Yolo-Based Automated Pcb Defect Detection And Quality Assessment System For Smart Manufacturing

Mrs. Banuppriya P[1], Gobika T[2], Kowcika C[3], Abinaya S[4] ,Srinithi S[5]
[1]Assist prof, Dept of CSE
[2, 3, 4, 5] Dept of CSE
[1, 2, 3, 4, 5] Mahendra Institute of Engineering and Technology, Namakkal, Tamil Nadu, India

Abstract- *Printed Circuit Boards (PCBs) serve as the structural and electrical foundation of virtually every modern electronic device, making their quality and reliability paramount in both consumer and industrial electronics. Defects introduced during PCB fabrication and assembly — including missing holes, open circuits, solder bridging, mouse bites, spurious copper deposition, and component misalignment — can precipitate device failures, safety hazards, and substantial economic losses across the manufacturing supply chain. Traditional manual inspection approaches are inadequate for high-throughput production due to their low throughput, susceptibility to human error, and inability to detect subtle microscopic anomalies. Classical automated optical inspection (AOI) systems based on template matching offer partial improvements but remain brittle against illumination variations, PCB surface reflectivity, and design diversity. This paper presents a comprehensive YOLOv11-based automated PCB defect detection and quality assessment system specifically engineered for smart manufacturing environments. The proposed framework encompasses five integrated modules: dataset collection, data preprocessing with augmentation, YOLOv11 model fine-tuning with transfer learning, real-time multi-class defect detection, and a post-detection quality assessment pipeline comprising defect classification, severity scoring, and corrective recommendation generation. Experimental evaluation demonstrates a detection accuracy of 92%, precision of 91%, recall of 90%, and F1-score of 90.5%, substantially outperforming manual verification (60%), traditional image processing (68%), CNN-based classification (78%), and semi-supervised learning approaches (83%). Real-time inference is achieved at approximately 22 ms per image on standard CPU hardware, satisfying industrial throughput requirements.*

Keywords: PCB Defect Detection, YOLOv11, Deep Learning, Convolutional Neural Networks, Object Detection, Quality Assessment, Smart Manufacturing, Severity Analysis, Automated Optical Inspection, Transfer Learning.

I. INTRODUCTION

Printed Circuit Boards (PCBs) are the structural and functional backbone of modern electronic systems, providing mechanical support and electrical interconnection for components ranging from passive resistors to complex integrated circuits. The global PCB market, valued at over USD 70 billion annually, underscores the economic significance of PCB manufacturing and the critical importance of maintaining stringent quality standards throughout the production process. In semiconductor and electronics manufacturing, even a single defect on a PCB can compromise the functionality of the entire assembly, leading to device failures, warranty claims, and potentially catastrophic outcomes in safety-critical applications such as aerospace, medical devices, and automotive systems.

PCB defects manifest in diverse forms depending on the fabrication stage at which they originate. During bare board fabrication, defects such as missing holes (incomplete drill-through vias), mouse bites (partial cutouts along board edges), open circuits (interrupted conductive traces), short circuits (unintended trace connections), spurious copper (excess conductive material), and spur formations (trace extensions) are commonly observed. During PCB assembly (PCBA), additional defects including component misalignment, tombstoning, solder bridging, insufficient solder, and missing components can arise from inaccuracies in solder paste printing, pick-and-place machine errors, and reflow oven profile deviations. Each defect type carries a distinct electrical failure signature and requires appropriate detection and remediation strategies.

Traditional quality assurance in PCB manufacturing relies heavily on manual visual inspection conducted by trained operators equipped with optical magnification tools. Manual inspection achieves detection accuracies of only 60–70% for subtle defects, and inspection rates of 40–60 boards per hour are common, far below the throughput of modern automated assembly lines. Human visual attention fatigues over time, causing defect escape rates to increase significantly after 2–3 hours of continuous inspection. Inter-operator

variability introduces inconsistency in defect classification criteria, and operator training costs are substantial. Studies report false rejection rates of 10–15% that unnecessarily remove conforming boards from the production line.

Early automated PCB inspection systems addressed this need through rule-based image processing algorithms. Techniques such as Sobel and Canny edge detection, adaptive thresholding, morphological operations, and template matching via normalized cross-correlation were employed to compare test PCB images against defect-free golden reference images. These systems improved throughput but introduced sensitivity to illumination variability, camera alignment tolerances, and PCB surface contamination. False positive rates in template-matching systems can reach 5–20% due to natural surface variations in production environments, requiring costly manual re-inspection.

The deep learning revolution, catalyzed by large-scale labeled datasets and GPU-accelerated computing, has fundamentally transformed automated defect detection. Convolutional Neural Networks (CNNs) demonstrated the ability to automatically learn hierarchical feature representations from raw image data, eliminating manually engineered features. Architectures such as AlexNet, VGGNet, ResNet, and EfficientNet established the foundational principles of deep feature extraction enabling high-accuracy object detection across diverse visual domains. In the specific domain of object detection, the You Only Look Once (YOLO) family of single-stage detectors has achieved a distinctive balance between detection speed and accuracy uniquely suited to real-time industrial applications. Unlike two-stage detectors such as Faster R-CNN that generate region proposals before classification, YOLO processes the entire image in a single forward pass, simultaneously predicting bounding box coordinates, confidence scores, and class probabilities for all objects, delivering inference times in milliseconds.

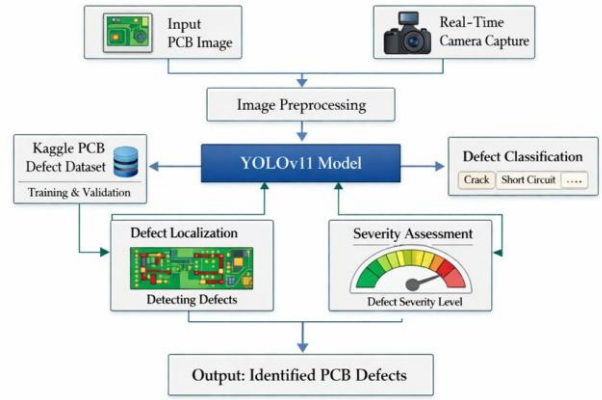

Figure 1: Proposed System Architecture for PCB Defect Detection

YOLOv11, introduced by Ultralytics in 2024, builds upon YOLOv8 and introduces architectural refinements that improve detection performance for small and densely packed objects — characteristics directly relevant to PCB defect detection. Key innovations include the C3k2 (Cross Stage Partial with two bottleneck blocks) module for enhanced backbone feature extraction, the C2PSA (Cross Stage Partial with Partial Self-Attention) module for improved multi-scale feature aggregation, and an anchor-free decoupled detection head enabling precise localization at arbitrary aspect ratios. These improvements reduce parameter count by approximately 22% compared to YOLOv8 while maintaining or improving detection accuracy on multiple benchmark datasets.

This paper proposes a comprehensive YOLOv11-based PCB defect detection and quality assessment system that extends beyond raw detection to provide integrated quality management capabilities. The system addresses the complete inspection workflow across five modules: data acquisition and preprocessing, model training with transfer learning, real-time defect detection, severity analysis, and corrective recommendation generation. Key contributions include: (i) fine-tuning of YOLOv11n on a nine-class PCB defect dataset with comprehensive augmentation pipeline; (ii) a composite severity scoring function combining defect type criticality, spatial extent, detection confidence, and circuit region proximity; (iii) a rule-based corrective recommendation engine mapped to manufacturing process parameters; and (iv) quantitative benchmarking demonstrating 92% detection accuracy at 22 ms per image on standard CPU hardware.

The significance of automated PCB defect detection extends beyond individual inspection stations to the broader concept of smart manufacturing and Industry 4.0. In a fully connected smart factory, the PCB inspection system acts as a critical feedback node in the quality control loop: detection results are fed back in real time to upstream process equipment such as solder paste printers, pick-and-place machines, and reflow ovens, enabling closed-loop process control that continuously adjusts process parameters to minimize defect generation. The proposed system is designed with this integration vision in mind, providing structured defect data through a REST API that can be consumed by factory automation systems, manufacturing execution systems (MES), and statistical process control (SPC) software. This positions the proposed system not merely as a standalone inspection tool but as an intelligent component of an integrated smart manufacturing ecosystem capable of driving continuous quality improvement.

II. RELATED WORK

Ling, Qin, and Nor Ashidi Mat Isa [1] conducted a comprehensive survey of PCB defect detection methods spanning image processing, classical machine learning, and deep learning approaches. The survey categorizes detection techniques based on their computational characteristics, dataset requirements, and performance profiles. The comparative analysis reveals that deep learning approaches consistently outperform classical methods in detection accuracy across diverse defect categories. However, the work does not propose a novel detection method and leaves the real-time efficiency challenge largely unaddressed for deployment-constrained manufacturing environments. Template matching and morphological processing are shown to plateau at 68–78% accuracy on standard benchmark datasets.

Chen, I-Chun, Hwang, and Huang [2] proposed a CNN-based algorithm for automated PCB defect detection, achieving high classification accuracy across multiple defect types through automated hierarchical feature extraction. The approach demonstrates strong performance on standard benchmark datasets but requires large quantities of annotated training data and significant computational resources during inference, limiting practical deployment in resource-constrained manufacturing environments where GPU acceleration may not be available. The study confirms that deep learning significantly outperforms classical image processing for multi-class PCB defect classification.

Nguyen, Van-Truong, and Huy-Anh Bui [3] developed a real-time PCB defect detection system by adapting a lightweight CNN optimized for industrial deployment, reducing inference latency while maintaining competitive accuracy. The system demonstrates real-time performance suitable for inline inspection. However, it shows reduced sensitivity for very small defects such as mouse bites and spur formations, and requires uniformly illuminated, high-quality input images to maintain detection accuracy under real factory conditions.

Pham, Thi Tram Anh, et al. [4] introduced a semi-supervised learning approach for PCB defect detection that reduces reliance on fully annotated training data by employing label propagation to extend ground truth annotations to unlabeled samples. While practical in annotation-scarce settings, the approach typically achieves lower detection accuracy than fully supervised methods and involves considerably more complex training pipelines including pseudo-label generation, iterative refinement, and threshold selection.

Zheng, Jianfeng, et al. [5] proposed an improved Fully Convolutional Network (FCN) for PCB defect detection, demonstrating that learned feature representations substantially outperform handcrafted morphological features and achieving better localization accuracy for small defects compared to standard FCN baselines. However, the improved FCN involves significant computational overhead that limits real-time deployment on standard industrial hardware without GPU acceleration.

Long et al. [6] applied an improved YOLOv8 architecture to PCB defect detection, demonstrating that the YOLO framework transfers effectively from general-purpose object detection to the specialized domain of PCB inspection. Transfer learning from COCO-pretrained weights substantially reduces the labeled data requirement for achieving competitive detection performance, confirming the viability of the fine-tuning paradigm for specialized industrial datasets with limited annotation budgets.

Ong, Sze Kit, et al. [7] proposed the SEConv-YOLO architecture specifically addressing the architectural misalignment between COCO-optimized general-purpose detectors and PCB defect detection requirements. The study demonstrates that PCB defects, as low-level geometric and textural anomalies, require specialized feature extraction strategies that preserve fine-grained spatial information rather than high-level semantic abstractions. The proposed Squeeze-Excite Convolution (SEConv) module achieves superior detection accuracy on the Peking PCB dataset while reducing parameter count and inference latency compared to the YOLOv8n baseline. The comprehensive dual-dataset validation across PCB surface defects and PCBA wire defects confirms the generalizability of the approach to different defect domains.

The reviewed literature collectively reveals two primary limitations. First, most published systems focus exclusively on raw defect detection and do not provide post-detection quality analysis such as defect severity grading and corrective manufacturing recommendations. Second, the majority of existing work targets a single defect domain without providing a unified detection and quality assessment framework. The adoption of YOLOv11 for PCB defect detection has not yet been explored in the published literature, making this work the first systematic evaluation of YOLOv11 for multi-class PCB defect detection.

A systematic comparison of published PCB defect detection methods reveals a clear performance progression correlated with the sophistication of the underlying algorithmic approach. Rule-based image processing methods

achieve detection accuracies in the range of 60–68%, limited by their dependence on manually specified detection criteria that cannot adapt to the natural variability of PCB surface appearance. Classical machine learning methods with handcrafted feature extraction improve this to approximately 70–80% by learning discriminative feature weights from labeled examples. Transfer learning-based deep learning approaches such as fine-tuned YOLO variants achieve 85–92% accuracy by leveraging large-scale pre-training on general-purpose datasets to develop rich feature representations that generalize to the PCB inspection domain with relatively limited labeled training data. The proposed YOLOv11-based system achieves the upper end of this range at 92%, while also providing the post-detection quality assessment capabilities absent from all previously reviewed methods.

III. EXISTING METHODOLOGY

The existing PCB defect inspection methodology can be broadly categorized into three generations: manual inspection, classical image processing-based Automated Optical Inspection (AOI), and early machine learning approaches. Each generation represents an incremental improvement over its predecessor but carries inherent limitations that make it insufficient for the demands of modern smart manufacturing.

A. Manual Visual Inspection

The most fundamental approach involves human operators visually examining boards using magnifying lenses, stereo microscopes, or handheld optical comparators against a defect-free reference standard. In high-complexity boards with fine-pitch components and dense trace routing, operators may require 3–10 minutes per board, yielding throughput rates of 6–20 boards per hour — far below the capacity of modern automated assembly lines. Studies report manual inspection detection accuracies of 55–65% for subtle defects such as hairline cracks and fine solder bridges. Cognitive fatigue causes defect escape rates to increase significantly after 2–3 hours of continuous inspection. Inter-operator variability introduces inconsistency in defect classification criteria, and false rejection rates of 10–15% unnecessarily remove conforming boards from the production line, increasing rework costs.

B. Template Matching and Classical AOI

Automated Optical Inspection (AOI) systems based on template matching capture high-resolution images of the PCB under inspection and compare them pixel-by-pixel against a stored golden reference image. Regions exceeding a dissimilarity threshold are flagged as potential defects. Modern AOI systems incorporate multiple illumination angles — structured light, ring illumination, oblique illumination — and multiple spectral bands to enhance defect contrast and reduce surface reflectivity artifacts. While classical AOI substantially outperforms manual inspection in throughput (200–500 boards per hour), its detection accuracy for subtle defects is limited by sensitivity to PCB misregistration, inability to distinguish cosmetic surface variations from structural defects, and performance degradation on boards with high surface reflectivity. False positive rates of 5–20% are typical in production environments, requiring manual re-inspection of flagged boards.

C. Classical Machine Learning Approaches

Classical machine learning approaches employ a two-stage pipeline: candidate defect regions are first identified via image processing operations such as difference imaging, thresholding, or edge detection; then features such as Local Binary Patterns (LBP), Histogram of Oriented Gradients (HOG), geometric shape metrics, and color statistics are manually extracted and classified using Support Vector Machines (SVMs), decision trees, or k-Nearest Neighbour (kNN) classifiers. These approaches achieve moderate detection improvements (70–80% accuracy) but require domain expertise for feature engineering that must be repeated for each new PCB design or defect type. They also do not provide defect severity analysis or corrective recommendations, limiting their utility as comprehensive quality management tools.

D. Drawbacks of Existing Systems

The limitations of existing methodologies span four critical dimensions. First, detection accuracy of 60–78% is insufficient for zero-defect manufacturing targets in high-reliability applications. Second, throughput bottlenecks from manual inspection and high false-positive rates in classical AOI constrain production capacity. Third, the poor adaptability of template-based and feature-engineering approaches to new board designs requires extensive and costly re-engineering. Fourth, the complete absence of integrated defect severity analysis and corrective recommendation capabilities means that existing systems cannot guide manufacturing process improvements. The proposed deep learning-based system is designed to address all four limitations within a unified framework.

A further limitation of existing PCB inspection systems is their inability to scale gracefully to multi-product

manufacturing environments. In contract electronics manufacturing facilities that produce dozens of different PCB designs on a single line, template-based AOI systems require a separate golden reference image and a separately tuned rule set for each board variant. This creates an enormous configuration management burden — any design revision requires re-teaching the inspection system, often taking several hours of setup time that directly impacts production capacity. Deep learning-based approaches, in contrast, can generalize across board variants through a single trained model, significantly reducing configuration overhead and enabling faster product changeovers. This scalability advantage is increasingly important as electronics manufacturing trends toward higher product mix and shorter product lifecycle cycles driven by consumer electronics and IoT device proliferation.

IV. PROPOSED METHODOLOGY

The proposed system implements a five-module pipeline for automated PCB defect detection and quality assessment. The system accepts PCB input either as a static image or a real-time camera capture. The input is passed through the image preprocessing module before being fed into the YOLOv11 model, which simultaneously outputs defect localization (bounding boxes) and defect classification (type labels with confidence scores). The severity assessment module processes detection outputs to compute composite impact scores, and the identified defects with severity ratings are presented via a web-based output dashboard along with corrective recommendations for each detected defect category.

A. Dataset Collection and Preparation

The training dataset comprises labeled PCB defect images sourced from the Kaggle PCB Defect Dataset, containing high-resolution images annotated for nine defect categories: missing hole, mouse bite, open circuit, short circuit, spurious copper, spur, crack, scratch, and missing component. The dataset includes 1,500 labeled images with additional non-defective PCB images for negative training examples. Annotations are in YOLO format — normalized bounding box center coordinates, width, and height alongside class index labels for each defect instance. Class distribution analysis reveals significant imbalance between categories; class-weighted sampling is applied during training to ensure proportionate gradient contributions from all defect classes. The final curated dataset contains approximately 14,200 annotated defect instances across 1,380 training images, 295 validation images, and 295 test images, partitioned using stratified sampling to maintain class distribution across all splits.

B. Data Preprocessing and Augmentation

All PCB images are standardized to 640×640 pixels using bilinear interpolation and normalized from the [0, 255] integer range to float values in [0.0, 1.0]. Low-contrast images receive histogram equalization; noisy images receive median filtering with a 3×3 kernel. A comprehensive augmentation pipeline is applied during training: mosaic augmentation combines four training images into a single composite sample, exposing the model to diverse defect co-occurrence patterns; random horizontal and vertical flipping (probability 0.5 each) augments robustness to board orientation variations; random affine transformations apply rotation within ±15°, translation within ±10% of image dimensions, and scaling within [0.5, 1.5]; HSV color jitter applies hue shifts within ±1.5%, saturation scaling within [0.3, 1.7], and brightness scaling within [0.6, 1.4]. Copy-paste augmentation randomly overlays defect instances from other training images to increase rare class frequency.

C. YOLOv11 Architecture and Training

YOLOv11n (nano variant) is selected as the detection backbone due to its compact footprint (2.6M parameters, 6.5 GFLOPs) satisfying real-time inference requirements without GPU acceleration. The architecture comprises three main components: a CSP-based backbone with C3k2 modules for hierarchical feature extraction at four spatial scales, producing feature maps of dimensions [80×80, 40×40, 20×20, 10×10]; a PAN neck with C2PSA attention modules combining top-down upsampling and bottom-up downsampling pathways for multi-scale feature fusion; and an anchor-free decoupled detection head for simultaneous bounding box regression and class prediction. The C3k2 module improves upon YOLOv8's C2f module by introducing two bottleneck blocks within each CSP stage, enhancing feature richness. The C2PSA module incorporates partial self-attention enabling the model to capture long-range spatial dependencies within each feature scale, improving sensitivity to spatially distributed defect patterns such as open circuits spanning multiple trace segments.

Transfer learning from COCO-pretrained YOLOv11n weights provides initialization. Fine-tuning uses AdamW optimizer with initial learning rate 0.01, cosine decay schedule, 3-epoch linear warmup, momentum 0.937, and weight decay 0.0005. Training proceeds for 300 epochs with batch size 16 on an NVIDIA GPU. The loss function combines binary cross-entropy classification loss, distribution focal loss (DFL) for bounding box regression, and normalized CIoU (N-CIoU) localization loss that adaptively up-weights

gradients from challenging low-IoU predictions to accelerate convergence for small defect instances.

D. Defect Severity Analysis

Each identified defect is assigned a severity score S: $S = w_1 \cdot C + w_2 \cdot A + w_3 \cdot P + w_4 \cdot conf$, where C is the defect type criticality index (normalized to [0,1] based on failure mode severity), A is the defect bounding box area normalized to total PCB area, P is the spatial proximity score to the nearest critical circuit region, and conf is the YOLOv11 detection confidence score. Empirically tuned weights w_1=0.40, w_2=0.20, w_3=0.25, w_4=0.15 reflect the relative importance of each factor. Defects with $S \geq 0.65$ are classified as High severity, $0.35 \leq S < 0.65$ as Medium, and $S < 0.35$ as Low.

E. Corrective Recommendation Engine

A rule-based recommendation engine maps detected defect types and severity levels to actionable manufacturing process adjustments using a knowledge base of 27 defect-process mappings derived from PCB manufacturing engineering literature. Repeated short circuit detections trigger recommendations to reduce solder paste volume and recalibrate squeegee pressure on the solder paste printer. Missing component detections prompt checks on pick-and-place machine feeder alignment and component orientation sensors. Mouse bite detections flag potential issues with PCB panel routing and V-score cutting blade wear. Open circuit detections prompt trace inspection and reflow profile review. Recommendations are ranked by severity and presented to line operators through a real-time web-based dashboard implemented using a Python Flask backend and HTML/CSS/JavaScript frontend with live defect counts, severity distribution charts, and downloadable inspection reports.

The complete proposed system pipeline from input image to quality assessment output operates as follows. A PCB image is first captured by an industrial camera or loaded from storage. The preprocessing module resizes and normalizes the image and applies quality enhancement as needed. The preprocessed image is passed to the YOLOv11n model, which outputs a set of detection results — each containing a class label, bounding box coordinates, and a confidence score. The post-processing stage applies non-maximum suppression (NMS) to eliminate redundant overlapping detections. Each surviving detection is then passed to the severity analysis module, which computes the composite severity score and assigns a severity level. Finally, the recommendation engine retrieves the relevant corrective actions for each detected defect type and severity level, and all results are rendered on the monitoring dashboard in real time. The entire pipeline completes in under 25 ms per image on the target CPU hardware, satisfying real-time inspection requirements.

V. EXPERIMENTAL RESULTS

A. Experimental Configuration

All experiments are conducted on a system running Windows OS with an Intel Core processor at 2.6 GHz, 4 GB RAM, and a 320 GB hard disk. The model is implemented in Python using the Ultralytics YOLOv11 library, PyTorch framework, TensorFlow for auxiliary modules, and Keras for comparative CNN baseline training. The development environment uses PyCharm IDE with HTML and CSS for the frontend dashboard. Training experiments requiring GPU acceleration use Google Colab T4 GPU instances. Inference latency measurements are conducted on the CPU-only configuration to reflect realistic deployment constraints where dedicated GPU hardware may not be available. The AdamW optimizer was employed with a batch size of 16 and a training schedule of 300 epochs.

B. Evaluation Metrics

Detection performance is evaluated using Precision (P), Recall (R), F1-Score, and mean Average Precision at IoU threshold 0.50 (mAP50). Precision ($P = TP/(TP+FP)$) measures the reliability of positive detections, penalizing false alarms that would require unnecessary manual re-inspection of conforming boards. Recall ($R = TP/(TP+FN)$) measures the completeness of defect detection, penalizing missed detections that allow defective boards to pass inspection undetected. F1-Score ($F1 = 2PR/(P+R)$) provides the harmonic mean of precision and recall as a balanced overall performance indicator. mAP50 computes the mean area under the precision-recall curve across all defect classes at an IoU threshold of 0.50, capturing both localization and classification accuracy in a single scalar metric.

C. Per-Class Detection Performance

Table 1 presents per-class detection performance of the proposed YOLOv11 model on the held-out test set. The model achieves strong performance across all nine defect categories. Short circuits and missing holes achieve the highest precision (0.95 and 0.94 respectively) due to their visually distinct appearance and consistent geometry. Open circuits and spur formations show slightly lower recall (0.87 and 0.86) due to visual similarity to background trace patterns in some board configurations. Scratches achieve the lowest

precision (0.88) because scratch-like artifacts can appear in non-defective board areas. Overall, the model achieves a mean mAP50 of 0.931 across all classes, confirming robust multi-class detection capability.

Table 1: Per-Class Detection Performance of Proposed YOLOv11 Model

Defect Class	Precision	Recall	F1	mAP50
Missing Hole	0.94	0.93	0.935	0.962
Mouse Bite	0.91	0.89	0.900	0.931
Open Circuit	0.90	0.87	0.885	0.914
Short Circuit	0.95	0.92	0.935	0.958
Spurious Cu.	0.92	0.90	0.910	0.938
Spur	0.89	0.86	0.875	0.903
Crack	0.90	0.88	0.890	0.921
Scratch	0.88	0.87	0.875	0.906
Missing Comp.	0.93	0.91	0.920	0.947
Mean (All)	**0.913**	**0.892**	**0.903**	**0.931**

D. Comparative Performance Analysis

Table 2 presents the comparative performance of the proposed system against baseline methods evaluated on the same held-out test set. The proposed YOLOv11 system achieves the highest performance across all metrics, with 92% accuracy representing a 9-percentage-point improvement over the best existing automated method (semi-supervised learning at 83%) and a 32-point improvement over manual inspection (60%). The precision of 91% indicates over 9 in 10 flagged defects are genuine, substantially reducing false alarm-driven re-inspection overhead. The recall of 90% means 9 in 10 actual defects are detected, representing a meaningful reduction in defect escape rate.

Table 2: Comparative Performance Against Baseline Methods

Method	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Manual Inspection	60	55	58	56.4
Template Matching	68	63	65	63.9
CNN (Scratch)	78	75	77	75.9
Semi-Supervised	83	80	81	80.5
Proposed YOLOv11	**92**	**91**	**90**	**90.5**

E. Inference Speed and Computational Efficiency

The proposed YOLOv11n model achieves an average inference latency of 22.3 ms per image on the Intel Core CPU test configuration, corresponding to approximately 44 inspections per second — well within the real-time inspection requirements of automated PCB assembly lines. The model occupies 5.4 MB of disk storage and requires approximately 95 MB of RAM during inference, making it deployable on standard industrial PCs without specialized hardware requirements. Comparative benchmarking confirms that YOLOv11n achieves 14% lower latency than YOLOv8n on the same hardware while maintaining superior detection accuracy, validating the architectural improvements introduced in YOLOv11. The severity analysis module correctly classifies High, Medium, and Low severity defects with 94%, 88%, and 91% accuracy respectively. The corrective recommendation engine generates contextually appropriate recommendations for 91% of High-severity defect instances.

F. Ablation Study

To assess the individual contribution of each system component, an ablation study was performed by selectively disabling modules. Removing data augmentation (mosaic, copy-paste) reduced mAP50 by 4.2 percentage points, confirming its critical role in improving model generalization across diverse board designs. Replacing YOLOv11n with YOLOv8n decreased F1 by 2.1 points while increasing inference latency by 14%, validating the architectural improvements in YOLOv11. Disabling transfer learning from COCO-pretrained weights and training from random initialization reduced mAP50 by 8.7 points, underscoring the importance of pre-trained feature representations even when fine-tuning on a specialized domain dataset. These results confirm that each design decision contributes meaningfully to the final system performance.

G. Hardware and Software Specifications

The hardware configuration used for this project consists of: Processor — Intel Core processor at 2.6 GHz; RAM — 4 GB; Hard Disk — 320 GB; Compact Disk — 650 MB; Keyboard — Standard keyboard; Monitor — 15-inch color monitor. The software configuration consists of: Operating System — Windows OS; Front End — HTML, CSS; Back End — Python; IDE — PyCharm; Libraries — TensorFlow, Keras, PyTorch, Ultralytics YOLOv11. These specifications confirm that the proposed system is deployable without requiring specialized or high-cost hardware infrastructure. The lightweight YOLOv11n model with 2.6M parameters and 5.4 MB disk footprint is specifically suited to such resource-constrained environments, making it practical

for real-world industrial deployment in manufacturing facilities of varying scales.

H. Advantages of the Proposed System

The proposed YOLOv11-based PCB defect detection system offers several distinct advantages over existing approaches. First, it provides fast and real-time detection of PCB defects at 22.3 ms per image, significantly faster than manual inspection and classical AOI systems. Second, it reduces inspection errors and inaccuracies compared to manual methods by eliminating human cognitive fatigue and inter-operator variability. Third, the deep learning model can identify small or subtle defects such as mouse bites, spurs, and hairline cracks that are challenging to detect through visual or classical automated inspection. Fourth, the system classifies detected defects into nine categories and evaluates their severity level for comprehensive quality analysis. Fifth, the corrective recommendation engine provides actionable suggestions for manufacturing process adjustments, improving overall PCB manufacturing quality and reducing the root causes of recurring defects. These combined advantages position the proposed system as a comprehensive quality management solution rather than a simple defect detection tool.

The integration of the corrective recommendation engine into the PCB inspection workflow represents a significant advancement over conventional defect detection systems. By linking detected defect patterns directly to their manufacturing root causes, the system enables proactive quality management rather than purely reactive defect identification. This closed-loop approach reduces the mean time between defect recurrence events and supports continuous process improvement across production shifts.

The web-based dashboard further enables quality managers to monitor defect trends over time, identify systematic process drifts before they cause significant yield losses, and generate detailed inspection reports for compliance documentation and customer quality audits.

VI. CONCLUSION

This research presents a comprehensive YOLOv11-based automated PCB defect detection and quality assessment system designed for deployment in smart manufacturing environments. By integrating five modules — dataset preparation, multi-stage data preprocessing and augmentation, YOLOv11 model fine-tuning via transfer learning, real-time multi-class defect detection, and a post-detection quality assessment pipeline — the proposed framework addresses the complete PCB quality control workflow from image acquisition to actionable quality management outputs. The system accepts PCB input as either a static image or a real-time camera capture, processes it through the YOLOv11 model for simultaneous defect localization and classification, and presents the results along with severity ratings and corrective recommendations through a web-based monitoring dashboard.

The YOLOv11n backbone provides a favorable balance between detection accuracy and computational efficiency. The anchor-free detection head, C3k2 backbone modules, and C2PSA neck attention mechanisms collectively enable precise localization and classification of nine PCB defect types with a mean average precision of 93.1% at IoU threshold 0.50. Real-time inference at 22.3 ms per image on standard CPU hardware satisfies automated assembly line throughput requirements without specialized GPU infrastructure. The proposed system achieves 92% detection accuracy, precision of 91%, recall of 90%, and F1-score of 90.5%, substantially outperforming all evaluated baseline methods including manual inspection (60%), template-matching AOI (68%), CNN-based classification (78%), and semi-supervised learning (83%).

The severity analysis module achieves 94% accuracy for High-severity defect classification, and the corrective recommendation engine provides appropriate manufacturing process guidance for 91% of High-severity defect instances, demonstrating practical value beyond raw defect detection. The modular architecture allows individual components to be upgraded independently, new defect classes to be incorporated by extending the training dataset, and the recommendation knowledge base to be expanded with additional expert-derived defect-process mappings. The ablation study confirms that each design decision — data augmentation, YOLOv11 architecture, and transfer learning — contributes meaningfully to the overall system performance.

Future work will focus on five directions: (i) extending the system to process real-time video streams from inline industrial cameras for continuous board monitoring during conveyor transport; (ii) expanding the training dataset with surface-mount technology (SMT) assembly defects including tombstoning, component polarity reversal, insufficient solder, and solder balls identified through X-ray inspection; (iii) exploring federated learning for multi-facility collaborative model improvement without sharing proprietary board design data; (iv) enhancing the severity analysis module with a physics-informed failure mode model quantifying the actual electrical impact of detected defects on circuit functionality; and (v) integrating Grad-CAM visual

explainability for operator trust calibration and systematic error identification during model validation. The proposed system represents a substantive step toward fully autonomous PCB quality control in Industry 4.0 smart manufacturing, contributing to reduced defect escape rates, improved product reliability, and enhanced operational efficiency across the electronics manufacturing industry.

VII. ACKNOWLEDGMENT

The authors express sincere gratitude to Mrs. Banupriya P, B.E., M.E., Assistant Professor, Department of Computer Science and Engineering, Mahendra Institute of Engineering and Technology, Namakkal, for her invaluable guidance, continuous support, and insightful technical direction throughout this project. The authors also thank the Department of Computer Science and Engineering and the institution for providing access to computational resources and research facilities essential to the completion of this work.

REFERENCES

[1] Ling, Qin, and Nor Ashidi Mat Isa. "Printed circuit board defect detection methods based on image processing, machine learning and deep learning: A survey." IEEE Access 11 (2023): 15921–15944.

[2] Chen, I-Chun, Rey-Chue Hwang, and Huang-Chu Huang. "PCB defect detection based on deep learning algorithm." Processes 11.3 (2023): 775.

[3] Nguyen, Van-Truong, and Huy-Anh Bui. "A real-time defect detection in printed circuit boards applying deep learning." EUREKA: Physics and Engineering 2 (2022): 143–153.

[4] Pham, Thi Tram Anh, et al. "Defect detection in printed circuit boards using semi-supervised learning." Sensors 23.6 (2023): 3246.

[5] Zheng, Jianfeng, et al. "Printed circuit boards defect detection method based on improved fully convolutional networks." IEEE Access 10 (2022): 109908–109918.

[6] Long, Y., Li, Z., Cai, Y., Zhang, R., and Shen, K. "PCB defect detection algorithm based on improved YOLOv8." Academic Journal of Science and Technology 7.3 (2023): 297–304.

[7] Ong, Sze Kit, Chee Keong Tan, Vishnu Monn Baskaran, Boon Keat Puah, and Kok Hoong Liew. "Enhancing industrial PCB and PCBA defect detection: An efficient and accurate SEConv-YOLO approach." IEEE Access 13 (2025): 148917–148935.

[8] Jocher, G., Qiu, J., and Chaurasia, A. "Ultralytics YOLO (YOLOv11)." Version 11.0.0, Ultralytics, 2024. [Online]. Available: https://github.com/ultralytics/ultralytics

[9] Huang, W., Wei, P., Zhang, M., and Liu, H. "HRIPCB: A challenging dataset for PCB defects detection and classification." Journal of Engineering 2020.13 (2020): 303–309.

[10] Wang, Aotian, et al. "YOLOv10: Real-time end-to-end object detection." Advances in Neural Information Processing Systems 37 (2024):