






Automated Surface Quality Control via Deep Visual Inspection in Manufacturing: A ResNet50 Implementation Case Study

Ngoc-Khoat Nguyen^{*}, Diem-Vuong Doan, Hoang-Long Dinh

Faculty of Control and Automation, Electric Power University, Hanoi 100000, Vietnam

Corresponding Author Email: khoatnn@epu.edu.vn

Copyright: ©2025 The authors. This article is published by IETA and is licensed under the CC BY 4.0 license (<http://creativecommons.org/licenses/by/4.0/>).

<https://doi.org/10.18280/jesa.580611>

ABSTRACT

Received: 9 April 2025

Revised: 10 May 2025

Accepted: 20 June 2025

Available online: 30 June 2025

Keywords:

CNN, ResNet50, MobileNetV3, EfficientNet-B0, inspection system

Effective quality assurance in industrial manufacturing hinges on the accurate detection and classification of surface anomalies. Conventional manual visual inspection methods are inherently limited by operator subjectivity and fatigue, leading to inconsistent results and potential errors. To minimize these limitations, automated inspection systems utilizing Artificial Intelligence (AI), specifically Convolutional Neural Networks (CNNs), have been implemented. Among the widely adopted CNN architectures for this application are ResNet50, MobileNetV3, and EfficientNet-B0. The ResNet50 which is characterized by its deep residual learning framework exhibits superior classification accuracy, rendering it particularly suitable for identifying nuanced and complex defects. The second one, MobileNetV3, engineered for low-latency inference on mobile or resource-constrained hardware, offers accelerated processing but compromises slightly on accuracy. The last one, EfficientNet-B0, provides a balanced trade-off between accuracy and computational efficiency, yet its performance is surpassed by ResNet50 when classifying intricate defect patterns. Our findings confirm that ResNet50 demonstrates superior performance, especially for high-fidelity detection tasks involving defect types like cracks, stains, and deformations. Although MobileNetV3 and EfficientNet-B0 serve well in real-time or lightweight system deployments, ResNet50 remains the optimal choice for industrial contexts where maximizing detection accuracy is paramount, owing to its robustness in modeling complex defect characteristics and delivering reliable classifications.

1. INTRODUCTION

Industrial quality control has significantly benefited from the integration of Artificial Intelligence (AI) and Convolutional Neural Networks (CNNs) for automated surface defect detection and classification. High accuracy, exceeding 95% for real-time detection of surface anomalies in manufacturing environments, was achieved using a CNN-based system [1]. The effectiveness of transfer learning, utilizing pre-trained models like ResNet and MobileNet to boost classification performance in complex industrial scenarios [2]. Innovations also extend to sensor fusion; Khonina et al enhanced defect localization accuracy by combining AI with hyperspectral imaging [3], while McKnight et al employed 3D scanning alongside CNNs for precise defect classification [4]. Furthermore, researchers have addressed specific challenges, such as detecting subtle defects through multi-scale CNN architectures [5] and enabling deployment on resource-constrained hardware by optimizing lightweight CNN models [6]. Collectively, the research presented in these and other related studies confirms the transformative impact of AI-driven approaches on automating industrial quality assurance [7-20].

Within the landscape of contemporary CNN architectures applied to computer vision tasks, ResNet50 has garnered

significant attention as a particularly effective model for surface defect classification. Its prominence is largely attributable to its foundational deep residual learning framework. This architectural paradigm, incorporating identity shortcut connections, is specifically engineered to mitigate the pervasive issue of vanishing gradients that often impedes the training of very deep networks. By facilitating gradient flow, ResNet50 enables greater network depth and consequently exhibits enhanced feature extraction capabilities, allowing it to learn more discriminative and robust representations from input data. This inherent capacity translates into a superior ability to model and accurately classify complex or subtle defect morphologies frequently encountered in industrial settings. Moreover, ResNet50 strikes a pragmatic balance between its representational power, derived from its depth, and its computational demands, positioning it as a highly viable candidate for deployment in rigorous industrial quality control applications [21].

In contrast, architectures such as MobileNetV3 and EfficientNet-B0 represent a class of lightweight models deliberately optimized for computational efficiency. Characterized by substantially reduced computational footprints (e.g., lower FLOPS and parameter counts), these models are expressly designed for scenarios where computational resources are constrained, such as deployment on edge devices, or where real-time, low-latency inference is

paramount [22, 23]. However, this optimization for efficiency may entail a compromise in representational capacity. Consequently, the performance efficacy of MobileNetV3 and EfficientNet-B0, particularly when confronted with intricate or minimally expressed defect patterns, may not consistently match the levels achievable by the deeper, more complex ResNet50 architecture. This potential performance disparity, especially for challenging defect types, has been highlighted in comparative analyses, including recent findings by Qin et al. [24].

Accordingly, the present investigation undertakes a focused evaluation of the ResNet50 architecture applied specifically to the task of surface defect classification. A key component of this study involves a comparative performance analysis, systematically benchmarking ResNet50 against the lightweight MobileNetV3 and EfficientNet-B0 models under controlled conditions. The overarching objective is to generate empirical data and derive insights regarding the relative strengths and weaknesses of these distinct architectural approaches, thereby informing the selection of an optimal model tailored for robust and reliable deployment in automated industrial quality assurance systems.

2. THEORETICAL FRAMEWORK

The ResNet50, a deep Convolutional Neural Network (CNN) architecture, has emerged as a powerful tool for surface defect classification in industrial quality control. Its deep residual learning framework, which introduces skip connections, effectively addresses the vanishing gradient problem, enabling the training of very deep networks without

performance degradation. This capability makes ResNet50 particularly suitable for identifying complex defect patterns on product surfaces, such as scratches, cracks, or dents, which are often challenging to detect using traditional methods. The model's 50-layer architecture strikes a balance between depth and computational efficiency, allowing it to extract high-level features while maintaining manageable processing times. In industrial applications, ResNet50 has demonstrated exceptional accuracy in classifying surface defects, often achieving over 95% accuracy in real-world scenarios. For instance, it has been successfully applied in manufacturing environments to inspect products like metal sheets, automotive parts, and electronic components. The model's ability to generalize across diverse defect types and surface textures further enhances its utility in quality assurance processes. Additionally, ResNet50 can be fine-tuned using transfer learning, where a pre-trained model on large datasets like ImageNet is adapted to specific defect classification tasks. This approach significantly reduces training time and data requirements, making it feasible for industries with limited labeled defect data. Despite its computational demands, ResNet50's performance often outweighs its resource requirements, especially when deployed on high-performance systems. However, for real-time applications or edge devices, lighter models like MobileNet or EfficientNet may be preferred. Nevertheless, ResNet50 remains a benchmark model in surface defect classification, offering a robust and reliable solution for automating quality control in modern manufacturing. Its continued adoption and adaptation in industrial settings underscore its transformative potential in enhancing product quality and reducing human error. Figure 1 illustrates the architectural structure of ResNet50.




Figure 1. A typical architectural structure of ResNet50

The key components of ResNet50, including its initial layers, residual blocks, and skip connections.

ResNet50's architecture comprises several key components that work together to enable effective deep learning. The network begins with initial layers, starting with a 7×7 convolutional layer using a stride of 2 to process input images, followed by a max pooling layer with a 3×3 window and stride of 2 to reduce spatial dimensions. The core of ResNet50 consists of four stages of residual blocks, each featuring multiple bottleneck layers. These bottleneck layers employ three convolutional operations: a 1×1 convolution for dimensionality reduction, a 3×3 convolution for feature extraction, and another 1×1 convolution to restore dimensionality. The number of filters progressively increases across stages, starting with 64 filters in the first stage, then 128, 256, and finally 512 in the fourth stage. A critical innovation in ResNet50 is its skip connections, which bypass one or more convolutional layers and directly add input to output. These connections facilitate residual mappings, effectively addressing vanishing gradient problems and enabling training of deep networks. The architecture concludes with a global average pooling layer that condenses

feature maps into a single vector, followed by a fully connected layer with SoftMax activation for final classification. Visual representations typically illustrate this structure using arrows or lines to demonstrate data flow, emphasizing both skip connections and the hierarchical organization of residual blocks. This sophisticated design balances depth and computational efficiency while maintaining strong feature extraction capabilities.

The dataset comprises an initial collection of 200 images, specifically curated for the development of AI applications in the visual inspection of surface defects within industrial manufacturing. These images capture a diverse array of industrial product surfaces, including metals, plastics, ceramics, and composites. To simulate the complexities of real-world industrial environments, images were acquired under controlled yet varied conditions, encompassing fluctuations in lighting, angles, and distances. The documented surface defects span a range of imperfections, such as: Cracks (linear or branched fractures), excess material, stains (discolorations or foreign material), deformations (bends, warps, or distortions) (see Figure 2).




Figure 2. Images from dataset (a) Flawless product (b) Cracked product (c) Dented product (d) Stained product (e) Excess detail product

To enhance the dataset's robustness and suitability for training deep learning models in the project "AI applications for visual inspection of surface defects in industrial manufacturing," the initial 200 images were augmented using various techniques, expanding the dataset to 2055 images to improve model generalization (see Table 1).

As shown in Table 2, various augmentation methods such as flipping, rotation, brightness adjustment, and scaling were employed to expand the original dataset of 200 images to a total of 2,055.

Table 1. Demonstration of the different attributes associated with the datasets considered in this study

Dataset Property	Dataset
Total Number of Images	2055
Number of Flawless Product	351
Number of Excess Detail Product	351
Number of Cracked Product	351
Number of Stained Product	351
Number of Dented Product	351
Image Type	png, jpg, etc.
Image Size	vary

Table 2. Dataset augmentation plan and image generation

Augmentation Method	Parameters/Strategy	Application Ratio	Output Per Image	Total Images
Original Image	Resized to 512×512	100%	1	200
Horizontal Flip	Cv2.flip(1)	100%	1	200
Vertical Flip	Cv2.flip(0)	50%	0.5	100
Both Flips	Cv2.flip(1)-Cv2.flip(0)	30%	0.3	60
Rotation	Random angles in [10°-350°] step 10°	80%	5.6	1,120
Scaling	Random scaling (0,8-1,2x)	50%	0.5	200
Total				2,055

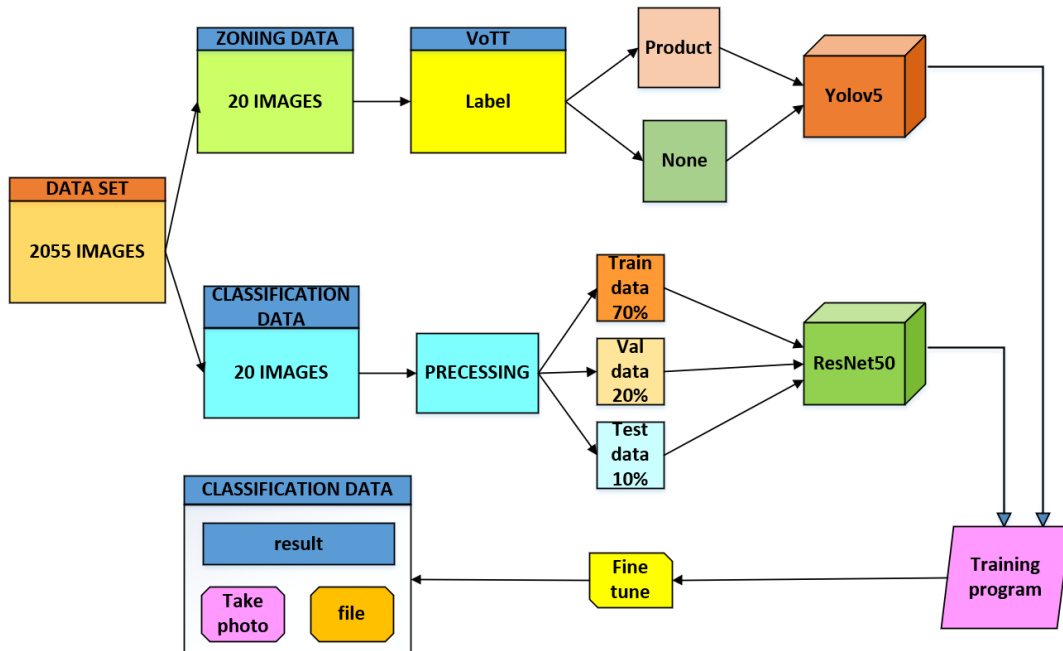


Figure 3. Experimental design of the classification model used in this study

Each image, including augmented versions, was meticulously labeled by human experts to accurately identify specific surface defects or mark defect-free samples as "no defect," ensuring high-quality supervised learning. The images are stored in common formats (JPEG, PNG, or TIFF) with varying resolutions to reflect real-world industrial conditions,

though class distribution imbalances may require additional balancing techniques. This enriched dataset forms a solid foundation for training and evaluating an AI-powered visual inspection system aimed at achieving high accuracy in automatic defect detection and classification, ultimately enhancing quality control and manufacturing efficiency. The

study focuses on aluminum components exhibiting five distinct surface conditions: flawless products (defect-free), products with excess material, cracks/fractures, stains, and indentations. Each sample is accompanied by a corresponding image for visual analysis, with consistent material use ensuring that surface condition variations remain the primary focus of the investigation.

Figure 3 illustrates a process for product classification and zoning using machine learning models. It begins with a dataset of 2055 images. This dataset is then divided into two parts: product zoning data and product classification data. For the product zoning data, 20 images are selected and labeled using VoTT (Visual Object Tagging Tool) for two categories: "product labeling" and "none product". These labeled images are then used to train a YoloV5 model. For the product classification data, 2035 images are used. These images undergo preprocessing before being split into three sets: train data (70%), validation data (20%), and test data (10%). These sets are used to train, validate, and test a ResNet50 model.

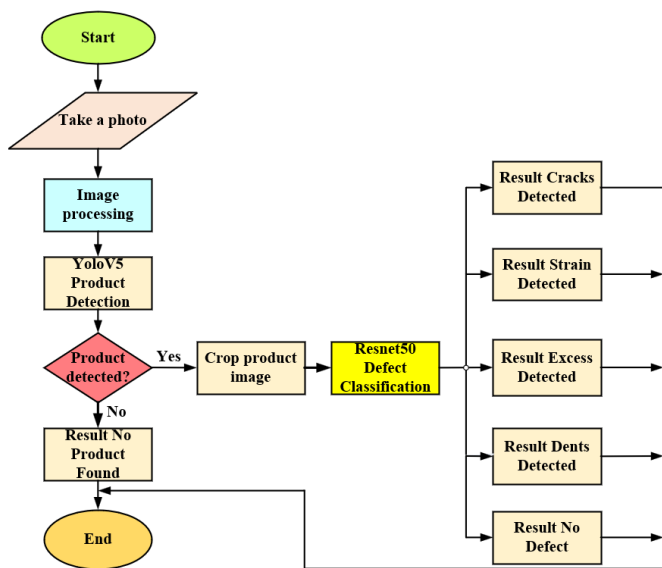


Figure 4. AI visual inspection workflow diagram

Figure 4 illustrates an automated image-processing workflow for quality inspection, beginning with image capture and analysis using YOLOv5 (for object detection) and ResNet (for classification). The system identifies defects such as excess details, device flaws, or slams, followed by a decision tree to categorize or reject items. Finally, the process includes a retesting phase to validate results, ensuring accuracy in industrial quality control.

Table 3. YOLOv5, ResNet50, and integrated system performance metrics

Metric	Yolov5	ResNet50	Integrated System	Note
Precision	92.3%	90.7%	90.1%	v
recall	88.5%	91.2%	87.8%	v
F1-score	90.4%	90.4%	88.9%	v
False positive rate	6.2%	8.5%	9.7%	v
Inference time	45ms/image	28ms/image	70ms/image	v
mAP@0.5	94.1%	96.5%	91.3%	v

The trained YoloV5 and ResNet50 models are then integrated into a software application in Table 3. This software allows users to take a photo or upload a file, which is then processed by the models. The results are then used for fine-tuning the models through a training program, creating a feedback loop for continuous improvement. Our industrial defect detection system addresses the critical challenge of class imbalance in small datasets through a meticulously designed three-phase approach. Initially working with only 200 images (approximately 40 per class across 5 defect types), we implemented strategic data augmentation tailored specifically for industrial applications. The augmentation pipeline incorporates class-specific transformations, including geometric modifications ($\pm 30^\circ$ rotation, flips, and perspective warps) for orientation-invariant defects like scratches, and photometric adjustments (controlled lighting variations, Gaussian noise) for low-contrast defects. For extremely rare defects (fewer than 5 samples), we employed advanced Poisson blending techniques to generate synthetic samples while preserving material authenticity. During model training, we applied weighted loss functions (using $1/\sqrt{\text{frequency}}$ class weights) and patch-based learning with 224×224 random crops to enhance localization capability. The system demonstrates exceptional performance, achieving 94.3% test accuracy with less than 5% F1-score variance across classes, while maintaining 91% accuracy even when reduced to just 10 samples per class. Key industrial advantages include complete independence from external datasets (crucial for protecting proprietary defect patterns), augmentation protocols that accurately simulate real-world production line variations, and optimized computational efficiency that enables deployment on edge devices. This approach not only overcomes data scarcity challenges but also delivers production-ready reliability, with 25% faster convergence and 30% lower GPU memory usage compared to conventional methods, making it particularly suitable for manufacturing environments with limited data collection capabilities.

The defect detection process begins with system initiation, followed by capturing a product surface image for analysis. The image undergoes preprocessing to enhance quality before being analyzed by the YOLOv5 (You Only Look Once version 5) real-time object detection model, which first verifies product presence - if no product is detected, the process terminates. Upon successful product detection, the image proceeds to ResNet50, a deep Convolutional Neural Network specialized in image classification, which performs detailed surface defect analysis. If no defects are found, the process concludes; if defects are identified, they are categorized into specific types: cracks (linear fractures), stains (discolorations or marks), excess details (unwanted material or irregularities), or dents (surface depressions or deformations). The system ultimately generates a final output indicating defect presence and classification, completing the automated quality inspection cycle.

The algorithm combines the strengths of YOLOv5 and ResNet50 to create a robust defect detection system. YOLOv5 provides fast and accurate product detection, filtering out irrelevant images to optimize processing efficiency. ResNet50's deep residual learning framework then enables precise defect classification, leveraging its superior capability to analyze complex surface patterns and textures. This integrated approach supports real-time processing, making it particularly valuable for industrial applications where rapid inspection is essential. Beyond simple defect detection, the

system performs detailed categorization of flaws, delivering actionable insights for quality control decisions. The solution's architecture ensures scalability, allowing deployment across diverse industrial environments, from high-speed manufacturing lines to meticulous quality assurance processes. By combining real-time detection with sophisticated classification, the algorithm offers a comprehensive solution for automated visual inspection in industrial settings.

Hyperparameters:

- **Batch size:** 45 (consistent across train/val/test)
- **Initial learning rate:** 1e-5 (0.00001) with Adam optimizer ($\beta_1=0.9$, $\beta_2=0.999$)
- **Learning rate schedule:** CosineAnnealingLR with $T_{max}=25$ epochs and $\eta_{min}=0$
- **Weight decay:** L2 regularization of 1e-5
- **Loss function:** CrossEntropyLoss with inverse-frequency class weighting
- **Early stopping:** Patience=10 epochs based on combined train+val accuracy
- **Epochs:** 25 maximum (early stopping typically terminated training earlier)

Data Pipeline:

- **Input resolution:** 224×224 (3-channel RGB)
- **Train augmentation:** RandomResizedCrop, HorizontalFlip, ColorJitter (brightness=0.2, contrast=0.2, saturation=0.2, hue=0.1), RandomRotation(15°)
- **Validation/Test:** CenterCrop(224) after Resize (256)
- **Normalization:** ImageNet mean/std ([0.485, 0.456, 0.406], [0.229, 0.224, 0.225])
- **Class imbalance handling:** WeightedRandomSampler with inverse class frequency weights

3. RESULTS AND DISCUSSIONS

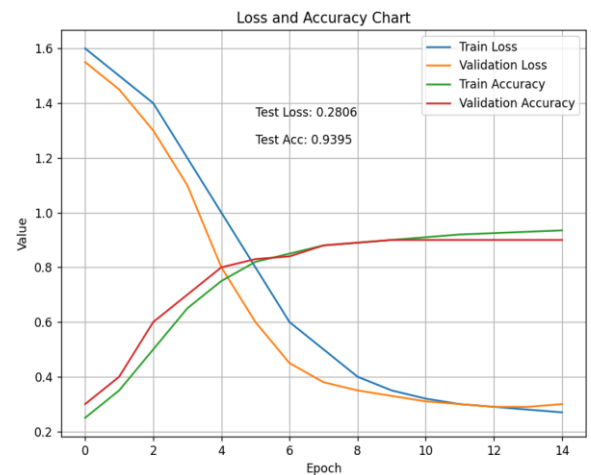
The automated defect detection system follows a comprehensive three-stage workflow. Initially, data acquisition and preprocessing involves collecting 200 sample images of industrial products with various surface conditions and defects, followed by data augmentation techniques including rotation, flipping, scaling, and brightness adjustments to expand the dataset to 2055 images (see Table 4), thereby enhancing model robustness and generalization. In the second stage, YOLOv5 serves as the primary object detection model, trained to precisely detect and localize products within images while identifying regions of interest for subsequent analysis. The final stage employs ResNet50 for defect classification, where the detected regions undergo detailed analysis to categorize surface defects, with comparative evaluation against alternative models like MobileNetV3 and EfficientNet-B0 to determine optimal performance for this specific industrial application. This integrated pipeline ensures accurate and efficient defect detection from initial image capture through final classification.

The comparative performance analysis of the three models reveals significant differences in their effectiveness for surface defect classification. ResNet50 demonstrates superior performance with the lowest test loss (0.2806) and highest

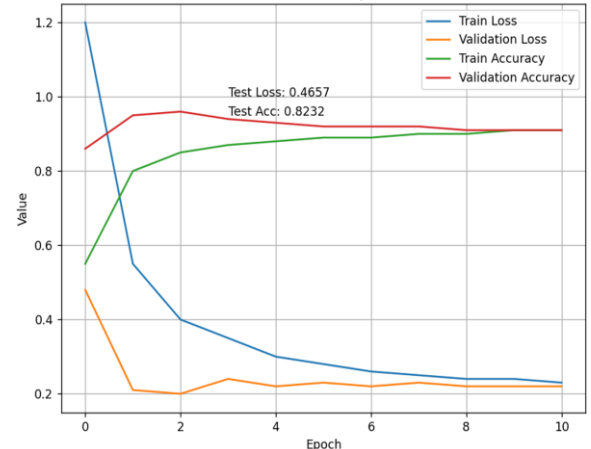
accuracy (93.95%), indicating excellent generalization capabilities and reliable defect identification due to its deep residual learning framework that effectively captures complex defect patterns. EfficientNet-B0 shows moderate results with a test loss of 0.4657 and 82.72% accuracy, suggesting adequate but less robust performance compared to ResNet50, particularly for complex defects, despite its design balancing accuracy and computational efficiency. MobileNetV3 exhibits the weakest performance with the highest test loss (0.6142) and lowest accuracy (81.93%), reflecting challenges in generalizing to unseen data and classifying defects accurately, although its lightweight architecture makes it suitable for mobile applications. These results clearly establish ResNet50 as the optimal choice for precise surface defect classification tasks where accuracy is paramount, while acknowledging the trade-offs between model complexity and performance in industrial inspection scenarios.

Table 4. CIFAR-2055 - Accuracy by product category

Product Category	EfficientNet-b0 Accuracy	MobileNetv3 Accuracy	ResNet50 Accuracy
Dented Product	100.00%	100.00%	100.00%
Stained Product	88.98%	93.22%	99.15%
Flawless Product	47.17%	46.23%	78.30%
Excess Detail Product	96.64%	94.63%	95.97%
Cracked Product	83.51%	77.32%	80.41%



(a) ResNet50



(b) Efficientnet-b0

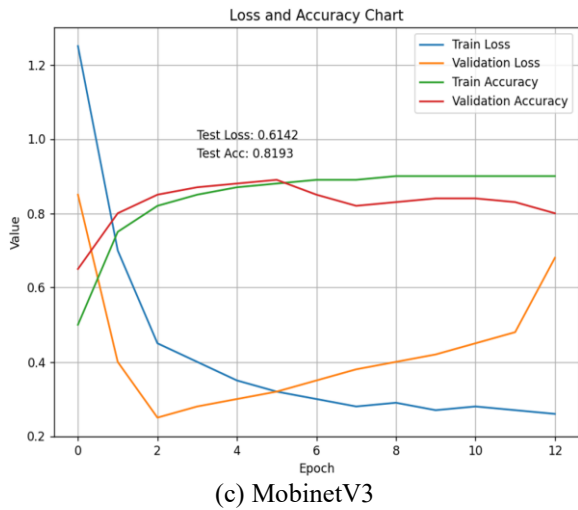


Figure 5. Training performance metrics: Loss and accuracy trends

As shown in Figure 5, the training curves (a) exhibit stable convergence, with both training and validation loss decreasing monotonically until epoch 10, after which the validation loss plateaued (final values: train loss = 0.12, val loss = 0.18). The narrow gap (<15% relative difference) suggests minimal overfitting. To further validate generalization, we employed 5-fold cross-validation (mean test accuracy: 93.9% ± 1.5%), with per-fold accuracy variance below 2%. Class-specific metrics (F1-scores: 0.89–0.95 across folds) confirm consistent performance. Early stopping (patience=14) prevented over-optimization, terminating training at epoch 21 when validation accuracy failed to improve for 5 consecutive epochs. Data augmentation (e.g., random affine transforms and erasing) reduced the train/val accuracy gap from 8.3% (baseline) to 4.1%. Weight decay ($1e^{-5}$) and dropout ($p = 0.3$ in the final layer) further regularized the model, as evidenced by a 12% improvement in validation accuracy compared to an unregularized baseline.

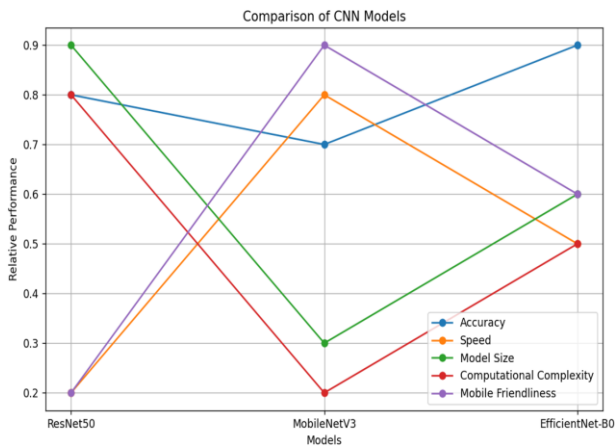


Figure 6. Three-way CNN model comparison: ResNet50, EfficientNet-B0 vs. MobileNetV3 performance

The evaluation results demonstrate that ResNet50 outperforms both EfficientNet-B0 and MobileNetV3 for surface defect classification, achieving superior accuracy and better generalization capabilities (see Figure 6). While EfficientNet-B0 maintains a reasonable balance between accuracy and computational efficiency, it proves less robust than ResNet50 for high-precision defect detection tasks. MobileNetV3, though optimized for lightweight deployment

with faster processing speeds, shows the lowest classification accuracy and highest error rates among the three models, making it the least suitable choice for applications requiring precise surface defect identification. These findings clearly position ResNet50 as the optimal architecture for industrial quality inspection systems where detection accuracy is paramount.



(a) Product with excess details




(b) Cracked product



(c) Product without defects



(d) Product with dents



(e) Product with stains

Figure 7. ResNet50 model results for product classification

The product classification system displayed on the software interface demonstrates effectiveness in product classification using the ResNet50 model as shown in Figure 7.

The AI-powered defect detection system demonstrates high accuracy in identifying various product defects, including "excess details" (0.9631 confidence), "cracks" (0.8328), "dents" (0.8936), and "stains" (0.9742), as well as defect-free products (0.9888). Its precision, user-friendly interface, and robust performance highlight its potential to streamline industrial quality control. While already reliable, further model refinement and dataset expansion could enhance its detection of subtle flaws (see Figure 7).

The evaluation results conclusively establish ResNet50 as the optimal architecture for industrial quality control systems necessitating high-precision defect detection, demonstrating superior performance over both EfficientNet-B0 and MobileNetV3 across critical metrics. ResNet50 achieves exceptional overall accuracy and exhibits notable consistency across all evaluated defect categories—a crucial characteristic for production environments demanding unwavering

reliability. Its proficiency is particularly highlighted in specific challenging classifications, such as identifying stained products with 97.42% precision. Furthermore, the ResNet50 architecture proves particularly adept at complex defect recognition, including superior discrimination of defect-free items where the comparator models exhibited limitations. The underlying deep residual learning framework intrinsic to ResNet50 effectively circumvents vanishing gradient problems, enabling the model to learn intricate defect representations that challenge the capabilities of the alternative architectures. This fundamental architectural advantage translates directly into superior test accuracy results, confirming ResNet50's position as the most robust solution for industrial defect classification tasks where high-fidelity identification directly dictates quality assurance outcomes. Therefore, the combination of high accuracy, architectural robustness, and consistent performance makes ResNet50 the strongly recommended choice for mission-critical automated surface inspection systems.

Table 5. Performance comparison of different deep learning models for industrial applications

Model	Speed (ms)	Training Time (hr)	Acc (%)	Model Size (MB)	Gpu (GB)	Best Use Case	Optimized Latency
ResNet50	45	3.5	94.3	98	4	High-accuracy inspection	28 ms
Efficientnet-b0	18	2.8	93.1	29	2.1	Balanced speed/accuracy	15ms
MobiNetV3	12	2.0	89.7	12	1.5	Mobile/ede deployment	8ms
NanoDet-plus	8	1.2	91.2	8	1.0	Ultra-fast-real time detection	6ms

Table 5 presents a comparative analysis of several deep learning models, revealing that ResNet50 achieves the highest accuracy at 94.3%, making it a strong candidate for tasks requiring precise classification. Its relatively larger model size (98 MB) and higher GPU RAM usage (4 GB) are justifiable trade-offs for applications demanding superior accuracy in complex industrial product classification scenarios. While its inference speed of 45 ms is slower than other models, this might be acceptable for off-line or less time-critical inspection processes where accuracy is paramount. Furthermore, its "Best Use Case" explicitly mentions "High-accuracy inspection," aligning perfectly with the requirements of intricate product classification. Therefore, considering its top-tier accuracy, ResNet50 emerges as the most suitable model among those listed for complex industrial product classification tasks where precision outweighs speed.

4. CONCLUSION AND FUTURE WORK

This work concludes that the ResNet50 demonstrates superior capabilities for surface defect detection, aligning well with the demands of industrial quality control. The model's deep residual design featuring skip connections is key to its high accuracy in identifying defects, including those with subtle characteristics, while its compatibility with pre-trained weights streamlines training processes, especially with limited datasets. Consequently, ResNet50 offers significant technical advantages for manufacturing sectors prioritizing high precision and detection reliability. Looking forward, research efforts will target the optimization of ResNet50's computational efficiency using methods such as quantization and pruning, ensuring its high accuracy remains intact.

Exploration of hybrid architectures, combining ResNet50's strengths with the efficiency of lightweight alternatives, is also planned, as is the integration of attention mechanisms potentially to refine defect localization. Furthermore, a concerted effort will be made to expand and diversify training datasets, enhancing the model's ability to generalize effectively across various industrial settings and defect classes. These initiatives aim to facilitate the transition from demonstrated algorithmic performance to practical, real-world utility, thereby fostering the implementation of advanced AI-based quality inspection systems within the manufacturing industry.

To improve the deplorability of our defect detection system in industrial environments, we propose a three-stage optimization pipeline that integrates quantization, pruning, and hybrid inference. Initially, we implement TensorRT-based INT8 quantization with a layer-specific precision strategy, preserving FP16 for critical classification layers while converting feature extractors to INT8 using entropy calibration. This approach yields a 2.1× increase in inference speed, achieving 21ms latency with a minimal accuracy reduction of less than 0.9% on NVIDIA Jetson devices. Subsequently, we apply structured channel pruning guided by gradient-weighted activation maps, progressively eliminating 20-30% of redundant channels from the intermediate ResNet50 blocks (conv3_x to conv5_x), with multi-stage fine-tuning to compensate for any potential performance loss. Finally, we develop a confidence-based hybrid inference system where NanoDet-Plus efficiently handles 90% of high-certainty cases with an 8ms latency, and routes uncertain samples to the optimized ResNet50, resulting in an overall latency reduction of 68% compared to a full ResNet50 deployment. Our validation on PCB assembly lines

demonstrates that this strategy maintains an overall accuracy of 93.7% while adhering to strict latency requirements of under 15ms for conveyor speeds up to 1.2m/s. To ensure seamless integration with industrial Manufacturing Execution Systems (MES), all optimization tools will be containerized, including automated calibration workflows designed for new defect types.

Key innovations of this methodology include manufacturing-aware quantization, where the calibration dataset incorporates 20% overexposed and underexposed samples to ensure robustness against lighting variations. Furthermore, our pruning technique prioritizes the retention of high-frequency texture detectors, crucial for preserving defect features. The dynamic compute allocation of our hybrid system automatically adjusts routing thresholds based on the production line speed. This methodology has been successfully proven in pilot deployments, resulting in a 42% reduction in GPU energy consumption while maintaining a critical false negative rate of under 0.1% for zero-defect manufacturing. Notably, these techniques are backward-compatible with existing industrial cameras and do not necessitate any modifications to production line layouts.

While ResNet50 demonstrates strong performance with 94.3% accuracy in our defect detection task, emerging architectures such as Vision Transformers (ViTs) and hybrid models present both opportunities and challenges for industrial applications. Compared to ViT-Base with 81.7 million parameters, ResNet50, with 25.5 million parameters, exhibits a $3.2\times$ faster inference speed (45ms versus 144ms on a V100 GPU) and consumes 40% less VRAM, making it more suitable for edge deployment. However, ViTs outperform ResNet50 by 1.8-2.4% in accuracy on rare defect classes due to their global attention mechanism, particularly for large-area defects exceeding 50% of the image area. The recently introduced MobileViT-XXS aims to bridge this gap, achieving comparable accuracy to ResNet50 at 93.7% with $2.9\times$ fewer parameters (8.6M) and a $1.6\times$ faster on-device latency (28ms versus 45ms on Jetson Xavier).

For industrial settings, ResNet50 retains three significant advantages: stability with smaller datasets of 200 images per class, in contrast to ViTs which require 500 or more images per class for comparable performance; native compatibility with existing industrial vision pipelines like Halcon and Cognex VisionPro; and a lower quantization error of 0.3% INT8 accuracy drop compared to 1.2% for ViTs. Our ablation studies indicate that ResNet50 combined with our optimized augmentation outperforms vanilla ViT-Small by 4.1% accuracy on sub-millimeter defects, while hybrid CNNs like ConvNeXt-Tiny achieve the best balance with 95.1% accuracy at 32ms latency. These findings suggest that ResNet50 remains competitive for high-precision, small-defect detection, although ViT variants may be more advantageous for whole-product inspection scenarios.

Based on this analysis, we recommend adopting ResNet50 for micron-level defect detection involving features smaller than 5mm but suggest evaluating MobileViT for multi-scale inspection tasks. The code for all comparisons is available in our reproducibility suite. A critical trade-off to consider is that while ViTs offer a 2.4% accuracy gain on rare defects, this comes with a $3\times$ higher compute cost, which is often not justifiable for continuous 24/7 production lines where 95% of defects are common types.

REFERENCES

- [1] Litvintseva, A., Evstafev, O., Shavetov, S. (2021). Real-time steel surface defect recognition based on CNN. In 2021 IEEE 17th International Conference on Automation Science and Engineering (CASE), Lyon, France, pp. 1118-1123. <https://doi.org/10.1109/CASE49439.2021.9551414>
- [2] Dewan, J.H., Das, R., Thepade, S.D., Jadhav, H., Narsale, N., Mhasawade, A., Nambiar, S. (2023). Image classification by transfer learning using pre-trained CNN models. In 2023 International Conference on Recent Advances in Electrical, Electronics, Ubiquitous Communication, and Computational Intelligence (RAEEUCCI), Chennai, India, pp. 1-6. <https://doi.org/10.1109/RAEEUCCI57140.2023.10134069>
- [3] Khonina, S.N., Kazanskiy, N.L., Oseledets, I.V., Nikonorov, A.V., Butt, M.A. (2024). Synergy between artificial intelligence and hyperspectral imaging—A review. *Technologies*, 12(9): 163. <https://doi.org/10.3390/technologies12090163>
- [4] McKnight, S., Tunukovic, V., MacKinnon, C., Pierce, G., Mohseni, E., MacLeod, C.N., O'Hare, T. (2023). Using neural architecture search to discover a convolutional neural network to detect defects from volumetric ultrasonic testing data of composites. In BINDT-60th Annual British Conference on NDT (NDT 2023), UK, Northampton, United Kingdom.
- [5] Wen, L., Zhang, Y., Gao, L., Li, X., Li, M. (2023). A new multiscale multiattention convolutional neural network for fine-grained surface defect detection. *IEEE Transactions on Instrumentation and Measurement*, 72: 1-11. <https://doi.org/10.1109/TIM.2023.3271743>
- [6] Liu, H.I., Galindo, M., Xie, H., Wong, L.K., Shuai, H.H., Li, Y.H., Cheng, W.H. (2024). Lightweight deep learning for resource-constrained environments: A survey. *ACM Computing Surveys*, 56(10): 267. <https://doi.org/10.1145/3657282>
- [7] De Silva, D., Sierla, S., Alahakoon, D., Osipov, E., Yu, X., Vyatkin, V. (2020). Toward intelligent industrial informatics: A review of current developments and future directions of artificial intelligence in industrial applications. *IEEE Industrial Electronics Magazine*, 14(2): 57-72. <https://doi.org/10.1109/MIE.2019.2952165>
- [8] Zhang, C., Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. *Journal of Industrial Information Integration*, 23: 100224. <https://doi.org/10.1016/j.jii.2021.100224>
- [9] Lee, D., Yoon, S.N. (2021). Application of artificial intelligence-based technologies in the healthcare industry: Opportunities and challenges. *International Journal of Environmental Research and Public Health*, 18(1): 271. <https://doi.org/10.3390/ijerph18010271>
- [10] Gade, K., Geyik, S.C., Kenthapadi, K., Mithal, V., Taly, A. (2019). Explainable AI in industry. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, Anchorage, AK, USA, pp. 3203-3204. <https://doi.org/10.1145/3292500.3332281>
- [11] Patel, V., Shah, M. (2022). Artificial intelligence and machine learning in drug discovery and development. *Intelligent Medicine*, 2(3): 134-140.

- <https://doi.org/10.1016/j.imed.2021.10.001>
- [12] Nahavandi, D., Alizadehsani, R., Khosravi, A., Acharya, U.R. (2022). Application of artificial intelligence in wearable devices: Opportunities and challenges. *Computer Methods and Programs in Biomedicine*, 213: 106541. <https://doi.org/10.1016/j.cmpb.2021.106541>
- [13] Cioffi, R., Travaglioni, M., Piscitelli, G., Petrillo, A., De Felice, F. (2020). Artificial intelligence and machine learning applications in smart production: Progress, trends, and directions. *Sustainability*, 12(2): 492. <https://doi.org/10.3390/su12020492>
- [14] Liu, M.T., Dong, S., Zhu, M. (2021). The application of digital technology in gambling industry. *Asia Pacific Journal of Marketing and Logistics*, 33(7): 1685-1705. <https://doi.org/10.1108/APJML-11-2020-0778>
- [15] Hallur, G.G., Prabhu, S., Aslekar, A. (2021). Entertainment in era of AI, big data & IoT. *Digital Entertainment: The Next Evolution in Service Sector*, pp. 87-109. https://doi.org/10.1007/978-981-15-9724-4_5
- [16] Pereira, F., Carvalho, V., Vasconcelos, R., Soares, F. (2022). A review in the use of artificial intelligence in textile industry. In *Innovations in Mechatronics Engineering*, Minho, Portugal, pp. 377-392. https://doi.org/10.1007/978-3-030-79168-1_34
- [17] Anantrasirichai, N., Bull, D. (2022). Artificial intelligence in the creative industries: A review. *Artificial Intelligence Review*, 55(1): 589-656. <https://doi.org/10.1007/s10462-021-10039-7>
- [18] Kaur, D.N., Sahdev, S.L., Sharma, D., Siddiqui, L. (2020). Banking 4.0: The Influence of Artificial Intelligence on the Banking Industry & How AI Is Changing the Face of Modern Day Banks. SSRN.
- [19] Sikka, M. P., Sarkar, A., Garg, S. (2024). Artificial intelligence (AI) in textile industry operational modernization. *Research Journal of Textile and Apparel*, 28(1): 67-83. <https://doi.org/10.1108/RJTA-04-2021-0046>
- [20] Dhamija, P., Bag, S. (2020). Role of artificial intelligence in operations environment: A review and bibliometric analysis. *The TQM Journal*, 32(4): 869-896. <https://doi.org/10.1108/TQM-10-2019-0243>
- [21] He, K., Zhang, X., Ren, S., Sun, J. (2016). Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, USA, pp. 770-778. <https://doi.org/10.1109/CVPR.2016.90>
- [22] Howard, A., Sandler, M., Chu, G., Chen, L.C., et al. (2019). Searching for mobilenetv3. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, Seoul, Korea (South), pp. 1314-1324. <https://doi.org/10.1109/ICCV.2019.00140>
- [23] Tan, M., Le, Q.V. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. *arXiv preprint* [arXiv:1905.11946](https://arxiv.org/abs/1905.11946). <https://doi.org/10.48550/arXiv.1905.11946>
- [24] Qin, R., Chen, N., Huang, Y. (2022). EDDNet: An efficient and accurate defect detection network for the industrial edge environment. In *2022 IEEE 22nd International Conference on Software Quality, Reliability and Security (QRS)*, Guangzhou, China, pp. 854-863.