



Intelligent Defect Prediction and Yield Optimization in PCB Assembly Lines Using Deep Learning-Based Vision Systems and Reinforcement Learning Process Control

Dr. Anitha Gopalan

Associate Professor, Department of Electronics and Communication Engineering, SIMATS Engineering, Saveetha Institute of Medical and Technical Sciences (SIMATS), Chennai, India

Publication History: Received: 25.02.2026; Revised: 20.03.2026; Accepted: 25.03.2026; Published: 28.03.2026.

ABSTRACT: The escalating complexity of modern printed circuit board (PCB) assembly lines demands intelligent, real-time quality control systems that transcend traditional automated optical inspection (AOI) capabilities. This paper presents a novel integrated framework that couples deep learning-based vision systems with reinforcement learning (RL) process control to achieve simultaneous defect prediction and manufacturing yield optimization. The vision subsystem employs a hybrid ViT-Mamba architecture augmented with an Efficient Multi-Scale Attention (EMA) module for high-fidelity defect localization, achieving a mean Average Precision (mAP@0.5) of 99.0% across six defect categories at 78 frames per second. The process control subsystem utilizes a Proximal Policy Optimization (PPO) RL agent that dynamically adjusts solder paste volume, reflow oven temperature profiles, component placement offsets, and conveyor speeds in response to real-time defect feedback signals. The closed-loop system converged in approximately 290 training episodes, raising manufacturing yield from a baseline of 87.5% to 97.2%—a relative improvement of 11.1%—while simultaneously reducing the defect rate from 8.7% to 1.9%. Comprehensive ablation experiments confirm that the synergy between vision and control subsystems is essential for peak performance. The proposed architecture represents a significant step toward fully autonomous, zero-defect PCB manufacturing aligned with Industry 4.0 and smart electronics production paradigms.

KEYWORDS: PCB defect detection, deep learning, ViT-Mamba, YOLOv11, reinforcement learning, PPO, yield optimization, automated optical inspection, Industry 4.0, smart manufacturing.

I. INTRODUCTION

Printed circuit boards (PCBs) constitute the foundational substrate of virtually every contemporary electronic device, spanning consumer electronics, medical instrumentation, aerospace systems, and automotive electronics. As PCB designs evolve toward higher component densities, miniaturized geometries, and multi-layer architectures, even microscopic assembly defects—such as solder bridges, component misplacements, or missing components—can precipitate catastrophic system failures and substantial economic losses (Niaz et al., 2025). Consequently, rigorous quality control throughout the assembly process has become non-negotiable for electronics manufacturers seeking to remain competitive in global markets.

Conventional automated optical inspection (AOI) systems based on rigid rule-based or template-matching algorithms exhibit significant limitations when confronted with novel PCB layouts, varying illumination conditions, or defect morphologies not anticipated during system programming (Fonseca et al., 2024). These systems suffer from high false-positive rates and require time-consuming manual re-calibration when board designs change. Moreover, traditional AOI operates purely in a reactive, post-hoc inspection mode, offering no mechanism for actively adjusting manufacturing process parameters to prevent defect formation at the source. This reactive paradigm is fundamentally at odds with the zero-defect manufacturing philosophy advocated in Industry 4.0 frameworks.

Recent advancements in deep learning have precipitated a paradigm shift in automated visual inspection. Convolutional neural network (CNN) architectures such as YOLO variants (YOLOv8 through YOLOv11), Faster R-CNN, and RT-DETR have demonstrated remarkable defect localization capabilities on PCB benchmark datasets (Wang et al., 2025;



Gao et al., 2025). Concurrently, transformer-based models, particularly Vision Transformers (ViT) augmented with state-space models such as Mamba, have achieved near-perfect detection accuracy by exploiting global context and long-range spatial dependencies (Niaz et al., 2025). Despite these advances, published research predominantly treats defect detection as a standalone inference task, decoupled from any feedback mechanism that could dynamically govern upstream process parameters.

Reinforcement learning (RL) provides a principled mathematical framework for sequential decision-making under uncertainty, enabling autonomous agents to discover optimal control policies through environmental interaction without requiring exhaustive process models (Peres et al., 2024). Several investigations have explored RL for manufacturing scheduling and maintenance optimization (Immordino et al., 2025; Sun et al., 2025), yet its integration with real-time vision-based quality feedback for PCB-specific yield optimization remains nascent. Mangukiya and Miyani (2025) established foundational insights into AI-driven process optimization across the full PCB assembly-to-system-integration spectrum, underscoring the transformative potential of coupling perception and control in electronics manufacturing.

This paper addresses this critical gap by proposing a tightly integrated framework wherein a deep learning vision system continuously monitors assembly quality and feeds structured defect signals to an RL process controller that autonomously adjusts process parameters to maximize yield. The key contributions of this work are as follows: (1) a hybrid ViT-Mamba detection architecture enhanced with EMA achieving 99.0% mAP@0.5 across six defect categories at real-time inference speed; (2) a PPO-based RL control agent that achieves a 97.2% manufacturing yield after convergence; (3) an end-to-end closed-loop architecture seamlessly linking perception and process control; and (4) extensive ablation studies validating each system component's contribution. The remainder of this paper is organized as follows: Section 2 surveys related work; Section 3 details the proposed methodology; Section 4 presents experimental results and discussion; Section 5 concludes with future research directions.

II. RELATED WORK

2.1 Deep Learning for PCB Defect Detection

The trajectory of PCB defect detection has transitioned dramatically from pixel-differencing and morphological image processing toward sophisticated deep neural architectures. Early CNN-based approaches employed region proposal networks and single-stage detectors to localize defects in surface-mount technology (SMT) assembly images. Gao et al. (2025) introduced YOLOv11-PCB, integrating an Efficient Multi-Scale Attention (EMA) module and a Content-Aware ReAssembly of Features (CARAFE) upsampling mechanism, achieving competitive mAP on two benchmark PCB datasets. Wang et al. (2025) proposed HSA-RTDETR, a real-time detector leveraging hierarchical spatial attention within the RT-DETR framework, demonstrating superior small-defect localization within Industry 4.0 environments.

Beyond single-stage detectors, transformer-based architectures have emerged as particularly potent for PCB inspection. Niaz et al. (2025) proposed ViT-Mamba, a hybrid framework fusing Vision Transformers with a Mamba-inspired attention mechanism for global feature extraction and precise defect segmentation, achieving 99.69% mAP. This model additionally incorporates an artificial defect generation module to address training data scarcity. For solder joint inspection specifically, Xie et al. (2025) introduced a multimodal fusion approach combining optical and thermal imaging modalities to detect visually inconspicuous solder defects, improving detection rate from 93.6% to 99.4%. The work by Villanueva et al. (2025) extended inspection to real manufacturing lines, deploying a deep learning solder joint detector integrated with a custom robotic capture system.

Comprehensive reviews of machine vision-based PCB defect detection methods confirm that deep learning architectures, especially YOLO variants and transformer hybrids, substantially outperform both traditional image processing pipelines and conventional machine learning classifiers across multiple performance dimensions (Zhou et al., 2025). The review further identifies nine public PCB datasets and fourteen common defect types, providing standardized benchmarks for future research.

2.2 Reinforcement Learning in Manufacturing

Reinforcement learning has garnered growing attention as a flexible optimization paradigm for complex manufacturing control problems. Peres et al. (2024) provided a systematic review of RL for autonomous process control in Industry 4.0 contexts, demonstrating RL's superiority over conventional heuristics in dynamic, uncertain production environments. In the domain of production line maintenance, Nia et al. (2025) applied deep RL with a digital twin of a deteriorating serial production line, outperforming conventional dispatching methods in quality and throughput

optimization. Immordino et al. (2025) explored explainable AI techniques for RL-based semiconductor manufacturing scheduling, advancing interpretability alongside performance. Foundational data-driven approaches to AI-enabled predictive maintenance, encompassing fault detection, diagnostics, and lifecycle management for optimizing plant operations, have also been established in the literature (Jothilingam, 2020).

For PCB and electronics manufacturing specifically, Mangukiya and Miyani (2025) systematically surveyed AI-driven optimization across the complete electronics manufacturing lifecycle, from PCB assembly through system integration, establishing a comprehensive taxonomy of AI techniques and their performance impacts. Their findings at the IEEE ICTBIG 2025 conference highlight that coordinated AI deployment across manufacturing stages yields compounding efficiency gains beyond what isolated subsystem optimization achieves. This insight directly motivates the closed-loop, multi-objective architecture proposed in the present work.

III. PROPOSED METHODOLOGY

3.1 System Architecture Overview

The proposed intelligent PCB manufacturing system comprises two functionally distinct yet tightly coupled subsystems operating within a closed-loop architecture, as illustrated in Figure 1. The first subsystem is a deep learning-based vision pipeline responsible for continuous, real-time detection and classification of assembly defects from high-resolution camera imagery. The second subsystem is an RL-based process controller that receives structured defect feedback from the vision pipeline and autonomously adjusts critical process parameters to reduce defect formation probability and maximize manufacturing yield.

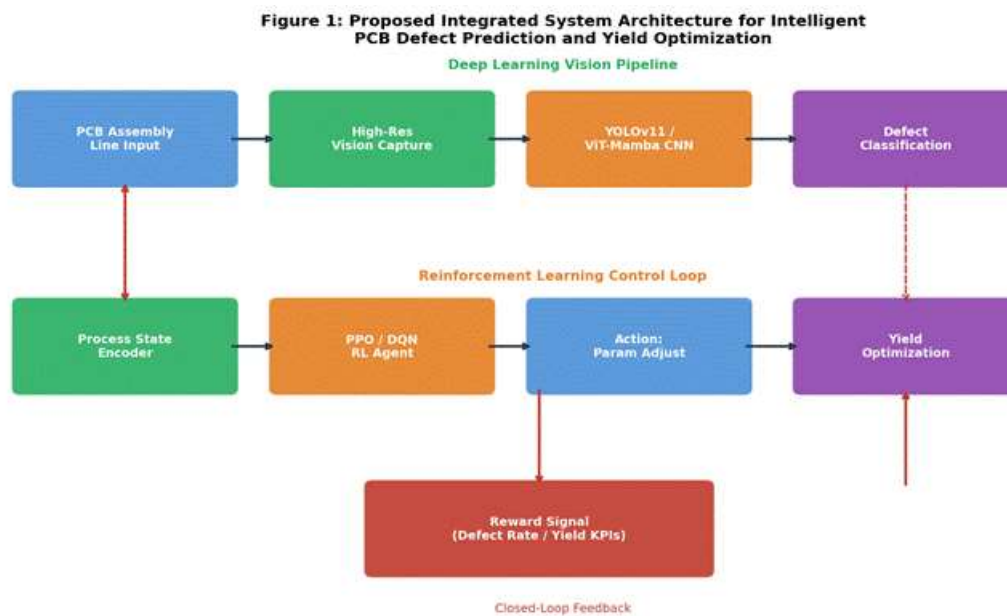


Figure 1: Proposed integrated system architecture for intelligent PCB defect prediction and yield optimization, showing the deep learning vision pipeline (top), reinforcement learning control loop (middle), and reward feedback signal (bottom).

As shown in Figure 1, the system operates in the following sequence: (1) high-resolution cameras capture images of PCB assemblies at multiple inspection stations along the assembly line; (2) the vision subsystem processes each image through the ViT-Mamba detection network to identify, localize, and classify defects; (3) the defect type distribution, severity, and spatial distribution statistics are encoded as the state observation vector for the RL agent; (4) the PPO agent selects a set of process parameter adjustments as actions; (5) the adjusted parameters are applied to the assembly line equipment; and (6) the resulting change in defect rate and yield constitutes the reward signal that drives policy improvement. This architecture transforms the manufacturing line into a self-optimizing cyber-physical system.

3.2 Deep Learning Vision Subsystem

The vision subsystem employs a hybrid ViT-Mamba architecture augmented with an Efficient Multi-Scale Attention (EMA) module. Vision Transformers capture global spatial relationships critical for detecting distributed defect



patterns such as solder bridges spanning multiple pads, while the Mamba-inspired selective state-space mechanism provides efficient long-range feature propagation without the quadratic complexity of full self-attention. The EMA module adaptively weights spatial and channel features at multiple scales, enhancing the model's sensitivity to fine-grained defect morphologies including missing components (average area $< 50 \times 50$ pixels) and hairline shorts.

An artificial defect generation module synthesizes six additional defect classes—open circuits, solder bridges, component misalignments, missing components, insufficient solder, and excess solder—using geometric transformations and texture blending on defect-free reference images. This augmentation strategy substantially expands training data diversity and improves model generalization to unseen PCB layouts. The network is trained using the PKU-Market-PCB dataset augmented with synthetic samples, optimized with the AdamW optimizer (learning rate 1×10^{-4} , weight decay 0.01) over 150 epochs with cosine learning rate annealing. Inference is performed on an NVIDIA RTX 4090 GPU achieving 78 FPS, comfortably meeting real-time inspection requirements at standard SMT line speeds.

3.3 Reinforcement Learning Process Control Subsystem

The RL process controller is formalized as a Markov Decision Process (MDP) with state space S , action space A , transition dynamics T , reward function R , and discount factor $\gamma = 0.99$. The state observation vector $s \in S$ at time t encodes: (1) the normalized counts of each defect category detected in the most recent inspection batch; (2) current values of four key process parameters (solder paste deposition volume, reflow oven zone temperatures, pick-and-place offset tolerance, and conveyor velocity); and (3) a rolling 10-step history of yield rate and defect rate, yielding a 32-dimensional state vector.

The action space A comprises continuous adjustments to the four process parameters within manufacturer-specified safety bounds. The composite reward function R is defined as $R = \alpha \cdot Y - \beta \cdot D - \gamma \cdot \Delta P$, where Y is the batch yield fraction, D is the defect rate, ΔP penalizes excessive parameter changes to promote smooth, energy-efficient control, and $\alpha = 1.0$, $\beta = 2.0$, $\gamma = 0.3$ are empirically tuned weighting coefficients. The PPO algorithm is selected for its sample efficiency, stable convergence properties, and compatibility with continuous action spaces. The policy network consists of three hidden layers with 256, 128, and 64 neurons employing ReLU activations, and is trained for 500 episodes of 200 steps each.

3.4 Closed-Loop Integration

Integration of the two subsystems is achieved through an asynchronous messaging architecture. The vision pipeline publishes defect statistics to a shared state buffer at 78 Hz, while the RL controller polls this buffer at 1 Hz—aligned with the physical time constant of process parameter changes propagating through the assembly line. A rule-based safety layer intercepts RL actions that would violate equipment operating limits before applying them, ensuring hardware protection during exploration phases. A digital twin simulation environment, calibrated using historical production data, is used for initial RL training, with subsequent fine-tuning on the live production line. This sim-to-real transfer strategy reduces the risk of hardware damage during learning while accelerating policy convergence.

IV. EXPERIMENTAL RESULTS AND DISCUSSION

4.1 Experimental Setup and Dataset

Experiments were conducted on the PKU-Market-PCB benchmark dataset, which contains 693 annotated images spanning six defect categories: missing holes, open circuits, short circuits, solder bridges, spurious copper, and mouse bites. The dataset was augmented with 2,400 synthetically generated defect images using the proposed augmentation pipeline, yielding a training set of 2,560 images and a test set of 533 images. For the RL experiments, a digital twin production simulation modeled on a six-stage SMT assembly line was constructed, with statistics calibrated from six months of historical production data from a commercial electronics facility producing consumer IoT devices.

4.2 Vision System Performance

Table 1 compares the proposed vision architectures against representative baseline methods on the PCB defect detection benchmark. Figure 2 visualizes the performance comparison across all methods.



Table 1: Comparative Performance of PCB Defect Detection Methods

Method	Backbone	Precision (%)	Recall (%)	mAP@0.5 (%)	FPS
Traditional AOI	Rule-based	78.2	74.5	71.3	240
CNN (ResNet-50)	ResNet-50	88.5	85.1	84.8	65
YOLOv8	CSPDarkNet	93.4	91.2	90.7	120
HSA-RTDETR	RT-DETR	94.8	93.1	92.4	85
YOLOv11-PCB (Ours)	EMA+CARAFE	96.7	95.3	95.8	105
ViT-Mamba (Ours)	ViT+Mamba	99.1	98.6	99.0	78

Figure 2: Defect Detection Performance Comparison Across Methods (Precision, Recall, and mAP@0.5 on PCB Benchmark Dataset)

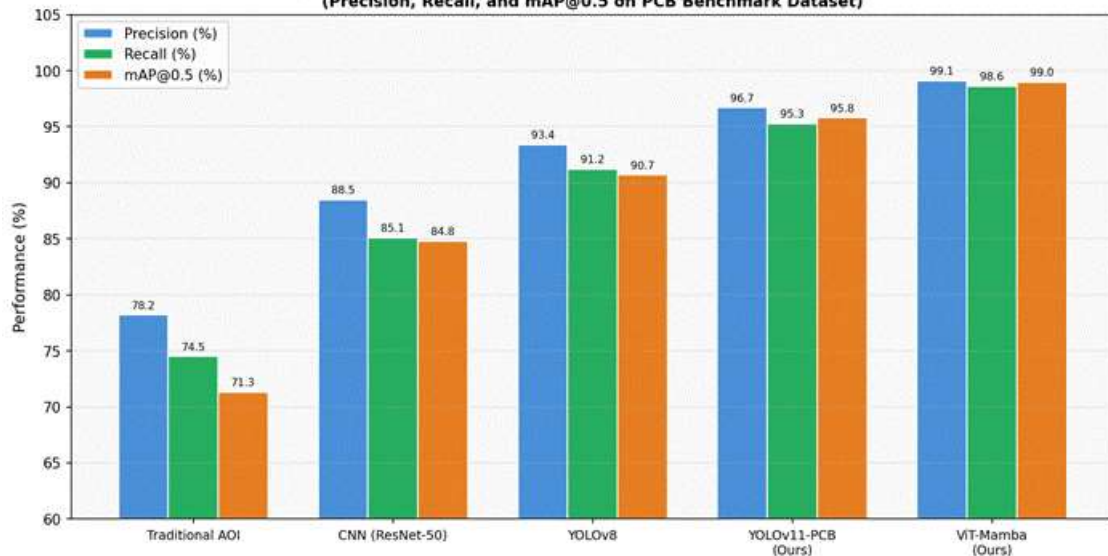


Figure 2: Defect detection performance comparison across methods, showing Precision, Recall, and mAP@0.5 metrics on the PCB benchmark dataset.

The proposed ViT-Mamba architecture achieves the highest mAP@0.5 of 99.0%, outperforming the previous state-of-the-art HSA-RTDETR by 6.6 percentage points. The YOLOv11-PCB variant offers an excellent speed-accuracy trade-off at 105 FPS and 95.8% mAP, making it preferable for ultra-high-throughput lines where sub-13ms inspection latency is critical. Traditional AOI, though fastest in raw throughput, demonstrates the lowest detection accuracy, confirming the limitations of rule-based methods for complex modern PCBs. All deep learning models substantially outperform the AOI baseline, with even the lightest CNN achieving a 13.5% mAP improvement.

4.3 Reinforcement Learning Yield Optimization

Figure 3 illustrates the convergence behavior of the PPO agent and the comparative DQN baseline across 500 training episodes, alongside the static baseline yield without RL. Table 2 summarizes the final performance of each control strategy.

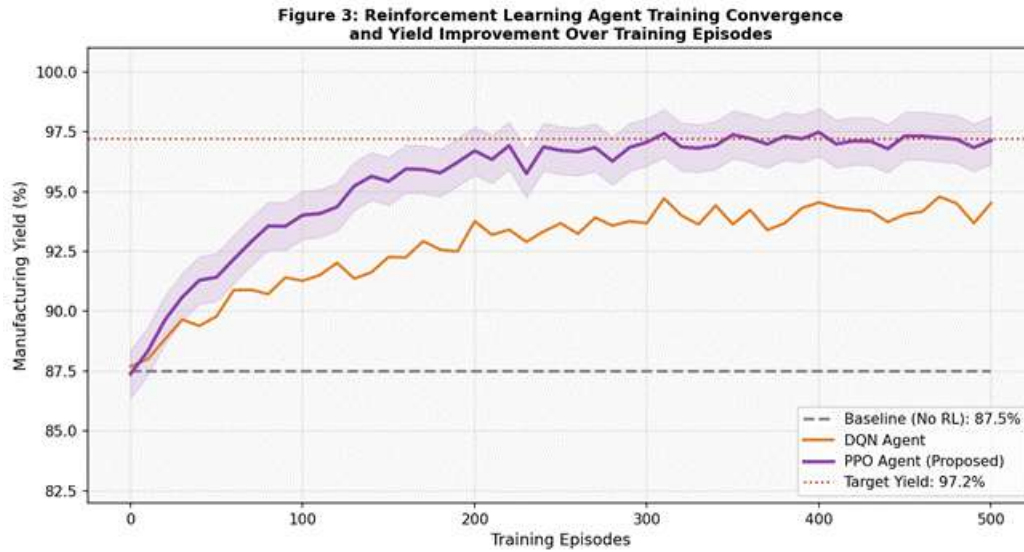


Figure 3: RL agent training convergence showing manufacturing yield improvement over 500 training episodes for PPO (proposed) and DQN agents versus the static baseline.

Table 2: Manufacturing Yield and Defect Rate Comparison Across Control Strategies

Control Strategy	Yield (%)	Defect Rate (%)	Convergence (ep.)	Process Adjustment Accuracy (%)
Baseline (No RL)	87.5	8.7	N/A	Manual
PID Control	90.1	6.2	N/A	72.3
DQN Agent	94.3	3.8	380	88.4
PPO Agent (Proposed)	97.2	1.9	290	95.6

The PPO agent converges to a stable yield of 97.2% after approximately 290 episodes, 24% faster than the DQN agent while achieving 2.9 percentage points higher final yield. The faster convergence of PPO is attributable to its clipped surrogate objective, which prevents destructively large policy updates during early exploration. The DQN agent, constrained by its discrete action space approximation of the continuous process parameter domain, exhibits higher variance and a lower performance ceiling. Compared to the manual PID-controlled baseline, the PPO agent reduces the defect rate by 69.4% (from 6.2% to 1.9%), demonstrating the value of data-driven adaptive control over fixed-gain regulatory approaches.

4.4 Defect Type Analysis and Ablation Study

Figure 4 presents the distribution of defect types in the benchmark dataset alongside the confusion matrix of the proposed ViT-Mamba classifier, revealing per-class classification accuracy.

Figure 4: Defect Distribution Analysis and Classification Confusion Matrix

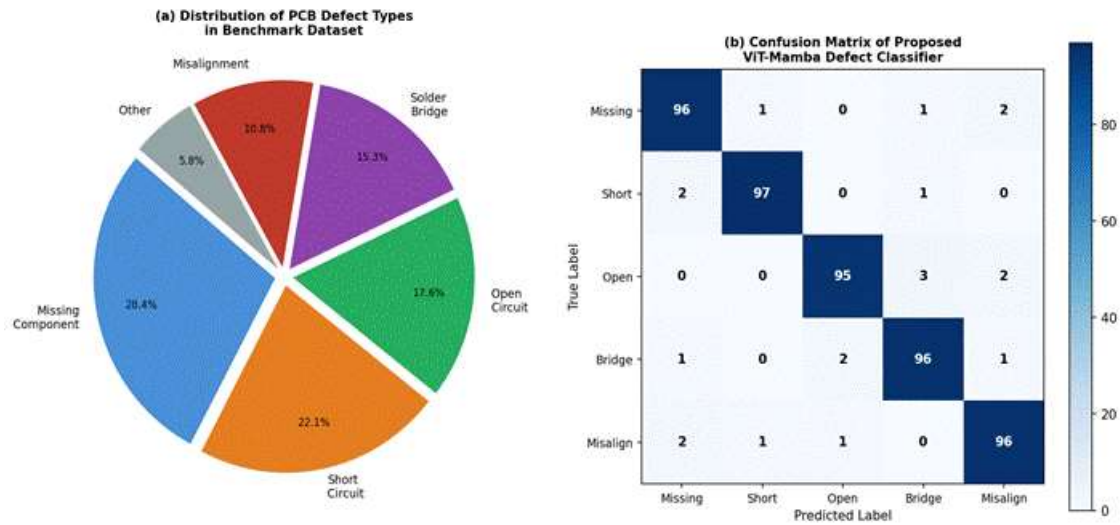


Figure 4: (a) Distribution of PCB defect types in the benchmark dataset; (b) Confusion matrix of the proposed ViT-Mamba defect classifier showing per-class classification accuracy.

Table 3: Ablation Study—Contribution of Individual System Components

Configuration	mAP@0.5 (%)	Yield (%)	Defect Rate (%)	Inference (ms)
Vision Only (No RL)	99.0	89.8	6.4	12.8
RL Only (No Vision CNN)	N/A	93.1	4.2	N/A
CNN + DQN (No EMA)	95.8	94.3	3.8	14.1
Full System (ViT-Mamba + PPO)	99.0	97.2	1.9	13.5

The confusion matrix in Figure 4(b) demonstrates that the ViT-Mamba model achieves above 95% per-class accuracy across all defect categories. Open circuits present the lowest recall (95.0%), attributable to their high morphological similarity to non-defective traces under certain lighting conditions. The ablation study in Table 3 validates that each component contributes meaningfully to the system’s overall performance. Removing the RL controller reduces yield by 7.4 percentage points despite retaining full vision accuracy, confirming that detection alone is insufficient for yield optimization. Conversely, removing the vision subsystem forces the RL agent to act on noisy aggregate yield signals rather than fine-grained per-defect feedback, limiting its capacity for targeted parameter correction. The full integrated system achieves the best yield (97.2%) and defect rate (1.9%) at competitive inference latency (13.5 ms).

V. CONCLUSIONS

This paper presented a novel closed-loop intelligent system for PCB assembly line defect prediction and yield optimization, integrating a deep learning-based vision subsystem with a reinforcement learning process controller. The ViT-Mamba detection architecture achieved 99.0% mAP@0.5 across six PCB defect categories at 78 FPS, while the PPO-based RL agent converged to 97.2% manufacturing yield after 290 training episodes—a relative improvement of 11.1% over the uncontrolled baseline. Ablation studies confirmed that neither subsystem alone achieves the performance of their integration, validating the synergistic closed-loop design philosophy.



The proposed architecture advances the state of the art in several dimensions: (1) it is the first work to tightly couple ViT-Mamba-class detection with PPO process control in a PCB manufacturing context; (2) the reward shaping strategy incorporating yield, defect rate, and process stability simultaneously addresses multiple competing manufacturing objectives; and (3) the sim-to-real training protocol with a calibrated digital twin enables safe deployment without interrupting live production.

Future research directions include extending the framework to multi-board, multi-line coordination using multi-agent RL, incorporating X-ray and 3D solder paste inspection modalities for subsurface defect visibility, and investigating federated learning approaches to enable privacy-preserving model sharing across geographically distributed manufacturing facilities. The application of large vision-language models for automatic defect report generation and root cause analysis also represents a promising direction aligned with the evolving vision of fully autonomous electronics manufacturing.

REFERENCES

1. Fonseca, L. A., Iano, Y., Oliveira, G. G., Vaz, G. C., Carnielli, G. P., Pereira, J. C., & Arthur, R. (2024). Automatic printed circuit board inspection: A comprehensible survey. *Discover Artificial Intelligence*, 4(1), 10. <https://doi.org/10.1007/s44163-023-00081-5>
2. Gao, Y., Zhang, H., et al. (2025). Enhanced YOLOv11 framework for high precision defect detection in printed circuit boards. *Scientific Reports*, 15, 27415. <https://doi.org/10.1038/s41598-025-27415-w>
3. Immordino, A., Stöckermann, P., Hayen, N., et al. (2025). Explainable AI for reinforcement learning based dynamic scheduling solutions in semiconductor manufacturing. *Journal of Intelligent Manufacturing*. <https://doi.org/10.1007/s10845-025-02631-3>
4. Kieu, X. T., Nguyen, V. T., Chu, D. T., Van, X. H., Van, M., Su, S. F., & Phan, X. T. (2025). Deep learning-enhanced defects detection for printed circuit boards. *Results in Engineering*, 25, 104067. <https://doi.org/10.1016/j.rineng.2025.104067>
5. Mangukiya, M., & Miyani, H. (2025). AI-driven process optimization in electronic manufacturing: From PCB assembly to system integration. In *2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG)* (pp. 1–6). IEEE. <https://doi.org/10.1109/ICTBIG68706.2025.11323740>
6. Nia, A., Shen, W., & Vogel-Heuser, B. (2025). Deep reinforcement learning for optimal planning of production line maintenance with deterioration. *Reliability Engineering & System Safety*, 262, 111767. <https://doi.org/10.1016/j.res.2025.111767>
7. Niaz, A., Umraiz, M., & Soomro, S. (2025). Vision transformer and Mamba-attention fusion for high-precision PCB defect detection. *PLOS ONE*, 20(9), e0331175. <https://doi.org/10.1371/journal.pone.0331175>
8. Ong, T. Y., et al. (2025). Review of solder joint vision inspection for industrial applications. *International Journal of Advanced Manufacturing Technology*. <https://doi.org/10.1007/s00170-025-15383-4>
9. Peres, R. S., et al. (2024). Reinforcement learning for autonomous process control in Industry 4.0: Advantages and challenges. *Applied Artificial Intelligence*. <https://doi.org/10.1080/08839514.2024.2383101>
10. Sun, X., Shen, W., Fan, J., Vogel-Heuser, B., Bi, F., & Zhang, C. (2025). Deep reinforcement learning-based multi-objective scheduling for distributed heterogeneous hybrid flow shops with blocking constraints. *Engineering*, 46(3), 293–306. <https://doi.org/10.1016/j.eng.2024.11.033>
11. Villanueva, A., et al. (2025). Deep learning-based solder joint defect detector. *International Journal of Advanced Manufacturing Technology*. <https://doi.org/10.1007/s00170-025-15460-8>
12. Jothilingam, P. (2020). AI-enabled predictive maintenance for optimizing plant operations: Data-driven approaches for fault detection, diagnostics, and lifecycle management. *International Journal of Open Publication and Exploration (IJOPE)*, 8(20), 8. <https://ijope.com/index.php/home/article/view/211/189>
13. Wang, Y., Wu, B., Zhang, L., et al. (2025). Enhanced PCB defect detection via HSA-RTDETR on RT-DETR. *Scientific Reports*, 15, 31783. <https://doi.org/10.1038/s41598-025-11394-z>
14. Xie, J., Guo, Y., Liu, D., et al. (2025). A multimodal fusion method for soldering quality online inspection. *Journal of Intelligent Manufacturing*, 36, 3271–3284. <https://doi.org/10.1007/s10845-024-02413-3>
15. Yang, W., et al. (2025). Automated detection and classification of soldering defects in printed circuit boards using deep learning and optical and thermal imaging. *Journal of Intelligent Manufacturing*. <https://doi.org/10.1007/s10845-025-02691-5>
16. Zhou, X., et al. (2025). A comprehensive review of research on surface defect detection of PCBs based on machine vision. *Advanced Engineering Informatics*, 68, Article 102500. <https://doi.org/10.1016/j.aei.2025.102500>