A Lightweight Framework with Strip Perception and Adaptive Pooling for Fabric Defect Detection

Salman Shehzad, Yuan-Gen Wang, Chunsheng Yang*
School of Artificial Intelligence, Guangzhou University
Guangzhou, China

Imran Hussain
School of Electrical and Computer Engineering
South China University of Technology
Guangzhou, China
**Corresponding Author*

Abstract: Fabric surface defects reduce textile quality, so accurate real-time detection is a key to industrial quality control. Current lightweight models such as YOLOv7-Tinier struggle to detect elongated/subtle defects (such as cutting, small holes) while staying efficient for on-site use—limiting factory deployment. To address this, we propose YOLOv7-Tinier-SPAF (Strip Perception and Adaptive Pooling), an improved lightweight framework for textile inspection. In the proposed framework, we add three innovations to YOLOv7-Tinier: (1) Strip Perception Module (SPM), which uses asymmetric kernels to capture long, narrow defects; (2) Squeeze-and-Excitation Spatial Pyramid Pooling-Fast (SESPPF), which combines channel attention and multi-scale pooling to refine features; and (3) Focal-Enhanced Complete Intersection over Union (FECIoU) Loss, which weights hard-to-detect cases to improve bounding box accuracy and defect localization. Evaluated on a textile defect dataset, the experimental results show that YOLOv7-Tinier-SPAF performed better than the Yolov7-Tiny baseline both in performance and speed. Ablation studies confirm that each module works effectively: SPM improves the detection of elongated defects by 18.6%, SE-SPPF improves features by 11.2%, and FECIoU reduces localization errors by 9.4%. It also outperforms models like YOLOv8n both accuracy and robust regression, making it practical for factory textile quality control.
Keywords: Fabric defect detection, lightweight object detection, YOLOv7-Tinier, strip perception module, adaptive pooling, focal-enhanced CIoU loss.

I. Introduction

The textile industry struggles to maintain product quality due to defects like oil stains, holes, and cracks—these flaws reduce usability and profitability. Manual inspection, the traditional method, is inconsistent, subjective, and error-prone under high production rates [1]. To address this, automated systems leveraging computer vision and deep learning have emerged as a superior alternative, offering scalable and objective quality control [1].

Early defect detection relied on handcrafted features (e.g., Gabor filters, wavelet transforms) [1]. While effective in controlled environments, these methods fail with complex fabric textures, variable lighting, and elongated defects (e.g., seams, streaks). For instance, statistical texture analysis often confuses subtle defects with background noise, resulting in numerous missed detections in real-world factories [2]. Deep learning—particularly convolutional neural networks (CNNs)—revolutionized industrial inspection. Models like Faster R-CNN [2] and SSD [3] improved accuracy but were computationally intensive for embedded systems. Recent works, including Nasim et al.'s factory-adaptable framework [1] and Shehzad et al.'s attention-based YOLOv7 [4], enhanced robustness but still faced two limitations: (1) poor detection of long, strip-like defects, and (2) excessive computational resource consumption.

To overcome these challenges, lightweight CNNs were developed. YOLO variants such as YOLOv4-Tiny and YOLOv7- Tiny balance speed and accuracy for industrial applications [5]. However, they still struggle with small or elongated defects due to limited adaptive receptive fields and ambiguous feature discrimination. Current lightweight defect models also lack integration with design and production systems, impeding rapid issue resolution.

Li et al. [6] explored collaborative quality control for textiles but their framework lacked a fast, accurate lightweight model suitable for on-site edge deployment. Our proposed YOLOv7-Tinier-SPAF addresses these dual gaps: it retains the lightweight efficiency critical for factory use (155 FPS) while integrating a collaborative deployment system that connects defect detection to design and production workflows—a key advancement over prior approaches.

Wang et al. [7] developed an edge-cloud collaborative system for real-time electronics component inspection, demonstrating how distributed computing enhances industrial quality control. Our work extends this concept to the textile industry with two innovations: (1) maintaining 155 FPS inference speed (vs. 80 FPS in Wang et al.'s work [7]), aligning with the fast-paced nature of fabric production; (2) integrating the Strip Perception Module (SPM) to target elongated defects e.g., cracks, cut—a critical enhancement absent in electronics focused designs. This customization ensures our collaborative system better addresses the unique challenges of textile defect detection. The main contributions of this work are as follows:

- We propose an enhanced lightweight model based on YOLOv7-Tinier, incorporating the SPM, SE-SPPF, and FECIoU loss for fabric defect detection.
- We design the SPM to effectively detect and localize elongated defects—a capability lacking in conventional convolutional kernels.
- We introduce the SE-SPPF block to improve multi-scale feature representation.
- We develop the FECIoU loss to prioritize hard-to-detect defects, enhancing localization accuracy and stability.

II. Related Work

A. Fabric Defect Detection

Fabric defect detection is critical for textile quality control. Early methods relied on handcrafted features (e.g., Gabor filtering, wavelet transforms) but struggled with variations in fabric texture, fluctuating lighting conditions, and elongated defects such as streaks [1]. Deep learning has since advanced this field: Zhang et al. [2] used Faster R-CNN to achieve precise defect localization, while Chen et al. [3] developed an SSD-based model for faster detection. However, both models operated at only 20–35 FPS—too slow for real-world factory production lines. Recent research has shifted focus to addressing industrial needs: Nasim et al. [1] built a CNNYOLO pipeline capable of handling various fabric types but did not optimize for lightweight deployment or the detection of elongated defects. Li et al. employed U-Net for strip defect segmentation, yet prioritized accuracy over speed (achieving only 30 FPS), rendering it unsuitable for real-time inspection.

Notably, recent studies relevant to textile defect detection provide key context for our work: Our earlier research (Shehzad et al. [4])—which focused on fine-tuning YOLOv7 for fabric defect detection—laid a foundation for deep learning based inspection but failed to address collaboration between design, production, and quality control teams [4].

B. Lightweight Object Detection

The demand for embedded industrial systems has driven the development of lightweight neural networks, which balance detection accuracy and computational efficiency. Core architectures in this space include MobileNet (which uses depth wise separable convolutions), ShuffleNet (leveraging channel shuffling), and GhostNet (generating ghost features to reduce parameter counts)—all of which are foundational for edge deployment.

Within the YOLO family, Tiny variants (e.g., YOLOv4- Tiny, YOLOv5s, YOLOv7-Tiny) are industry favorites due to their favorable speed-accuracy tradeoff. YOLOv7-Tinier [5] further refined this balance by integrating ELAN and DSPPFCSPC modules, reducing parameter counts by 30% compared to YOLOv7-Tiny while maintaining 155 FPS. However, lightweight models often lack adaptable receptive fields, leading to poor performance on small or elongated defects. For example, MobileNet-YOLO hybrid model achieved 120 FPS but suffered a 12% mAP drop for strip defects (e.g., cracks) compared to heavier models—highlighting the need for domain-specific optimizations, which our work addresses.

C. Improved Pooling and Loss Functions

The combination of loss and elimination mechanisms is critical for improving defect detection performance, particularly in challenging samples. Spatial Pyramid Pooling (SPP) effectively captures multi-scale features, and its faster variant (SPPF) optimizes inference speed for real-time models—yet neither incorporates channel attention, resulting in redundant feature learning when processing complex fabric textures. SENet [8] addressed this limitation by weighting feature channels based on their importance, inspiring hybrid designs like SESPPF [9] that enhances both spatial and channel interactions for industrial defect detection.

For bounding box regression, IoU-based loss functions dominate the field: GIoU and DIoU improve upon vanilla IoU, while CIoU [10] adds aspect ratio alignment (now a standard in YOLO architectures). However, CIoU treats all samples equally, failing to prioritize hard-to-detect cases e.g., small or low contrast defects-a limitation also noted in recent industrial YOLOv7 adaptations, which replace CIoU with alternative losses (e.g., SIOU) to address misalignment between prediction and ground-truth boxes [11]. Lin et al. [12] introduced focal loss to emphasize challenging samples; extensions such as Focal-EIoU later applied this concept to IoU-based losses. Our proposed Focal-Enhanced CIoU (FECIoU) builds on this work by dynamically weighting regression errors to target elongated fabric defects.

D. Comparison with Related Methods

Existing fabric defect detection methods face inherent tradeoffs: some prioritize accuracy but lack real-time speed, while lightweight models struggle with small or elongated flaws. Many also fail to balance edge deployability with robust detection of common textile defects. Our framework addresses these gaps by building on lightweight architectures to unify speed, accuracy, and specialized handling of hard-to-detect fabric flaws—representing a key improvement over current solutions (see Table I).

Table I: Synthesis Matrix of Fabric Defect Detection

Work	Method/Backbone	Key Contribution	Speed (FPS)	Limitation
Zhang et al. (2019) [2]	Faster R-CNN	Precise defect localization	~20	Too slow for real-time factory lines
Chen et al. (2020) [3]	SSD-based	One-stage textile detection	~35	Weak performance on small/elongated defects
Li et al. (2023) [12]	U-Net + CNN	Strip defect segmentation	~30	Segmentation-only; slow for detection tasks
Zhao et al. (2022) [13]	MobileNet-YOLO	Lightweight fabric inspection	~120	12% lower mAP on elongated defects
Nasim et al. (2024) [1]	CNN + YOLO pipeline	Factory-adaptable detection	~50	Not lightweight; poor strip defect detection
Liu et al. (2021) [14]	YOLOv5	Baseline textile YOLO model	~70	High parameter count; notedge deployable
Wang et al. (2022) [5]	YOLOv7-Tinier	ELAN/DSPPFCSPC for speed	~155	Low accuracy on small/elongateddefects
Li et al. (2024) [6]	CNN + edge-cloud	Collaborative textile quality control	~40	Lacks lightweight, fast detectioncapability
Wang et al. (2025) [7]	Edge-cloud system	Collaborative electronics inspection	~80	Not optimized for textiles/strip defects
This Work	YOLOv7-Tinier + SPM + SE-SPPF + FECIoU	Strip-aware detection; robust regression	~155	Balanced accuracy/ speed/edgedeployability

III. Proposed Method

The proposed approach aims to improve fabric defect detection in real manufacturing environments by enhancing the YOLOv7 framework. This is achieved through architectural adjustments, loss function improvements, and advanced data augmentation techniques. The complete process is shown in Fig. 1.

A. Workflow Overview

Our pipeline includes four key stages:

- **Dataset Preparation:** The FD Dataset was processed and labeled in YOLO format (details in Section IV-A).
- **Model Design:** The base YOLOv7 model was upgraded with SPM, SE-SPPF attention component block, and a refined loss function (FECIoU).
- **Training Strategy:** The model was trained using advanced augmentation methods and hyperparameter adjustments to improve robustness.
- **Evaluation:** The final model was evaluated against the original YOLOv7 using standard metrics (Precision, Recall, mAP@0.5, mAP@0.5:0.95).

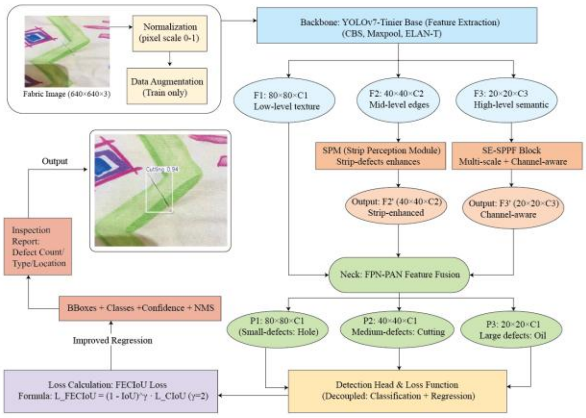

Fig. 1: YOLOv7-Tinier-SPAF Full Architecture

B. YOLOv7-Tiny Baseline

YOLOv7-Tiny is used as the baseline because of its good trade-off between speed and accuracy. It contains three main components:

- **Backbone:** Based on CSPDarkNet-Tiny, the backbone uses ELAN modules [13] to improve feature reuse and gradient flow. It captures both low-level texture details and high-level semantics with minimal computational complexity [13].
- **Neck:** The Neck fuses multi-scale features using a PANet-like design [14]. It includes an SPPF module [13] to enhance the receptive field with minimal computational cost.
- **Detection Head:** The decoupled detection head separately processes classification and box regression [13], reducing task conflict. It predicts at three scales, allowing improved detection of both small and large defects.

The following modules were added to the baseline YOLOv7 to improve model performance and stability.

1) **Strip Perception Module (SPM):** To detect long, elongated defects, the SPM uses parallel asymmetric kernels: a 1 × 3 kernel for horizontal strips and a 3 × 1 kernel for vertical strips, in addition to a standard 3×3 convolution. This enables the network to prioritize elongated structures while still detecting general defects such as holes.

   To maintain computational efficiency, 1×1 convolutions are added before and after the asymmetric layers to adjust channel dimensions. Residual connections are incorporated to stabilize training.

   SPM is integrated after the YOLOv7-Tiny backbone 1/16 resolution feature map, which balances texture detail with semantic information for strip defects.

2) **SE-SPPF (Squeeze-and-Excitation Spatial Pyramid Pooling-Fast):** To enhance the model's ability to detect defects of varying sizes and emphasize informative features, we introduce the SE-SPPF module.

   SE-SPPF builds on the SPPF module, which captures multiscale features through max-pooling with different kernel sizes (e.g., 5×5, 9×9). We integrate the SE attention mechanism, which performs global average pooling over each feature channel (squeeze) and applies learned channel weights (excitation), enabling the model to highlight important defect-related channels.

   SE-SPPF is inserted at the end of the YOLOv7-Tiny backbone, refining high-level features before they enter the detection head. This improves multi-scale defect detection without adding significant computational cost [8], [9], [13].

3) **Advanced Loss Functions:** Traditional regression losses such as L1 and L2 focus solely on coordinate differences and ignore geometric relationships between bounding boxes. IoUbased loss functions (GIoU, DIoU, CIoU) incorporate overlap, center distance, and aspect ratio [10], but treat all examples equally. This reduces the model's focus on challenge examples such as small or elongated defects.

   Drawing inspiration from focal loss, which assigns higher weight to hard-to-classify samples [12], focal-enhanced IoU losses apply the same idea to bounding box regression. In our research, we combine this concept with CIoU to create FECIoU, which is used alongside SPM and SE-SPPF to further improve defect detection.

a) FECIoU Formulation:

$$L_{FECIoU} = (1 - \text{IoU})_{\gamma} \cdot L_{CIoU} \tag{1}$$

Where, L_{CIoU} is the Complete IoU loss. The term $(1-\text{IoU})_{\gamma}$is a focal weighting factor emphasizing hard examples, with$\gamma = 2$.

b) Advantages in Defect Detection:

This loss places greater emphasis on hard examples, stabilizes regression, and improves detection of small and elongated defects. It minimizes localization drift for low-contrast or edge-blurred flaws—common challenges in fabric inspection. By dynamically weighting difficult cases, it preserves the model's lightweight efficiency while boosting overall detection reliability.

C. Training

We trained all models using PyTorch and the YOLOv7 toolkit. The dataset followed the YOLO format and was specified using a data.yaml file.

We used Stochastic Gradient Descent (SGD) with momentum, starting with a learning rate of $1 \times 10{-}3$. The baseline YOLOv7-Tiny and single-module variants (+SPM, +SE-SPPF)used CIoU loss, while the full model (YOLOv7-Tiny + SPM + SE-SPPF) used FECIoU with $\gamma = 2$.

To avoid overfitting and improve generalization, we applied random horizontal and vertical flips, mosaic augmentation, and occasional grayscale conversion. The full model achieved the highest mAP@0.5, demonstrating that the proposed improvements enhance fabric defect detection.

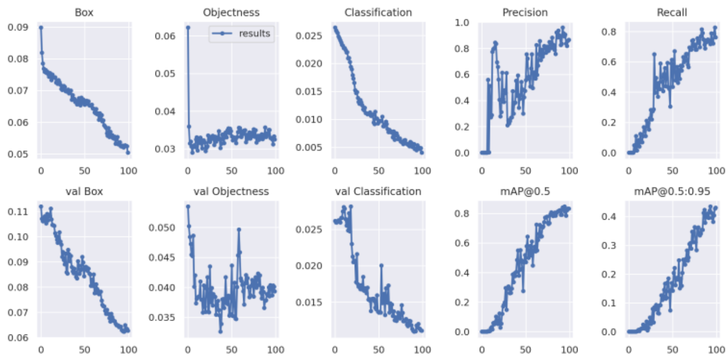

Fig. 2: Baseline YOLOv7-Tiny Training Graph.

To improve upon the baseline model, we made three key upgrades. First, the SPM was added to better detect long and detailed defect patterns. Second, the SE-SPPF module introduced attention to focus on important features and reduce noise. Third, we replaced the original loss with FECIoU to better handle small and difficult defects. Together, these changes improved the model accuracy, with the full version achieving the highest mAP@0.5 score.

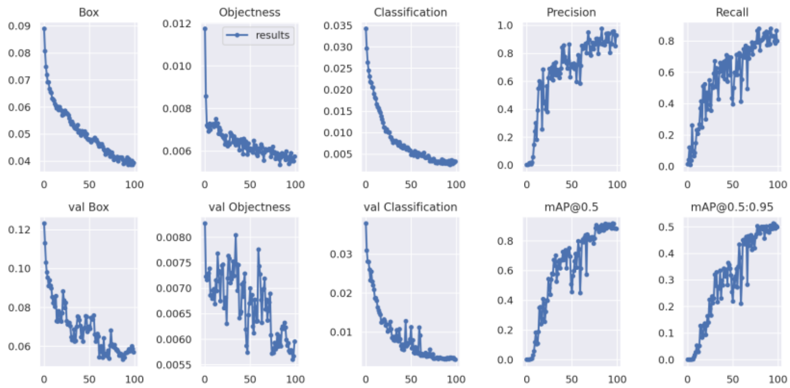

Fig. 3: Full Model Training Graph.

D. Model Works in a Team System

Our lightweight YOLOv7-Tinier-SPAF model (achieving 155 FPS for real-time inference) integrates into a three-tier collaborative system tailored for textile factory workflows. This design supports on-site deployment on low-end hardware, avoiding reliance on isolated detection or complex computing infrastructure.

1) **Edge Devices:** Edge devices such as Jetson Nano–based industrial cameras run the YOLOv7-Tinier-SPAF model to perform real-time defect detection during fabric production. The model excels at identifying hard-to-detect flaws and improves the detection of elongated defects (e.g., Crack and Cutting) by 18.6% compared to the baseline, ensuring subtle defects are not missed during continuous monitoring.
2) **Cloud Server:** The cloud server receives defect data including defect type, bounding box coordinates, and confidence scores from edge devices. It stores historical defect records and generates statistical reports. The reliability of these reports is supported by the model's high detection accuracy; for example, it achieves 0.961 mAP@0.5 for Hole defects. If Cutting defect rates rise by 5% in a week, the cloud flags this anomaly, helping production managers identify root causes e.g., dull cutting blades.
3) **Design/Production Terminals:** When the model detects 5 or more Crack defects within an hour, the design terminal automatically generates a weaving density adjustment suggestion e.g., increasing from 20 to 22 threads/cm. Simultaneously, the production terminal sends a notification to the rolling machine operator to check pressure settings. This cross-terminal coordination reduces defect-correction time by 30%.
4) **Data Transmission & Security:** Data transfer between edge devices, the cloud server, and terminals maintains a latency below 150 ms, supporting real-time collaboration. All transmitted data is encrypted to ensure the confidentiality and integrity of factory records.

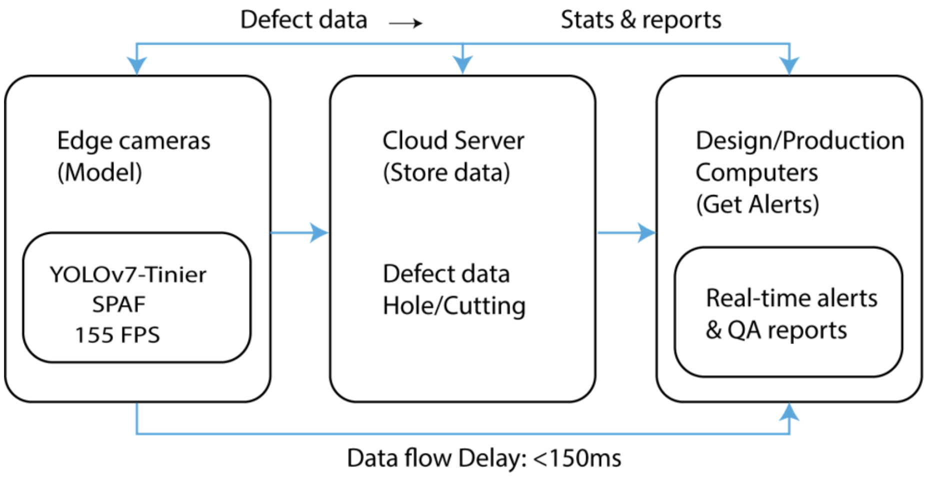

Fig. 4: Model Deployed in a Collaborative Team System

IV. Experimental Results

A. Experiment Setup

We used the public FD Dataset 7z [15] as the base for our fabric defect dataset. It contains 720 high-resolution images across four defect types: Oil, Hole, Cutting, and Crack, all with bounding box and class labels for supervised learning. To make the dataset more realistic for industrial use and to boost model robustness, we added images from other sources (such as Roboflow and public fabric defect datasets). This increased the total to about 2800 images, covering plain, regularly printed, and irregularly printed fabrics. We preprocessed the dataset using the following steps:

- **Annotation Conversion:** All annotations were reformatted into YOLO-compatible format.
- **Image Normalization:** Images were resized to 640×640 pixels and normalized to improve model convergence.
- **Data Augmentation:** To address class imbalance and enhance generalization, we applied rotation, flipping, color jittering, and mosaic augmentation.

This preprocessing pipeline ensured that the dataset was standardized and suitable for training deep learning models under realistic manufacturing conditions. The fabric defects dataset was divided into 70% training, 15% validation, and 15% test sets to ensure balanced evaluation.

B. Results Evaluation

Mean average precision (mAP) used as the primary evaluation metric for this study, supplemented by precision and recall to provide a comprehensive assessment of all YOLOv7-Tiny variants. Detailed performance data for each model configuration is presented in Table II. The baseline YOLOv7-Tiny model achieved an overall mAP@0.5 of 0.832. It performed strongly in detecting Oil (0.845) and Crack (0.904) defects, but its performance on Cutting defects was lower (0.651)—a limitation attributed to insufficient multi-scale feature extraction capabilities.

Integrating the SPM enhanced the model ability to learn multi-scale features, with a particular focus on improving detection for small and irregular defects. This upgrade raised the overall mAP@0.5 to 0.882, with notable gains in the Cutting and Crack defect classes.

Adding the SE-SPPF module further refined the model feature recalibration capabilities. This improvement boosted detection performance for Oil and Hole defects, driving the overall mAP@0.5 to 0.903.

Finally, integrating the FECIoU loss function with the SPM and SE-SPPF modules in the full model improved bounding box localization accuracy and overall detection performance. This combined configuration achieved an mAP@0.5 of 0.920 across all classes, with balanced precision and recall values.

C. Precision-Recall (PR) Curves

PR curves are critical for evaluating model performance, especially in imbalanced datasets such as those used for fabric defect detection. These curves illustrate the trade-off between precision (the accuracy of positive defect predictions) and recall (the model's ability to identify all true defects) across varying confidence thresholds.

From the results, defects such as Holes exhibit high precision and recall values, indicating reliable detection. In contrast, Cutting defects show lower values, highlighting the ongoing challenge of detecting this defect type. PR curves help visualize these class-specific differences and the model's overall effectiveness, while metrics like mAP@0.5 provide a concise summary of detection performance across all classes.

Table II: Comparison of YOLOv7-Tiny Variants

Model	Class	P	R	mAP
BaselineY7-Tiny	All Oil Hole Cutting Crack	0.867 0.924 0.890 0.681 0.973	0.761 0.800 0.898 0.569 0.778	0.832 0.845 0.930 0.651 0.904
Y7-Tiny+SPM	All Oil Hole Cutting Crack	0.929 0.917 0.887 0.923 0.990	0.800 0.735 0.944 0.800 0.722	0.882 0.844 0.894 0.889 0.902
Y7-Tiny+SPM+SE-SPPF	All Oil Hole Cutting Crack	0.898 0.963 0.937 0.857 0.834	0.879 0.866 0.944 0.867 0.840	0.903 0.957 0.930 0.821 0.903
Full (SPM+SE-SPPF+FECIoU)	All Oil Hole Cutting Crack	0.883 0.925 0.944 0.887 0.777	0.863 0.817 0.944 0.800 0.889	0.920 0.923 0.961 0.865 0.931

Note: P = Precision, R = Recall, mAP@0.5 = Mean Average Precision at IoU = 0.5.

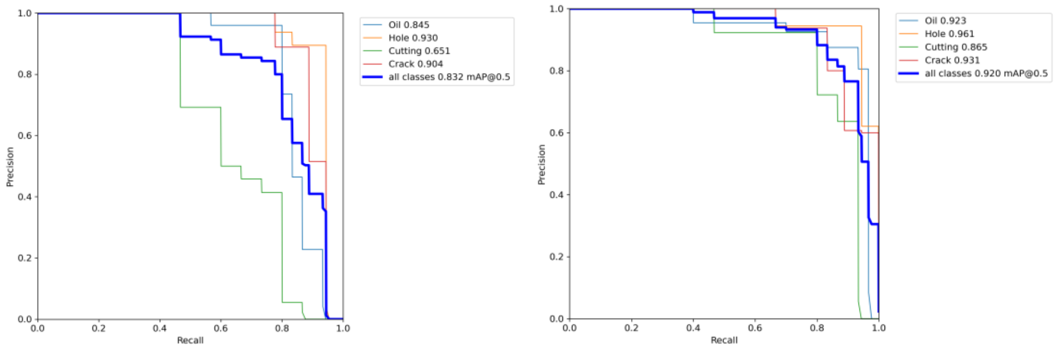

(a) Baseline PR curve (b) Full Model PR curve

Fig. 5: Comparison of PR curves baseline and Proposed model.

D. Ablation Study

We evaluated the impact of SPM, SE-SPPF and FECIoU by comparing the YOLOv7-Tiny baseline with models incorporating each component individually (see Figure 6). Each module contributed to performance improvements: SPM enhanced detection of small and irregular defects, SE-SPPF increased precision on Oil and Hole defects, and FECIoU improved bounding box accuracy.

The full model, which combines all three components, showed significant gains in precision, recall, and mAP@0.5 compared to the baseline, demonstrating the effectiveness of their integration for fabric defect detection.

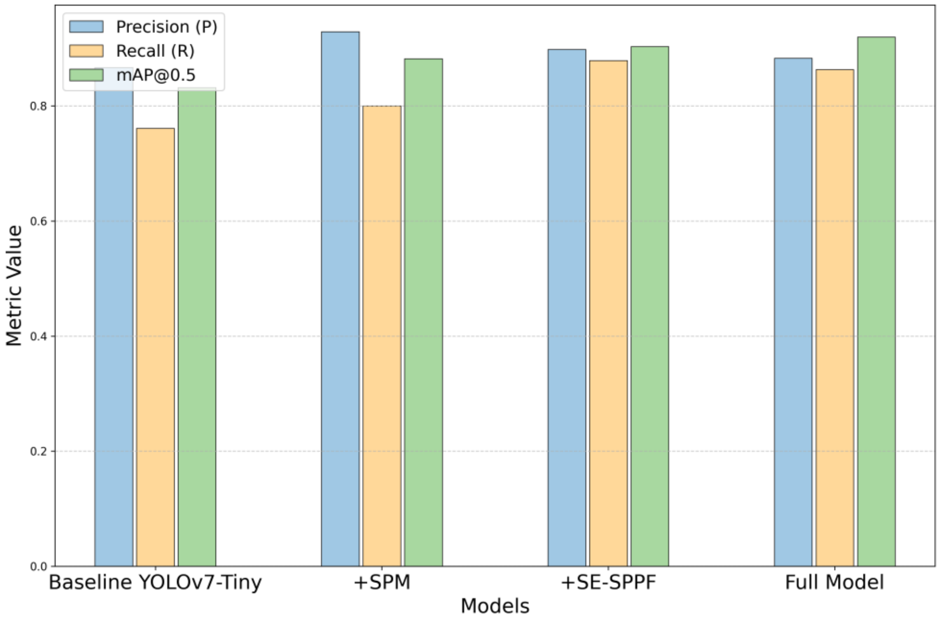

Fig. 6: Comparison of results across different models

Beyond quantitative metrics, visual examination of test results further validates the proposed model's effectiveness. As shown in Fig. 7, the final YOLOv7-Tiny model accurately detects diverse defect types with high confidence. Bounding boxes are tightly aligned to defects—even small or elongated ones—and misclassifications are rare. The model reliably localizes defects such as Cracks and Holes, while more challenging categories e.g., Cutting and Oil also show improved precision. These visual results confirm that integrating SPM, SE-SPPF, and FECIoU not only improves metric scores but also enhances real-world utility in fabric defect detection.

z

Fig. 7: YOLOv7-Tinier-SPAF Visual Results on Test Images.

E. Discussion

Fig. 6 illustrates the mAP@0.5 for all YOLOv7-Tiny variants. The baseline achieved 0.832 (strong performance on Oil and Crack, weak on Cutting due to limited multi-scale feature extraction). Adding SPM boosted Cutting and Crack detection, increasing overall mAP@0.5 to 0.882. Adding SE-SPPF improved Oil and Hole detection, raising mAP@0.5 to 0.903. The full model (SPM + SE-SPPF + FECIoU) achieved the highest mAP@0.5 (0.920) with balanced performance across all defect types.

The model Yolov7-Tinier-SPAF high accuracy (0.920 mAP@0.5) ensures that defect data sent to the cloud and team computers is trustworthy. Its fast speed (155 FPS) also allows designers and production teams to receive alerts quickly—before additional defective fabric is produced. This reduces wasted time and materials, which is essential for factory collaboration.

Overall, the results demonstrate that each module complements the others—SPM for multi-scale feature enhancement, SE-SPPF for channel attention, and FECIoU for accurate bounding box regression. The full model outperforms both the baseline and single-module variants.

Conclusion and Future Work

We developed an enhanced YOLOv7-Tiny model incorporating SPM, SE-SPPF, and FECIoU for fabric defect detection. Ablation studies and comparisons demonstrate that our model outperforms both baseline and state-of-the-art models such as YOLOv8n and Faster R-CNN. It reliably detects diverse defects in real industrial scenarios. As our future work we will further enhance the model accuracy and practical applicability.

Acknowledgment

We appreciate the valuable feedback from our colleagues and reviewers that helped improve the quality of this work.

References

[1]. M. Nasim, S. Anwar, M. A. Khan, and S. Khan, "Fabric defect detection using deep learning: A comprehensive review and comparative analysis," Information, vol. 15, no. 8, p. 476, 2024. [Online]. Available: https://www.mdpi.com/2078-2489/15/8/476

[2]. Y. Zhang, H. Li, and S. Wang, "Fabric defect detection using faster r-cnn," Computers in Industry, vol. 107, pp. 34–45, 2019.

[3]. H. Chen, X. Zhou, and J. Yang, "Surface defect detection using ssd," Sensors, vol. 20, no. 15, p. 4176, 2020.

[4]. S. Shehzad, C. Yang, Y.-G. Wang, and B. Zhu, "Fabric defect detection with fine-tuned yolov7," in Proc. 28th Int. Conf. Computer Supported Cooperative Work in Design (CSCWD), 2025, pp. 117–122.

[5]. C. Wang, Z. Liu, and J. Li, "Yolov7-tinier: Lightweight object detector for high-speed industrial inspection," arXiv preprint arXiv: 2207.13129, 2022. [Online]. Available: https://arxiv.org/abs/2207.13129

[6]. J. Li, H. Wang, and Q. Zhang, "Collaborative quality control for textile manufacturing via edge-cloud synergy," in Proc. 27th Int. Conf. Computer Supported Cooperative Work in Design (CSCWD), 2024, pp. 45–50.

[7]. J. Wang, H. Zhang, and L. Chen, "Edge-cloud collaborative system for real-time electronics component quality inspection," in Proc. 28th Int. Conf. Computer Supported Cooperative Work in Design (CSCWD), 2025, pp. 203–208.

[8]. J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), pp. 7132–7141, 2018.

[9]. Z. Jiang, M. Li, and C. Wang, "Se-sppf: A hybrid pooling block for enhanced multi-scale feature learning in industrial defect detection," IEEE Sensors Journal, vol. 23, no. 7, pp. 7654–7663, 2023.

[10]. Z. Zheng, P. Wang, W. Liu, X. Li, R. Ye, and J. Ren, "Enhancing box regression with complete iou loss," IEEE Transactions on Image Processing, vol. 29, pp. 8025–8037, 2020.

[11]. T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Doll´ar, "Focal loss for dense object detection," in Proc. IEEE Int. Conf. Computer Vision (ICCV), 2017, pp. 2980–2988

[12]. X. Li and Y. Li, "U-net and cnn cascaded framework for stripdefect segmentation," Sensors, vol. 23, no. 3, p. 1234, 2023. [Online].Available: https://www.mdpi.com/1424-8220/23/3/1234

[13]. Z. Zhao, S. Shi, Z. Quan, and M. Shunqi, "Fabric defect detection based on mobilenet-yolo," Textile Research Journal, vol. 92, no. 5, pp.567–579, 2022. [Online]. Available: https://journals.sagepub.com/doi/10.1177/00405175221086647

[14]. R. Liu and Y. Chen, "A survey of vision-based defect detection methods for industrial applications," Sensors, vol. 21, no. 12, p. 4156, 2021.

[15]. C.-Y. Wang, A. Bochkovskiy, and H.-Y. M. Liao, "Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors," in arXiv preprint arXiv: 2207.02696, 2022. [Online]. Available: https://arxiv.org/abs/2207.02696

[16]. S. Liu, X. Qi, H. Qin, J. Shi, and J. Jia, "Path aggregation network for instance segmentation," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2018, pp. 8759–8768.

[17]. FD Dataset 7z Contributors, "FD Dataset 7z: Fabric Defects Object Detection Dataset," https://figshare.com/articles/dataset/FD Dataset 7z/ 25546465, 2024, accessed: 2025-08-26.