vision Read API. To mitigate common OCR errors on metallic surfaces, rule-based postprocessing is applied: (i) character disambiguation corrects frequent confusions ($O \leftrightarrow 0$, $I \leftrightarrow 1$, $B \leftrightarrow 8$, $S \leftrightarrow 5$, $Z \leftrightarrow 2$) and (ii) noise filtering removes isolated single-character detections and non-alphanumeric symbols likely representing surface artifacts.

2.3 Target component and material properties

In this section, we describe the aerospace component used for system validation, detailing its functional role, material composition, and the specific optical challenges that motivated the proposed recognition approach.

2.3.1 Component specification

The test specimen is an Inconel 718 hanger bracket serving as a suspension element between the combustor liner and outer casing in aerospace gas turbine engines. This component accommodates thermal expansion while maintaining liner positioning in the high-temperature section. Laser-engraved alphanumeric identification codes on the component surface enable manufacturing traceability and quality control.

2.3.2 Material characteristics and vision challenges

Inconel 718, a nickel-based precipitation-hardened superalloy with exceptional high-temperature strength and corrosion resistance, presents multiple optical inspection challenges: (i) intense specular reflection and low inherent contrast; (ii) variable surface conditions from oxide scale and process residues; (iii) geometric distortion from component curvature; and (iv) shallow engravings with low signal-to-noise ratio. The integrated detection–OCR pipeline with targeted preprocessing addresses these material-specific challenges.

2.4 System integration and implementation

The complete inspection system integrates hardware, control logic, and vision algorithms into a production-ready platform. Herein, the system architecture, motion control strategy, timing sequence, and data management protocols are described.

To optimize processing efficiency and address the specific optical sensing challenges of Inconel 718 components, the system employs a hybrid edge–cloud architecture. The processing workload is distributed across three stages.

- (1) Edge Computing (Host PC & PAC): The local host executes the YOLOv8-n model for real-time character localization and applies targeted preprocessing (grayscale conversion and adaptive thresholding) exclusively to ROI. This edge-based processing strategy provides two critical advantages. First, it ensures that only high-quality, glare-mitigated image crops are generated, effectively addressing the low-contrast and high-reflectance characteristics typically found in nickel-based superalloy surfaces. Second, it acts as a data filter, transmitting only relevant image segments rather than full-resolution streams, thereby conserving network bandwidth and enhancing stability in the production environment.
- (2) Cloud-based Character Recognition: The preprocessed ROI crops are transmitted to the Microsoft Azure Computer Vision Read API via a secure HTTPS interface. This cloud-based approach leverages advanced deep learning models with superior generalization capabilities for reading laser-engraved characters on metallic surfaces, achieving higher accuracy than lightweight local OCR implementations when handling diverse fonts, surface textures, and character degradation patterns.
- (3) Postprocessing and System Integration: Upon receiving the OCR results from the cloud service, the local host immediately executes rule-based disambiguation algorithms (distinguishing ‘O’ from ‘0’) and spatial sequence reconstruction to consolidate fragmented

character segments into complete identification codes. The validated results are transmitted to the PAC via digital input/output (I/O) signals to trigger subsequent automation sequences and complete production data logging in the manufacturing execution system.

2.4.1 Hardware and control architecture

The integrated system (Fig. 1) was subjected to continuous operation testing over 500 inspection cycles spanning approximately 2 h. During this period, the PAC maintained consistent synchronization with the vision host through digital I/O communication, with zero control system failures or motion errors. The modular architecture successfully coordinates four motion stages (approach, inspection, departure, home return) with vision processing stages (acquisition, detection, preprocessing, OCR, logging) through digital input/output (I/O) handshaking. The standardized control interface enables straightforward migration to alternative PLC or edge computing platforms with minimal software modifications, supporting flexible deployment across diverse production lines.

The system architecture is illustrated in Fig. 2. A PAC serves as the central coordinator, interfacing with positioning sensors, controlling the linear slide motion, and managing the vision subsystem. The vision pipeline (Fig. 3) comprises four sequential stages: character region detection via YOLOv8, ROI cropping and preprocessing, cloud-based OCR recognition, and result output with data logging. The mechanical platform, shown in Fig. 4, consists of a linear guide system for precise component positioning during inspection.

2.4.2 Inspection cycle timing

The system operates with event-driven triggering at 3–5 s intervals. The host computer initiates single-frame acquisition and inference, communicating with the PAC through digital I/O signals. The total inspection cycle time T_{cycle} is decomposed into six sequential components as defined in

$$T_{cycle} = T_{move} + T_{settle} + T_{capture} + T_{infer} + T_{log} + T_{return}, \quad (8)$$

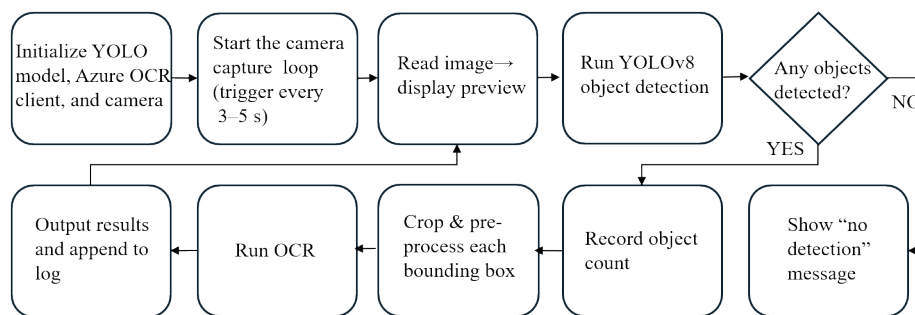


Fig. 3. Vision processing workflow for periodic single-frame detection and OCR.

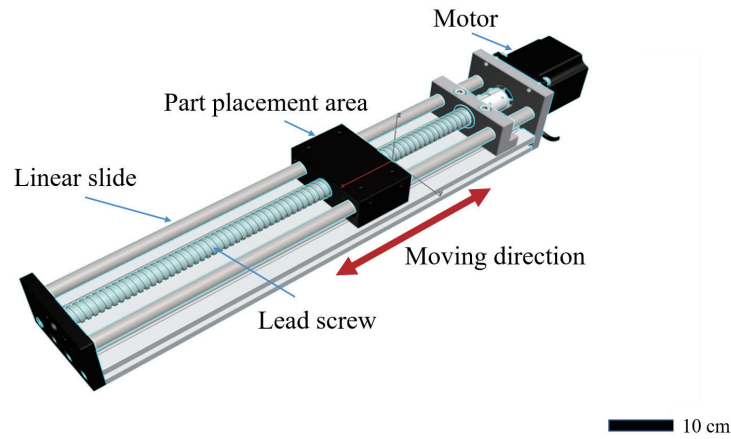


Fig. 4. (Color online) Three-dimensional schematic of the linear-slide inspection platform.

where T_{move} represents component travel to the inspection position, T_{settle} is the mechanical stabilization period, $T_{capture}$ denotes image acquisition time, T_{infer} encompasses vision processing (detection and OCR), T_{log} includes result display and data logging, and T_{return} represents the return motion to home position.

The end-to-end vision latency T_{e2e} , primarily determined by T_{infer} , was measured in approximately 0.35 s per component. The complete operational sequence proceeds as follows: approach positioning → image capture → YOLOv8 detection → ROI cropping and preprocessing → OCR recognition → result display and logging → departure → return to home position.

2.4.3 Motion control configuration

To optimize imaging stability while minimizing cycle time, the linear slide motion profile employs four sequential stages with distinct velocity and positioning parameters (Table 1). The inspection stage utilizes a reduced feed rate (1000 mm/min) to minimize vibration, while rapid traverse rates (15000–50000 mm/min) are employed for noncritical positioning.

2.4.4 Data logging and traceability

Each inspection cycle generates a timestamped record containing the image identifier, detected bounding box coordinates, recognized character strings, and object count. Data are logged in text format for quality control and process traceability. The user interface provides real-time feedback on the detection status for operator monitoring and subsequent audit trails.

3. Results

The proposed system is evaluated through controlled experiments by which YOLOv8 detection performance, OCR recognition accuracy, and integrated system efficiency are assessed. Experimental results validate the feasibility of the proposed system for industrial deployment under production-floor conditions.

Table 1
PAC motion recipe used in this study.

Stage	<i>X</i> position (mm)	Feed <i>F</i> (mm/min)	Description
Approach	1400	25000	Rapid positioning to inspection start
Inspection	1600	1000	Slow traverse with stabilization dwell
Departure	2400	15000	Accelerated exit after recognition
Home return	0	50000	Rapid return to initial position

3.1 Object detection results

3.1.1 Training and convergence behavior

The dataset was partitioned into training (70%), validation (20%), and testing (10%) subsets, with the test set intentionally including diverse lighting conditions and surface glare to assess generalization. YOLOv8-nano models were trained for 50, 100, 200, and 300 epochs to identify the optimal training duration. Figures 5(a)–5(d) present $mAP@0.5$ convergence curves for all training runs. The 300-epoch model achieved the highest validation performance ($mAP@0.5 \approx 0.93$) with stable convergence and was selected as the final detector. Compared with the 50-epoch baseline ($mAP@0.5 \approx 0.78$), $mAP@0.5$ of the 300-epoch model improved by approximately 0.15, demonstrating the benefit of extended training for small-scale, low-contrast character localization. To clarify the training dynamics presented in Fig. 5, the model performance is evaluated using $mAP50$ (mean average precision at an IoU threshold of 0.5), which is the standard metric for object detection tasks. In the graphs, the solid blue line represents the actual recorded metric value at each training epoch. In contrast, the dashed blue line represents the smoothed moving average, which helps visualize the overall trend. The convergence behavior shows rapid improvement during the initial training phase, with the performance stabilizing at around epoch 50. The flattening of both the actual and smoothed curves confirms that the model has reached convergence without significant fluctuations, validating the effectiveness of the selected 100-epoch training duration.

3.1.2 Validation/test performance

Table 2 summarizes detection performance on the test set. The final model achieves $mAP@0.5 = 0.93$, indicating high localization accuracy at the 50% IoU threshold. The $mAP@0.5:0.95$ metric (0.62) demonstrates consistent performance across stricter IoU thresholds. High precision (0.93) and recall (0.91) confirm robust detection with minimal false positives while capturing 91% of character regions.

3.2 Character recognition accuracy

OCR performance was evaluated under various imaging conditions. Under optimal conditions (stable illumination, proper focus), the system achieved approximately 96% string-

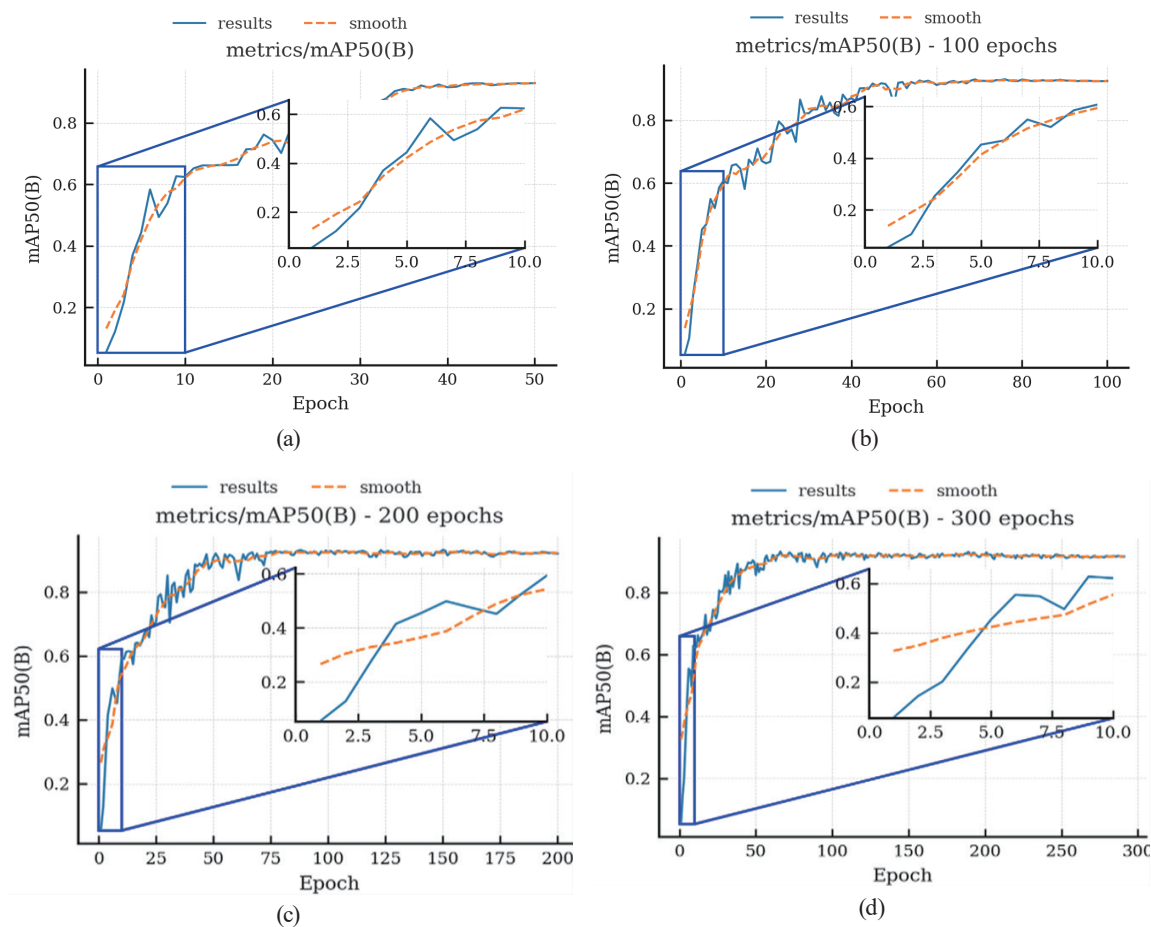


Fig. 5. (Color online) Training convergence curves (mAP@0.5 vs epoch): (a) 50 epochs; (b) 100 epochs; (c) 200 epochs; (d) 300 epochs.

Table 2
Yolov8-nano detection performance on test set (300 epochs).

Metric	Value	Interpretation
$mAP@0.5$	0.93	Localization accuracy at $IoU = 0.5$
$mAP@0.5:0.95$	0.62	Average across $IoU = 0.5-0.95$
Precision	0.93	Correct detections/total detections
Recall	0.91	Detected regions/ground truth regions

level accuracy, where complete identification codes were recognized without character errors. Under suboptimal conditions, including mild defocus, reduced illumination, or slight component misalignment, accuracy decreased to 85–90%. Representative inspection results are shown in Figs 6–8. Figure 6 presents the virtual inspection scene with a target-marking region on a 3D component model. Figure 7 shows the performance under motion blur (YOLOv8 detection bounding boxes). Figure 8 shows successful recognition under clear conditions (accurate character localization).

Recognition performance is susceptible to the angle of light incidence and surface reflectivity characteristics of the metallic substrate. Two primary optical mechanisms contribute to

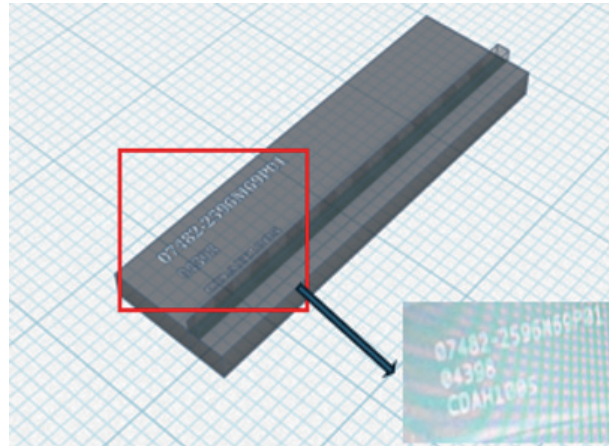


Fig. 6. (Color online) Virtual inspection scene showing target-marking region.



Fig. 7. (Color online) YOLOv8 detection results on motion-blurred image.



Fig. 8. (Color online) YOLOv8 detection results demonstrating accurate character localization.

recognition failures. (i) Stroke degradation caused by specular reflection: Because of the high-gloss nature of Inconel 718 alloy surfaces, when the incident light angle aligns closely with the camera's viewing angle, intense specular reflection creates localized surface saturation (hotspots). This optical washout effect interrupts the visual continuity of laser-engraved character strokes, rendering them undetectable by the OCR algorithm. (ii) Contrast loss due to illumination geometry: The shallow depth of laser-engraved characters (typically 50–100 μm) requires oblique illumination to generate sufficient shadow contrast for character edge detection. When illumination approaches perpendicular incidence to the surface, the resulting low

contrast-to-noise ratio makes the character texture indistinguishable from the background surface, leading to recognition failure.

Primary failure modes include (i) specular glare causing local overexposure and stroke degradation, (ii) insufficient engraving depth resulting in low contrast, and (iii) component tilt that introduces geometric distortion. System limitations arise from illumination variability under fixed camera positioning, limited training dataset diversity, and dependence on cloud OCR service availability and version consistency.

3.3 System efficiency analysis

The vision processing duty cycle quantifies computational resource utilization within each inspection cycle. With measured end-to-end vision latency $t_{e2e} \approx 0.35$ s and inspection cycle times $T_{cycle} = 3, 4, \text{ or } 5$ s, the duty cycle is calculated according to

$$\text{Duty cycle (\%)} = \frac{t_{e2e}}{T_{cycle}} \times 100\% . \quad (9)$$

For cycle times of 3, 4, and 5 s, corresponding duty cycles are 11.7, 8.8, and 7.0%, respectively. These low values indicate that vision processing is not the system bottleneck; the majority of cycle time is allocated to mechanical motion and stabilization, ensuring image quality while maintaining low computational power and thermal load.

4. Discussion

We presented an integrated pipeline—ROI-only preprocessing + YOLOv8 detection + cloud OCR + rule-based correction—for the real-time recognition of aerospace parts made of Inconel 718 under smart-factory conditions. On our prepared dataset with diverse illumination and surface states, the final model (YOLOv8-n, 300 epochs) attains $mAP@0.5 \approx 0.93$, $mAP@0.5:0.95 \approx 0.62$, $Precision \approx 0.93$, and $Recall \approx 0.91$ on the test set; the OCR string accuracy is $\approx 96\%$ under stable lighting and $\approx 85\text{--}90\%$ with varying lighting. The end-to-end latency is $t_{e2e} \approx 0.35$ s/part; relative to a trigger period of $T_{cycle} = 3\text{--}5$ s, the duty cycle is $11.7\% / 8.8\% / 7.0\%$, indicating that the vision stage is not the bottleneck. Compared with the 50-epoch baseline ($mAP@0.5 \approx 0.78$), the 300-epoch model improves $mAP@0.5$ by ≈ 0.15 , confirming the benefit of longer training for small characters and low-contrast engravings. The developed system enables the automated inspection of specialty aerospace parts and shows potential for deployment in real manufacturing lines to reduce manual effort and idle time, improve line efficiency, and support sustainable production.

5. Conclusions

In this study, we developed an automated vision system for character recognition on Inconel 718 aerospace components, by integrating YOLOv8 detection, cloud OCR, and rule-based correction to address material-specific optical challenges. Key contributions include the following.

1. The complete system achieves 93% detection precision and 91% recall, with 96% OCR accuracy under stable conditions and 85–90% under typical industrial lighting. The processing time of 0.35 s per component represents only 7–12% duty cycle, confirming that vision processing is not a bottleneck.
2. Extended training (300 epochs) improves detection accuracy by 19% over the baseline, demonstrating the benefit of adequate training for low-contrast metallic surface features.
3. ROI-focused preprocessing and character disambiguation rules effectively mitigate specular reflection and geometric distortion challenges on nickel-based superalloy surfaces.
4. PAC-integrated architecture with timestamped logging enables quality control, process traceability, and continuous improvement for industrial deployment.

Future improvements should include multiview imaging, edge-based OCR, and expanded training datasets to enhance system versatility and reduce external dependences.

Acknowledgments

The authors gratefully acknowledge financial support from the National Science and Technology Council, Taiwan (Grant No. NSTC 113-2222-E-992-007-). The authors thank the Aerospace Industrial Development Corporation (AIDC) for technical consultation and the National Kaohsiung University of Science and Technology (NKUST) for technical assistance and administrative support. The authors would like to express their gratitude to all members of the Mechatronics and Controls Laboratory at NKUST for their assistance.

References

1. P. Jiang, D. Ergu, F. Liu, Y. Cai, B. Ma: *Procedia Comput. Sci.* **199** (2022) 1066. <https://doi.org/10.1016/j.procs.2022.01.135>
2. M. Hussain: *IEEE Access*. **12** (2024) 42816. <https://doi.org/10.1109/ACCESS.2024.3378568>
3. Z. Wang, X. Li, Y. Liu, and J. Zhang: *Sensors* **24** (2024) 6268. <https://doi.org/10.3390/s24196268>
4. J. Gui, X. Lei, W. Wang, and H. Li: *Electronics* **13** (2024) 3129. <https://doi.org/10.3390/electronics13163129>
5. B.-L. Jian, J.-P. Hung, C.-C. Wang, and C.-C. Liu: *Sens. Mater.* **32** (2020) 4217. <https://doi.org/10.18494/SAM.2020.3101>
6. J. Hu, D. Yu, S. Li, and Z. Zhang: *Sensors* **24** (2024) 4088. <https://doi.org/10.3390/s24134088>
7. Z. Ying, H. Wang, and Q. Liu: *Electronics* **13** (2024) 2435. <https://doi.org/10.3390/electronics13132435>
8. H. Peng, J. Bayon, J. Recas, and M. Guijarro: *Algorithms* **18** (2025) 286. <https://doi.org/10.3390/a18050286>
9. Q. Tang, Y. Lee, and H. Jung: *Sustainability* **16** (2024) 2161. <https://doi.org/10.3390/su16052161>
10. H.-C. Lin, C.-C. Tsai, C.-H. Chang, and M.-L. Chen: *Sustainability* **15** (2023) 12915. <https://doi.org/10.3390/su151712915>
11. M. Al-Hasan, M. F. Al-Zubaidi, and S. A. Hameed: *Technologies* **12** (2024) 164. <https://doi.org/10.3390/technologies12090164>
12. M. Rehman, M. H. Malik, M. U. Ghani, and A. Chughtai: *Appl. Sci.* **15** (2025) 1257. <https://doi.org/10.3390/app15031257>

- 13 J. Sadek, A. Vlachidis, V. Pickering, M. Humbel, D. Metilli, M. Carine, and J. Nyhan: *Int. J. Digit. Humanit.* **6** (2024) 237. <https://doi.org/10.1007/s42803-024-00091-4>
- 14 Y. Liu, X. Wang, X. Chen, and Z. Zhang: *IEEE Sens. J.* **24** (2024) 4863. <https://doi.org/10.1109/JSEN.2023.3346470>
- 15 Y. Zuo, J. Wang, L. Sun, and H. Li: *IEEE Sens. J.* **24** (2024) 9197. <https://doi.org/10.1109/JSEN.2024.3359218>
- 16 C. Li, Z. Zhao, and F. Yang: *Sensors* **24** (2024) 5906. <https://doi.org/10.3390/s24185906>
- 17 D. Matuszczyk and F. Weichert: *Sensors* **23** (2023) 1619. <https://doi.org/10.3390/s23031619>