Selective Re-Inference for PCB Defect Detection viaImage-Specific Template-Difference Density

Yuning Xie[1*], Feng Xie[1], Chao Li[1], Zexi Deng[1]

[1]Department of Artificial Intelligence, Guangdong Industry Polytechnic University, Guangzhou 510300, China

[*]Corresponding author: 512206062@qq.com

Abstract: Small and subtle printed circuit board (PCB) defects can be missed by direct full-image detection, whereas dense slicing improves local resolution at the cost of additional detector views and more false positives. We propose a selective re-inference framework that uses image-specific template-difference density to allocate a single additional local detector pass. Given an inspected image and its paired defect-free template, the method computes a template-difference map, aggregates the difference responses over candidate 512×512 windows, selects the highest-density region of interest (ROI), applies the same detector to this crop, and fuses the remapped local detections with the full-image detections. On a locked pair-level split of the DeepPCB dataset, the proposed method improves F1 from 0.9453 to 0.9535 relative to direct full-image detection. Under the same two-view detector-inference budget, it achieves a higher F1 score and fewer false positives than random, grid, and mismatched-difference ROI controls. These results indicate that the gain is not attributable solely to adding a local crop, but to placing that crop using the image-specific paired template difference. The method is best interpreted as a recall-oriented selective re-inference strategy for paired PCB inspection.

Keywords: PCB defect detection; template-guided inspection; selective re-inference; template-difference density; ROI selection

I Introduction

Automated printed circuit board (PCB) defect detection is an important task in industrial visual inspection because local manufacturing defects can affect circuit reliability and product yield [1], [2]. Modern object detectors enable efficient full-image PCB inspection, but this setting remains difficult when defects are small, weakly contrasted, or spatially concentrated. Since the full board must be resized to a fixed detector input resolution, small defects may occupy only a few pixels after resizing and can therefore be missed by direct full-image detection.

Local re-inference is a natural way to improve the visibility of small defects. Dense slicing increases local resolution by applying the detector to multiple cropped views, and it has been widely used for small-object detection [3], [4]. However, dense slicing also increases the number of detector passes and can introduce duplicated or spurious detections. This motivates a more constrained setting for industrial inspection: allocating a single additional local detector pass to the region where it is most likely to recover missed defects without substantially increasing false positives (FP).

This paper studies selective local re-inference in the paired PCB inspection setting. In many PCB datasets and production workflows, each inspected image is associated with a corresponding defect-free template image. The spatial difference between the inspected image and its template provides an image-specific cue about where potential defects may appear. We do not use the difference map as a final defect mask. Instead, we use accumulated template-difference density as an inference-time prior for selecting one local region of interest (ROI), and then let the same detector perform local re-inference on that region.

We propose an image-specific template-difference-density guided selective re-inference framework. The detector is first applied to the full inspected image. A template-difference density map is then computed between the inspected image and its paired template, and the highest-density 512×512 ROI is selected for one additional local detector pass. The resulting local detections are mapped back to the original image coordinates and fused with the full-image detections. Since the detector checkpoint, inference thresholds, coordinate remapping, and fusion rule are kept unchanged across ROI methods, the comparison isolates the effect of ROI placement.

Overall, this work makes a focused contribution to ROI-guided PCB defect detection under a fixed detector-inference budget. It introduces a simple, detector-agnostic strategy that uses image-specific template-difference density to decide where to allocate one additional local detector pass. The method is evaluated with random ROI, grid ROI, mismatched-difference ROI, low-difference ROI, and extent-center ROI controls, showing that density-based ROI placement provides a more favorable trade-off between missed-defect recovery and false-positive control than alternative one-crop ROI placement rules.

II Related Work

A. PCB Defect Detection

PCB defect detection has been widely studied with both classical inspection pipelines and deep-learning-based detectors. Early detector-based studies demonstrated that one-stage detectors can be applied to PCB defect localization [1]. More recent work has mainly focused on improving the detector architecture itself. Representative directions include multi-scale feature fusion and attention for tiny defects [2], lightweight PCB-specific detectors [5], [12], and improved YOLO variants for PCB or PCBA inspection [6]-[11]. These methods are valuable because they strengthen the feature representation or detection head. In contrast, this work does not propose a new detector architecture. We keep the detector fixed and study an inference-stage question: where should one additional local detector pass be allocated under a fixed extra local-inference budget?

B. Template-Based Visual Inspection

Template-based visual inspection uses a reference or defect-free image to reveal abnormal regions in an inspected image. Traditional template-based methods often rely on subtraction, thresholding, registration, morphology, or connected-component analysis to directly identify suspicious regions [13], [14]. Such approaches can be sensitive to alignment error, illumination variation, and non-defect structural differences. Our method uses template difference in a different role. Instead of treating the difference map as the final defect prediction, we use its accumulated local density only to select a region for detector re-inference. The final detection results are still produced by the same object detector and the same fusion rule.

C. Cropped Inference and ROI Allocation

Cropped or sliced inference improves the relative scale of small objects by applying a detector to local image regions [3], [4], [15]. This strategy is useful when small defects lose visual detail after full-image resizing. However, dense slicing requires multiple detector views per image and can increase duplicated or spurious detections. ROI-guided inference aims to allocate computation more selectively, either through learned proposals, attention mechanisms, or handcrafted priors. In the paired PCB setting, the inspected image and its defect-free template provide a natural image-specific cue for ROI placement. The proposed method follows this selective-inference view: it uses the paired template-difference density to choose one local crop, rather than applying dense slicing or learning a new proposal module.

III Method

A. Problem Formulation

Given an inspected PCB image I and its corresponding defect-free template image T, the goal is to detect defect instances in I. A fixed detector f_θ first produces full-image detections.

$$\mathrm{B}_{full} = f_\theta(I) \tag{1}$$

Selective re-inference chooses one local window W for an additional detector pass. The objective is to recover defects missed by B_{full} while limiting the increase in false positives. In this work, the window is selected from the template-difference density rather than from ground-truth annotations or detector feedback.

B. Overview of Selective Re-Inference

Figure 1 illustrates the proposed framework with two connected modules: template-difference-density ROI generation and selective re-inference with fusion. For readability, the figure uses compact notation; the formal definitions are provided in the following subsections.

The upper module shows how the additional local crop is selected. Given an inspected PCB image I and its paired defect-free template T, the method first computes a pixel-level template-difference map $D(x,y)$. The processed difference response $D'(x,y)$ is then accumulated over candidate windows to obtain a window-level density score $S(W)$. The highest-scoring window W^* is selected as the single 512×512 ROI and cropped from the inspected image. This ROI-generation branch determines only where the additional local detector pass is allocated; it does not use ground-truth annotations and does not directly output defect predictions.

The lower module shows how the selected crop is used for re-inference. The full inspected image I and the selected ROI crop are fed independently into the same fixed detector, yielding full-image detections B_{full} and ROI detections B_{roi}. The ROI detections are remapped back to the original image coordinates by Φ_{W^*} and fused with the full-image detections using the same NMS-based fusion rule. Therefore, the detector architecture, detection thresholds, coordinate remapping, and fusion rule are kept unchanged; only the source and placement of the additional local crop differ among the compared ROI methods.

C. Template-Difference Density Map

The pixel-level template-difference map is computed from the inspected image and its paired template.

$$D(x,y) = \left| I(x,y) - T(x,y) \right| \tag{2}$$

In our implementation, the difference response is smoothed and thresholded to reduce noise from image alignment, illumination variation, or non-defect texture changes.

$$D'(x,y) = \begin{cases} D_s(x,y), & D_s(x,y) \ge \tau_d \\ 0, & D_s(x,y) < \tau_d \end{cases} \tag{3}$$

where D_s denotes the smoothed difference response and τ_d is the difference threshold. This processing step suppresses scattered low-intensity differences before density scoring. Importantly, D' is used only to choose the ROI. It is not used as a defect prediction mask and does not use ground-truth annotations from the locked test set. The smoothing and thresholding parameters were fixed before locked-test evaluation.

The smoothing kernel size and threshold τ_d were selected on the validation split and kept fixed for all locked-test evaluations.

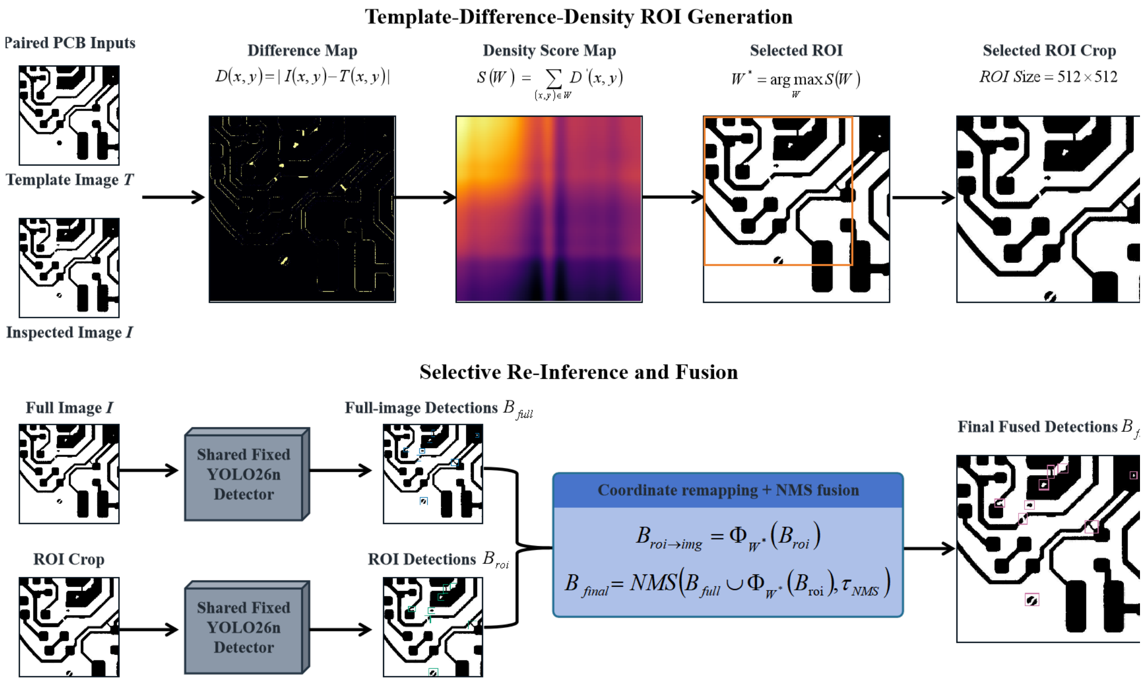

Figure 1. Overview of the proposed template-difference-density guided selective re-inference framework. The upper module generates a single local ROI from paired PCB inputs by computing a template-difference map, accumulating the processed difference response over candidate windows, and selecting the highest-density 512×512 crop. The lower module applies the same fixed detector to the full image and the selected ROI crop. ROI detections are remapped back to the original image coordinates and fused with full-image detections using the same NMS rule.

D. Density-Based ROI Selection

For each candidate window W, the difference-density score is defined as the accumulated processed difference response inside the window.

$$S(W)=\sum_{(x,y)\in W} D'(x,y) \tag{4}$$

The selected ROI is the window with the highest score.

$$W^* = \underset{\in \mathrm{W}}{\arg\max}\, S(W) \tag{5}$$

where W denotes the set of candidate 512×512 windows. The ROI size is fixed to 512×512 in the experiments. Density-based scoring differs from using the center of a thresholded difference extent. A large difference extent may be caused by sparse responses or alignment artifacts, whereas a high-density window favors concentrated local evidence. This design is intended to select a region that is more suitable for detector re-inference.

E. Local Re-Inference and Fusion

The selected ROI W^* is cropped from the inspected image and passed through the same detector.

$$\mathrm{B}_{\mathrm{roi}} = f_\theta(\mathrm{Crop}(I, W^*)) \tag{6}$$

The ROI detections are then transformed back to the original image coordinate system.

$$\mathrm{B}_{\mathrm{roi}\rightarrow\mathrm{img}} = \Phi_{W^*}(\mathrm{B}_{\mathrm{roi}}) \tag{7}$$

For a crop without scale change, Φ_{W^*} adds the ROI offset to the local box coordinates. The full-image detections and mapped ROI detections are then fused by class-wise non-maximum suppression (NMS).

$$\mathrm{B}_{\mathrm{final}} = \mathrm{NMS}\left(\mathrm{B}_{\mathrm{full}} \cup \Phi_{W^*}(\mathrm{B}_{\mathrm{roi}}), \tau_{\mathrm{NMS}}\right) \tag{8}$$

The same coordinate remapping and fusion rule are used for random ROI, grid ROI, mismatched-difference ROI, low-difference ROI, extent-center ROI, and the proposed difference-density ROI. Therefore, the performance difference among these methods reflects ROI selection rather than detector architecture or post-processing changes.

F. Detector-Inference Budget Definition

This paper evaluates ROI placement under a fixed detector-inference budget rather than under wall-clock runtime. Direct detection uses one full-image detector pass. Each same-budget ROI method uses one full-image detector pass and one additional local ROI detector pass. Dense standard slicing uses multiple local detector passes and is therefore treated as a high-budget reference rather than a same-budget ROI-selection baseline. Since the proposed method also requires template differencing and density scoring for ROI generation, we do not claim a wall-clock speed advantage; the main comparison is controlled by the number of detector passes and local crops.

IV Experiments

A. Dataset and Evaluation Protocol

We evaluated the proposed method on the DeepPCB dataset. Each sample contains an inspected PCB image and a corresponding template image, which makes the dataset suitable for evaluating template-guided local re-inference. To avoid leakage between paired samples, the dataset was split at the image-pair level into training, validation, and locked-test subsets. Table I summarizes the resulting split and defect-instance distribution.

Table I. Pair-level dataset split and defect-instance distribution

Defect Class	Train	Val	Test
Open	1351	300	291
Short	1056	225	225
Mousebite	1372	291	302
Spur	1131	247	247
Spurious_copper	1038	216	220
Pin_hole	1051	228	222
Total instances	6999	1507	1507
Image pairs	1050	225	225

The validation set was used for checkpoint selection and for fixing the operating protocol. The locked test set was used only once for final reporting and was not used for threshold tuning, ROI-parameter selection, or model selection.

B. Implementation Details and Metrics

YOLO26n was used as the fixed detector in all experiments. This detector was selected as a lightweight and reproducible baseline for isolating the effect of ROI selection, rather than as a newly optimized detector architecture. All compared methods used the same trained detector checkpoint. Table II summarizes the key training, inference, and evaluation settings.

Table II. Key training, inference, and evaluation settings of the fixed YOLO26n detector

Item	Setting
Detector	YOLO26n initialized from pretrained *yolo26n.pt*
Input size	Full-image 640×640; ROI crop 512×512
Training	100 epochs; batch size 16; early stopping patience 50
Optimization	Ultralytics auto optimizer; learning rate 0.01; momentum 0.937; weight decay 0.0005

续表2

Item	Setting
Checkpoint	*best.pt* selected by validation performance
ROI budget	1 full-image pass + 1 local ROI pass
Inference / fusion	confidence 0.001; internal NMS IoU 0.45; operating confidence 0.30; fusion IoU 0.55

Operating-point metrics, including precision, recall, F1, FP, FN, and recovered direct-FN, were reported at a confidence threshold of 0.30. We also report mAP@0.50 and COCO-style mAP@[0.50:0.95] as localization-sensitive secondary metrics. A direct-FN instance is a ground-truth defect that is not matched by any full-image detection under the evaluation IoU threshold of 0.50. It is counted as recovered if the final fused detection set contains a matched detection for the same ground-truth instance. Random ROI results are reported as mean and standard deviation over 50 predefined seeds. Mismatched-difference ROI results are reported as mean and standard deviation over 50 random permutations, where each inspected image uses a difference map from another image rather than its own paired template.

C. Direct Baseline and Same-Budget Control Analysis

Table III reports the direct full-image baseline and the ROI-based methods. Direct detection uses one full-image detector pass and is included as the one-view baseline. All ROI-based methods use the same two-view detector-inference budget: one full-image detector pass and one additional 512×512 local detector pass.

Direct detection produced 65 false negatives on the locked test set. Adding one random ROI reduced false negatives but increased false positives substantially. Grid ROI showed a similar pattern, suggesting that one additional local crop can improve recall but may also introduce additional false detections. In contrast, the proposed difference-density ROI reduced FN from 65 to 40 while keeping FP nearly unchanged relative to direct detection.

The main comparison supports the value of image-specific ROI placement. Compared with direct detection, difference-density ROI recovered 25 of 65 direct false negatives and improved F1 from 0.9453 to 0.9535, while FP increased only from 102 to 103. Compared with random ROI, the proposed method achieved a higher F1 score and reduced FP by about 20 detections. Compared with mismatched-difference ROI, the proposed method also achieved higher F1 and lower FP. This result shows that the improvement is not explained by the form of a difference-map-derived ROI alone; the correspondence between the inspected image and its own template-difference map is important.

The stability controls further support this interpretation. Across 50 random ROI seeds and 50 mismatched-difference permutations, no control trial exceeded the true difference-density ROI in F1 or achieved a lower FP count. These win-count results reduce the possibility that the proposed method only benefits from a favorable random crop or a favorable mismatched difference map.

The proposed method does not improve every metric. Its mAP@[0.50:0.95] is lower than that of direct detection and grid ROI. This result is consistent with the recall-oriented purpose of the method: local re-inference can recover missed defects and improve the trade-off between missed-defect recovery and false-positive control, but the additional local boxes and coordinate remapping do not necessarily improve strict high-IoU localization metrics for very small defects.

D. Difference-ROI Construction Ablation

Table IV examines whether the high-density paired difference region is specifically useful. Low-difference ROI selects a low-density region and serves as a reverse-prior control. Difference extent-center ROI uses the center of the thresholded difference extent and tests whether a simpler difference-based crop rule is sufficient. Difference-density ROI achieves the best F1, the lowest FP, the lowest FN, and the highest recovered direct-FN count among these difference-guided variants.

These results show that simply using any difference-based crop is not sufficient. The low-difference region and the extent-center region both recovered some missed detections, but they introduced more false positives and achieved lower F1 scores than the proposed density-based rule. The accumulated local difference response therefore provides a more useful ROI selection cue than either a reverse difference prior or the center of the global difference extent.

Table III. Direct baseline and same-budget comparison of ROI placement strategies

Method	Direct	Direct + Random ROI	Direct + Grid ROI	Direct + Mismatched-difference ROI	Direct + Difference-density ROI
ROI source	-	Random	Fixed grid	Wrong difference map	Template difference density
Trials	-	50 seeds	-	50 permutations	-
Precision	0.9339	0.9224 ± 0.0040	0.9218	0.9228 ± 0.0037	**0.9344**
Recall	0.9569	0.9726 ± 0.0018	0.9695	0.9717 ± 0.0014	**0.9735**
F1	0.9453	0.9468 ± 0.0024	0.945	0.9467 ± 0.0021	**0.9535**
FP	**102**	123.38 ± 6.80	124	122.48 ± 6.35	103
FN	65	41.22 ± 2.66	46	42.58 ± 2.06	**40**
mAP@0.50	0.9784	0.9826 ± 0.0011	0.9826	0.9825 ± 0.0009	**0.9834**
mAP@[0.50:0.95]	0.7786	0.7767 ± 0.0018	0.7792	0.7773 ± 0.0019	0.7719
Direct-FN recovered	0/65 (0.00%)	23.90 ± 2.64/65 (36.77%)	20/65 (30.77%)	22.74 ± 1.96/65 (34.98%)	**25/65 (38.46%)**

Table IV. Ablation of difference-guided ROI selection rules

Method	Direct + Low-difference ROI	Direct + Difference extent-center ROI	Direct + Difference-density ROI
Tested factor	Uses low-density difference region	Uses thresholded-difference extent center	Uses high-density paired difference region
Precision	0.9241	0.9184	0.9344
Recall	0.9688	0.9708	0.9735
F1	0.9459	0.9439	0.9535
FP	120	130	103
FN	47	44	40
Direct-FN recovered	19/65 (29.23%)	22/65 (33.85%)	25/65 (38.46%)

E. Qualitative Results

Figure 2 provides a qualitative example of the proposed selective re-inference process. Direct detection misses a small defect in the full-image view. The difference-density map highlights the local region with strong paired template response, and the proposed ROI covers the missed-defect region. Under the same one-crop budget, grid and mismatched controls do not provide the same image-specific placement. After local re-inference and fusion, the previously missed defect is recovered.

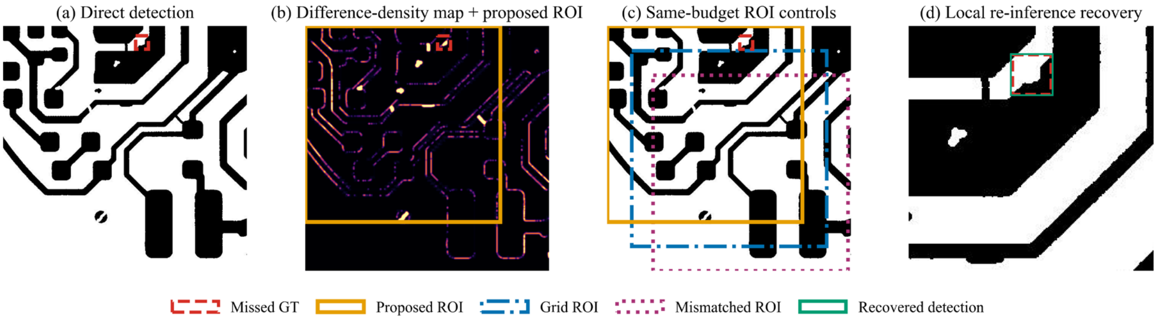

Figure 2. Qualitative comparison of direct detection, template-difference-density ROI selection, same-budget ROI controls, and local re-inference recovery. Ground-truth boxes are shown only for visualization and are not used during ROI generation, detector inference, or fusion.

F. Detector-Budget Discussion

Dense standard slicing was evaluated as a high-budget reference rather than as a same-budget ROI-selection baseline. In the current setting, dense slicing used five detector views per image. It achieved higher recall than the proposed single-ROI method, but it also produced 211 FP and an F1 score of 0.9232. In contrast, the proposed method used only one additional local detector pass, produced 103 FP, and achieved an F1 score of 0.9535. This comparison indicates that dense slicing can recover more defects by spending more local views, but its false-positive cost is substantially higher.

V Discussion

The experiments show that image-specific ROI placement matters for selective local re-inference. Under the same detector-inference budget, all ROI methods add one extra local crop, but the proposed difference-density ROI achieves a more favorable trade-off between missed-defect recovery and false-positive control than random, grid, and mismatched-difference controls. This result supports the central claim that the useful information is not merely the existence of one more crop, but the placement of that crop at a high-density template-difference region corresponding to the current inspected image.

The mismatched-difference control is important for interpreting the result. A mismatched difference map preserves the general form of a difference-guided ROI, but it breaks the sample-specific correspondence between the inspected image and its own template. Its performance is close to random ROI and clearly below true difference-density ROI. This indicates that the paired template relationship provides useful spatial information for local re-inference.

The density-rule ablation clarifies the ROI design choice. A thresholded difference extent can be affected by sparse responses, alignment noise, or large low-density regions. Selecting the highest-density window produces a more concentrated local crop and leads to a better final detection trade-off. This does not mean that geometric coverage alone explains the gain; rather, the selected crop must both cover relevant defect evidence and interact favorably with the detector and fusion step.

The method has several limitations. First, it assumes that a corresponding defect-free template image is available for each inspected image. Second, the quality of image alignment can affect the reliability of the difference map. Third, the current method selects only one ROI, which may be insufficient for images containing spatially separated defects. Fourth, the method targets recall-oriented defect discovery and false-positive control, not strict high-IoU localization; this explains why mAP@[0.50:0.95] is not the main advantage. Future work can study adaptive multi-ROI selection, more robust template registration, and joint optimization of ROI scoring and detection fusion.

VI Conclusion

This paper presented a template-difference-density guided selective re-inference strategy for PCB defect detection. The method uses the paired template image to compute an image-specific difference-density map, selects one high-density 512×512 ROI, performs one additional local detector pass, and fuses local detections with full-image detections. On the locked pair-level test split, the proposed ROI achieved a better F1 score and a more favorable trade-off between missed-defect recovery and false-positive control than random ROI, grid ROI, mismatched-difference ROI, low-difference ROI, and extent-center ROI controls under the same detector-inference budget. These findings indicate that template-difference density is an effective ROI prior for deciding where to allocate one additional local detector pass in paired PCB inspection. The method is not intended to replace dense slicing when maximum recall is required; instead, it provides a detector-budget-constrained selective re-inference alternative when only one additional local crop is allowed.

Funding

This work was supported by the 2025 Teaching Reform Project of the Guangdong Vocational Education Steering Committee for Electronic Information and Communication and the 2025 Teaching Reform Project of Guangdong Industry Polytechnic University (Grant No. JG202509).

References

[1] V. A. Adibhatla et al., "Defect detection in printed circuit boards using You-Only-Look-Once convolutional neural networks," Electronics, vol. 9, no. 9, p. 1547, 2020, doi: 10.3390/electronics9091547.

[2] K. C. Fung, K.-W. Xue, C.-M. Lai, K.-H. Lin, and K.-M. Lam, "Improving PCB defect detection using selective feature attention and pixel shuffle pyramid," Results Eng., vol. 21, Art. no. 101992, 2024, doi: 10.1016/j.rineng.2024.101992.

[3] F. C. Akyon, S. O. Altinuc, and A. Temizel, "Slicing aided hyper inference and fine-tuning for small object detection," in 2022 IEEE International Conference on Image Processing (ICIP), 2022, pp. 966-970, doi: 10.1109/ICIP46576.2022.9897990.

[4] F. O. Unel, B. O. Ozkalayci, and C. Cigla, "The power of tiling for small object detection," in 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 2019, doi: 10.1109/CVPRW.2019.00084.

[5] Y. Jiang, M. Cai, and D. Zhang, "Lightweight network DCR-YOLO for surface defect detection on printed circuit boards," Sensors, vol. 23, no. 17, p. 7310, 2023, doi: 10.3390/s23177310.

[6] M. Shen et al., "Defect detection of printed circuit board assembly based on YOLOv5," Sci. Rep., vol. 14, Art. no. 19287, 2024, doi: 10.1038/s41598-024-70176-1.

[7] Y. Tang, R. Liu, and S. Wang, "YOLO-SUMAS: improved printed circuit board defect detection and identification research based on YOLOv8," Micromachines, vol. 16, no. 5, p. 509, 2025, doi: 10.3390/mi16050509.

[8] P. Selvam et al., "YOLO-DefXpert: an advanced defect detection on PCB surfaces using improved

YOLOv11 algorithm," IEEE Access, vol. 13, pp. 143085-143101, 2025, doi: 10.1109/ACCESS.2025.3595048.

[9] Y. Zhao and Z. Jiang, "YOLO-WWBi: an optimized YOLO11 algorithm for PCB defect detection," IEEE Access, vol. 13, pp. 74288-74297, 2025, doi: 10.1109/ACCESS.2025.3564734.

[10] Z. Wei, F. Yang, K. Zhong, and L. Yao, "PCB-YOLO: enhancing PCB surface defect detection with coordinate attention and multi-scale feature fusion," PLOS One, vol. 20, no. 6, Art. no. e0323684, 2025, doi: 10.1371/journal.pone.0323684.

[11] G. Xiao, S. Hou, and H. Zhou, "PCB defect detection algorithm based on CDI-YOLO," Sci. Rep., vol. 14, no. 1, Art. no. 7351, 2024, doi: 10.1038/s41598-024-57491-3.

[12] C. Mo, Z. Hu, J. Wang, and X. Xiao, "SGT-YOLO: a lightweight method for PCB defect detection," IEEE Trans. Instrum. Meas., vol. 74, pp. 1-11, 2025, doi: 10.1109/TIM.2025.3563011.

[13] Z. Zeng, B. Liu, J. Fu, and H. Chao, "Reference-based defect detection network," IEEE Trans. Image Process., vol. 30, pp. 6637-6647, 2021, doi: 10.1109/TIP.2021.3096067.

[14] X. Liu, Y. Li, Y. Guo, and L. Zhou, "Printing defect detection based on scale-adaptive template matching and image alignment," Sensors, vol. 23, no. 9, p. 4414, 2023, doi: 10.3390/s23094414.

[15] H. Zhang, C. Hao, W. Song, B. Jiang, and B. Li, "Adaptive slicing-aided hyper inference for small object detection in high-resolution remote sensing images," Remote Sens., vol. 15, no. 5, p. 1249, 2023, doi: 10.3390/rs15051249.