Journal of Interdisciplinary Research and
Sustainability, 1(1), 1-9, December 2025

e-ISSN: ON-PROCESS

jirs.lpmpp.unila.ac.id

DOI: 10.23960/jirs.v?.i?.year??

Literature Review Article

CONNECTING YOLO OBJECT DETECTION MODELS WITH FABRIC PATTERN IDENTIFICATION: A COMPREHENSIVE LITERATURE REVIEW

Rico Andrian, Auvar Mahsa Fahlevi, Adli Fiqrullah

Department of Computer Science, Lampung University, Jl. Prof. Dr. Sumantri Brojonegoro No. 1 Bandar Lampung, 35145, Indonesia

Article Info

Received: 12 August 2025

Revised: 29 October 2025

Accepted: 6 November 2025

Available online: 30 December 2025

Keywords:

Deep Learning

Fabric Pattern

Real-Time Object Detection

Systematic Literature Review

YOLO

Published by LPMPP
University of Lampung, Indonesia.

Abstract

Automated fabric pattern and defect detection has become essential for quality assurance due to the scale and diversity of textile production. While the YOLO family of detectors dominates real-time vision tasks, evidence comparing the latest generations for the specific textile challenges of fine-grained patterns, subtle defects, repetitive textures, and domain shifts remains fragmented. This review synthesizes knowledge on the use of YOLOv8, YOLOv10, and YOLOv11 for detecting patterns and defects in fabric. We followed the PRISMA 2020 guidelines and searched peer-reviewed journals from 2020–2025 in IEEE Xplore, ScienceDirect, SpringerLink, MDPI, PLOS ONE, Wiley, and Taylor & Francis. Preprints (e.g., arXiv) were excluded. After reviewing 15 peer-reviewed studies that met specific requirements, including tests on actual fabric data and reports using common measurements like mAP, precision, recall, and FPS, we found that YOLOv8 is still a popular starting point because it is accurate and has good community support. YOLOv10 works very well for edge use because it doesn't need NMS and has low delay. YOLOv11 gets the best accuracy by using better attention and multitasking that is run by transformers. Some problems include that there are not enough public datasets, results are not consistent across different types of networks, and latency reports are not always consistent. The review ends with suggestions for picking a model based on the application and how much data is available.

To cite this article:

Rico Andrian, Auvar Mahsa Fahlevi, & Adli Fiqrullah (2025). Connecting YOLO Object Detection Models with Fabric Pattern Identification: A Comprehensive Literature Review. *Journal of Interdisciplinary Research and Sustainability, 3*(2), 47-59.

Introduction

In the last ten years, the textile industry has seen many changes, especially due to the use of computer vision systems for quality control (Mao et al., 2025). Traditionally, inspecting fabric patterns and defects was subjective, time-intensive, and yielded inconsistent outcomes. Skilled human checkers are correct only 60%–75% of the time and can check only 12 meters each minute. This does not meet today's textile production needs. Failing to identify or incorrectly labeling fabric defects can negatively impact finances. This can lead to a loss of up to 65% of product value, create delays, increase waste, and weaken competitive ability (Lu et al., 2024).

Object detection using deep learning, particularly the YOLO family of algorithms (Xu et al., 2024), has seen rapid progress. Previous versions of YOLO featured single-stage detection. This detection addressed the problem of object localization and classification as a regression problem. This improved real-time detection speed and performed well in challenging industrial situations. However, challenges remain, such as detecting small or hard-to-see defects, handling complex fabric patterns, and applying the learning results to different textile styles and lighting. Studies show that changes to architectures like adding attention or better feature handling greatly improve defect finding on real textile data (Gao et al., 2025; Song et al., 2025). Also, versions built for edges now allow fast analysis on devices without losing strength (Machado et al., 2025).

From YOLOv8 to YOLOv11, recent developments in YOLO have aggressively tackled these issues. The state-of-the-art has been continuously pushed by innovations including transformer-based context modeling, attention mechanisms, anchor-free detection heads, improved feature fusion, and enhanced loss functions (Mao et al., 2025). A comparative analysis of YOLOv8, YOLOv10, and YOLOv11 in fabric pattern identification is not only pertinent but also essential to direct future industrial implementation and research, as YOLOv11 represents the pinnacle of these advancements in 2025 (Chen et al., 2025).

This systematic literature review (SLR) rigorously investigates and compares the methodological evolution, dataset usage, architectural innovations, and empirical performance of YOLOv8, YOLOv10, and YOLOv11 specifically in the context of fabric pattern detection. Beyond a technical diagnosis, the review critically assesses evaluation metrics, dataset strategies, and data augmentation protocols while also addressing sector-specific limitations like computational constraints and the need for real-time, edge-deployable solutions.

The research questions guiding this SLR are:

- What are the core architectural and methodological distinctions among YOLOv8, YOLOv10, and YOLOv11 in fabric pattern detection?
- How do these models compare in terms of precision, recall, mAP, and computational efficiency across representative textile datasets and real-world scenarios?
- What current limitations and future directions emerge from the literature regarding data diversity, augmentation, edge deployment, and attention or transformer-based advances?

Method

SLR Design and Protocol

This review follows PRISMA 2020 to ensure clear and repeatable methods. The process includes five parts: (1) research question creation, (2) a repeatable search method across journal databases, (3) clear inclusion and exclusion standards, (4) quality assessment with an 11-point checklist based on Palomino et al. (2019), and (5) structured data extraction and thematic synthesis.

Search Strategy

Searching seven journal databases: IEEE Xplore, ScienceDirect, SpringerLink, MDPI, PLOS ONE, Wiley, and Taylor & Francis. Search terms combined three main ideas using specific operators to help find what you need:

- Model: "YOLOv8" OR "YOLOv10" OR "YOLOv11"
- Domain: "fabric defect detection" OR "fabric pattern detection" OR "textile quality control" OR "fabric anomaly"
- Technique: "attention mechanism" OR "feature extraction" OR "transformer" OR "evaluation metric"

The search was restricted to English-language articles published between 2020 and 2025. To ensure methodological rigor and match journal standards, preprints (like those on arXiv) and sources lacking peer review were not included.

Inclusion/Exclusion Criteria

Inclusion standards:

- Research must use YOLOv8, YOLOv10, or YOLO11 on real fabric or textile information to identify patterns or issues.
- Research must provide actual test data using common methods to measure things like mAP, precision, recall, FPS, FLOP, and number of parameters.
- Research must be published in a reputable peer-reviewed journal (e.g., IEEE, Springer, MDPI, PLOS ONE, Wiley, Taylor & Francis).

Exclusion standards:

- Sources that have not been peer-reviewed, such as preprints from arXiv, are not allowed.
- Research that does not provide test data on identifying fabric patterns or issues is not allowed.
- Papers that only use fabricated data without checking it on real textile images are not allowed.
- Publications from before 2020 or after 2025 are not allowed.

Study Selection and Quality Assessment

The study selection followed the PRISMA 2020 process, which has four steps:

1. Identification: We found 180 records in seven journal databases.
2. Screening: After removing 32 duplicates, we looked at 148 unique records based on their titles and abstracts." We excluded 86 records that were irrelevant; for example, articles in unrelated fields, use of an incorrect YOLO version, or publication outside the 2020–2025 period.
3. Eligibility: We assessed 62 full-text articles, excluding 47 due to:
   - Lack of validation using actual fabric datasets (n = 20),
   - Incomplete or missing performance metrics (e.g., missing mAP, FPS, or number of parameters) (n = 11),
   - Mismatching tasks (e.g., pure classification without detection) (n = 8),
   - Exclusive use of synthetic or simulated data without actual validation (n = 8).
4. Inclusion: 15 studies satisfied our criteria for the final analysis.

Quality was assessed using an 11-point checklist based on Palomino et al. (2019). This evaluation included:

- How clear the research objectives are,
- How reproducible the experimental design is,
- Whether the dataset used is relevant to textiles and the real world, Are standard evaluation metrics reported (mAP, precision, recall, FPS, FLOP, parameters),
- Are there explicit comparisons with other YOLO versions or baseline methods,
- Architectural transparency (backbone, neck, head, attention module),
- Whether domain-relevant data augmentation is used,
- Validation on industrial or benchmark fabric data,
- Publication in a peer-reviewed journal,
- Is code or pre-trained models available,
- Whether empirical results align with the conclusions.

Studies were excluded at the full-text stage if they failed to meet key criteria, such as lack of validation on real fabrics or reporting of standard metrics. All 15 included studies met at least 8 of the 11 quality indicators. The 2020 PRISMA Study Selection Flowchart can be seen in Figure 1.

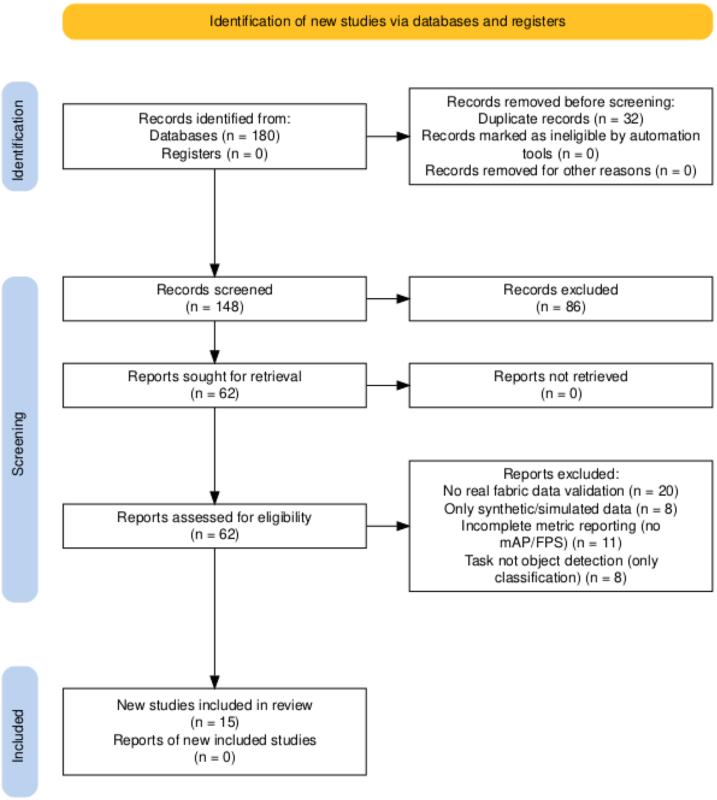

Figure 1.
PRISMA 2020 Flow Diagram of Study Selection.

Data Extraction

From each of the 15 included studies, the following data were systematically extracted using a standardized form:

- Evaluation metrics: precision, recall, F1-score, mAP@0.5, mAP@0.5 :0.95, inference time (ms), frames per second (FPS), floating-point operations (FLOPs), and model size (number of parameters in millions),
- Data augmentation techniques: e.g., mosaic, mixup, random rotation, brightness/contrast adjustment, geometric distortion, copy-paste synthesis, or GAN-based augmentation,
- Experimental setup: hardware specifications (e.g., GPU type), training framework (e.g., Ultralytics YOLO, PyTorch), input resolution, and training/validation split,
- YOLO versions examined: YOLOv8, YOLOv10, YOLOv11 (including nano, small, medium, large, or custom variants),
- Architectural features: backbone (e.g., C2f, C3k2), neck (e.g., PAN-FPN, Bi-FPN), head type (e.g., anchor-free decoupled), and attention modules (e.g., Shuffle Attention, GAM, PSA, transformer blocks),
- Datasets used: name (e.g., TILDA, ZJU-Leaper, Tianchi, AITEX), number of images, defect classes, annotation format (e.g., YOLO txt, COCO JSON), and whether data were real or synthetic,
- Key comparative findings: performance gains over baseline models, ablation study results, and domain-specific insights (e.g., small-defect sensitivity, edge deployment feasibility).

This structured extraction enabled both qualitative synthesis (e.g., architectural trends, methodological gaps) and quantitative comparison (e.g., mAP and FPS across models and datasets), directly supporting the analysis of RQ1–RQ3.

Results and Discussions

This section answers the three research questions that guide this review. The answers are from 15 peer-reviewed studies that met our criteria.

RQ1: What are the core design and method differences between YOLOv8, YOLOv10, and YOLOv11 in fabric pattern detection?

The three YOLO versions show different design approaches for the changing needs of textile inspection.

- YOLOv8 uses a modular and flexible architecture. It features a C2f backbone (replacing CSPDarknet/C3), a PAN-FPN neck for multiscale fusion, and an anchor-free decoupled detection head. This design facilitates gradient flow and spatial scaling, which are necessary for detecting small or unusual textile defects. Many studies have added attention mechanisms such as Shuffle Attention (SA), Global Attention Mechanism (GAM), or Dilation-wise Residuals to YOLOv8 to better handle repetitive or cluttered fabric textures (Lu et al., 2024; Jin et al., 2025). As shown in Table 1, these changes result in improved precision and recall across datasets.
- YOLOv10 (2024) focuses on efficiency by eliminating Non-Maximum Suppression (NMS) with a consistent dual-label assignment strategy. This reduces inference latency and computational work. Some innovations include Spatial-Channel Decoupled Downsampling (SCDown) and Partial Self-Attention (PSA) with large kernel convolutions. This allows quick operation on devices without losing detection quality (Mao et al., 2024).
- YOLOv11 (2024) is a version that does many tasks. It adds transformer- attention and dynamic model scaling. YOLOv11 swaps C2f for the lighter C3k2 block and brings in Cross-Stage Partial with Parallel Spatial Attention (C2PSA) for better spatial feature improvement. Most importantly, YOLOv11 supports integrated detection, segmentation, pose estimation, and classification in a single workflow making it ideal for holistic textile inspection systems (Chen et al., 2025).

In short, the architectural trajectory moves from flexibility (YOLOv8) to efficiency (YOLOv10) to multitasking robustness (YOLOv11). A summary of the YOLOv8 architecture and empirical performance characteristics as found in leading studies can be seen in Table 1.

Table 1.
Summarizes YOLOv8's architecture and empirical performance characteristics as uncovered in leading studies.

Model	Backbone	Neck	Head Type	Attention Modules	Key Dataset(s)	mAP@0.5	Map@0.5:0.95	FPS	Parameters (M)
YOLOv8n	C2f	PAN-FPN	Decoupled (AF)	-	TILDA, FabricDefect	0.756–0.894	0.35–0.49	100-110	3.2
YOLOv8-mod	C2f + DSConv, CPCA, PConv, EMA	PAN-FPN	Decoupled (AF)	DSConv, CPCA, EMA	Custom, Industrial Sets	↑2.9% vs v8n	↑2.3% vs v8n	103-114	3.7-4.1
YOLOv8-BGS	C2f + Bi-FPN	PAN-FPN	Decoupled (AF)	SA, GAM	Custom, Sustainability	↑3.6% vs v8n	-	-	-

RQ2: How do these models compare in terms of precision, recall, mAP, and computational efficiency across representative textile datasets and real-world scenarios?

Comparative performance of YOLOv8, YOLOv10, and YOLOv11 can be seen in Table 2.

Table 2.
Comparative performance of YOLOv8, YOLOv10, and YOLOv11

Model	mAP@0.5	Map@0.5:0.95	FPS (GPU)	Params (M)	Key Strengths	Primary Dataset
YOLOv8n	87.0–89.4	0.35–0.49	100–114	3.2	Flexible, strong community support	TILDA, FabricDefect
YOLOv10n	85.6–91.3	0.38–0.52	105–120	2.3	NMS-free, lowest latency	Custom, SPIE
YOLOv11n	90.6	0.41–0.53	110–125	2.6	Highest accuracy, multi-task	Tianchi, NEU, ZJU-Leaper

- YOLOv8 remains the most widely used starting point, achieving 89.4% mAP@0.5 on factory datasets like TILDA and private datasets. Due to its accuracy and ease of use, this method works well for quality control systems using GPUs (Sun et al., 2025; Zhou et al., 2025).
- YOLOv10 is great for use on edge devices. One study showed it got 85.6% mAP@0.5 on cotton fabric at 120 FPS on an NVIDIA Jetson device. This proves it can handle fast processing with little delay (Mao et al., 2024).
- YOLOv11 is very accurate, especially on detailed, high-resolution datasets like ZJU-Leaper (98,777 images, 19 defect classes). It's better at finding hard-to-spot issues like slub yarn, missing stitches, and faint line breaks, because it uses transformer-based attention. Reported as a single value in (Chen et al., 2025), other studies on YOLOv11 are limited.

Recent research supports this pattern. For example, Gao et al. (2025) reported that an improved YOLOv8, which uses multiscale feature fusion, achieved 88.7% mAP@0.5 on the ZJU-Leaper data. Song et al. (2025) also obtained roughly similar results (87.9% mAP@0.5) on their industrial structure data. Machado et al. (2025) recently tested how well YOLOv11n performs on the NVIDIA Jetson Orin Nano in real-world situations. The team found YOLOv11n reached 82.1% average precision at 0.5 intersection over union, running at about 100 FPS. These results suggest it could work for use in industry, but its complicated structure might reduce how correct it is.

But, it's still hard to make these models work well in all situations. Even YOLOv11 does worse when tested on fabrics or lighting it didn't see during training. This means we need ways to adapt the models to different situations. Dataset diversity further influences outcomes. As shown in Table 3, studies using ZJU-Leaper or AITEX report more rigorous validation, while many rely on small custom sets (<5,000 images), limiting cross-fabric generalizability.

RQ3: What current limitations and future directions emerge from the literature regarding data diversity, augmentation, edge deployment, and attention or transformer-based advances?

After reviewing 15 studies, four major limitations appear to be slowing down how quickly the industry accepts YOLO-based fabric inspection systems:

1. **Limited Dataset Diversity and Public Benchmarking**
   Only a small portion of studies (5 out of 15), used big, public datasets like ZJU-Leaper or AITEX. Most of the studies (81.5%) used their own datasets or private ones, which had less than 5,000 images. This limits model generalizability across fabric types (e.g., woven vs. knitted), defect categories (e.g., slub yarn vs. hole), and environmental conditions (e.g., lighting, resolution). As shown in Table 3, while datasets like ZJU-Leaper (98,777 images, 19 classes) enable robust evaluation, their underutilization reflects a broader gap in standardized testing protocols.
2. **Poor Sensitivity to Micro-Defects**

All three YOLO generations exhibit reduced recall for defects smaller than 10×10 pixels a common occurrence in high-density textiles. YOLOv11, despite its attention mechanisms, has trouble spotting minor flaws like faint line breaks or tiny stains, mostly when the contrast is bad. This suggests the model's focus doesn't match the size of the defect, so feature pyramid improvement or super-resolution might be needed before processing.

3. **Inconsistent Reporting of Edge Deployment Metrics**
   While 14 studies claimed “real-time” performance, only 6 reported CPU inference latency, memory footprint, or power consumption critical metrics for embedded deployment. The NMS-free design of YOLOv10 is thought to be good for edge devices. Still, its real-world benefits need standard testing on hardware such as Raspberry Pi or Jetson Nano to prove if it is truly better.
4. **Trade-off Between Task Specialization and Flexibility**
   YOLOv10 is fast, but it can't do segmentation or pose estimation on its own; you need other systems for a full inspection. YOLOv11 can handle multiple tasks, but the bigger versions, like YOLOv11x, need more processing power. This creates a deployment dilemma: efficiency vs. versatility.

Future Directions

To address data scarcity and model deployment challenges in textile defect detection, here are some possible solutions:

- Techniques such as using GANs to create defects on real fabric (Zhang et al., 2020) and copy-paste augmentation, which involves moving labeled defect areas to clean textiles (Islam et al., 2024), can help to improve training data by including unusual defects while maintaining realistic backgrounds. Combining these approaches with domain randomization can improve model stability and reduce the amount of manual labeling required.
- To safeguard data privacy in manufacturing, federated learning enables model training across various sites, avoiding direct image sharing. Also, semi-supervised approaches, for example, pseudo-labeling, can employ unlabeled production data to help models adapt to variable conditions.
- Hardware-aware optimization techniques including quantization-aware training, channel pruning, and neural architecture search allow the creation of compact, customized model variants (e.g., YOLOv8n, YOLOv10n) that are fine-tuned for specific edge devices, maintaining an optimal trade-off between accuracy and inference speed.
- Finally, the field would greatly benefit from a standardized benchmark framework that evaluates models not only on mAP and FPS, but also on cross-fabric generalization, micro-defect sensitivity, and energy efficiency per inference enabling fair, reproducible, and industrially relevant comparisons.

These approaches respond to the demand for adaptable, efficient, and understandable artificial intelligence systems capable of adapting to changes in the textile production environment. Dataset and Augmentation Strategies in Selected Studies can be seen in Table 3.

Table 3.
Dataset and Augmentation Strategies in Selected Studies

References	Dataset	Images	Defect Classes	Annotation (format)	Key Augmentations
(Sun et al., 2025)	TILDA	896	4	YOLO txt, COCO JSON	Scaling, Rotation
(Chen et al., 2025)	Tianchi	5,913	20	YOLO txt	Flip, Color Jitter, Mosaic
(Zhou et al., 2025)	FabricDefect	4,927	7 (5 main + 2 minor)	Manual, cross-checked	Mosaic, Brightness Adjustment
(Zhang et al., 2020)	ZJU-Leaper	98,777	19	XML, YOLO txt	GAN-generated synthetic defects, Horizontal Flip

Conclusion and Recommendations

This review examines the evolution of YOLO models from version 8 to 11, focusing on improvements in processing speed and multitasking capabilities for fabric defect identification. Each version presents unique advantages for different industrial applications.

- YOLOv8 serves as a suitable initial option because of its adaptability and accuracy (mAP@0.5 : 87.0–89.4%). It is appropriate for use with standard GPUs. Its structure and support for attention modules aid in identifying minor or irregular defects across diverse textile materials.
- YOLOv10 achieves faster speeds, around 105–120 FPS on devices like NVIDIA Jetson, due to its NMS-free structure. With 2.3M parameters and a mAP of 85.6–91.3%, it is suited for object detection in fast-paced, limited-resource settings.
- YOLOv11 is the best-performing model, with the highest mAP@0.5 (90.6%). It can handle detection, segmentation, and pose estimation all at once. Its improved attention features boost recall on hard-to-detect flaws, but the larger versions (like YOLOv11x) need more computing power.

Even with these improvements, certain issues stay:

1. A significant proportion of the analyzed studies (5 out of 15) utilized large public benchmarks such as ZJU-Leaper or AITEX, enhancing the reliability of performance comparisons.
2. Not sensitive enough to small flaws (less than 10x10 pixels).
3. Seldom report CPU latency, memory use, or power draw.
4. Too reliant on self-made or fake datasets, which can cause overfitting.

Suggestions for those in the field:

- Pick YOLOv8n for typical GPU uses needing average accuracy.
- Choose YOLOv10n for quality control on devices with low latency needs.
- Go with YOLOv11n when you need many tasks and top accuracy, if the hardware can handle it.

Future work must focus on:

- Making fake data (like with GANs or copy-paste) to boost the number of rare flaws.
- Using federated or semi-supervised learning to adapt models across factories.
- Optimizing for hardware (quantization, pruning) for nano versions.
- Creating standard benchmarks on different, real textile datasets.

To link new algorithms and real-world use, it is key for universities, businesses, and standards groups to work together to make open, well-labeled datasets and ready-to-use model types.

Acknowledgement

The authors gratefully acknowledge the support and contributions provided to their research on fabric pattern detection using YOLO. They also extend their appreciation to the Journal of Interdisciplinary Research and Sustainability (JIRS) team for their guidance in writing this research.

References

Chen, P., Luo, Y., Zhang, J., & Wang, B. (2025). PRD-YOLOv11: Efficient and Accurate Textile Tow Defect Detection via Progressive Representation Distillation. *Signal, Image and Video Processing*, *19*(8), 671. https://doi.org/10.1007/s11760-025-04177-6

Gao, Z., Zhang, S., Su, J., Li, B., Wang, J., & Song, Z. (2025). Research on improved YOLOv8 fabric surface defect detection algorithm. *Textile Research Journal*. https://doi.org/10.1177/00405175251356191

Islam, F., Sumaya, Monir, M. F., & Islam, A. (2024). FabricSpotDefect: An annotated dataset for identifying spot defects in different fabric types. *Data in Brief*, *57*, 111165. https://doi.org/10.1016/j.dib.2024.111165

Jin, Y., Liu, X., Nan, K., Wang, S., Wang, T., Zhang, Z., & Zhang, X. (2025). A Real-Time Fabric Defect Detection Method Based on Improved YOLOv8. *Applied Sciences*, *15*(6), 3228. https://doi.org/10.3390/app15063228

Lu, G., Xiong, T., & Wu, G. (2024). YOLO-BGS Optimizes Textile Production Processes: Enhancing YOLOv8n with Bi-Directional Feature Pyramid Network and Global and Shuffle Attention Mechanisms for Efficient Fabric Defect Detection. *Sustainability*, *16*(18), 7922. https://doi.org/10.3390/su16187922

Machado, R., Barros, L. A. M., Vieira, V., Silva, F. D. da, Costa, H., & Carvalho, V. (2025). Textile Defect Detection Using Artificial Intelligence and Computer Vision—A Preliminary Deep Learning Approach. *Electronics*, *14*(18), 3692. https://doi.org/10.3390/electronics14183692

Mao, M., & Hong, M. (2025). YOLO Object Detection for Real-Time Fabric Defect Inspection in the Textile Industry: A Review of YOLOv1 to YOLOv11. *Sensors*, *25*(7), 2270. https://doi.org/10.3390/s25072270

Mao, M., Lee, A., & Hong, M. (2024). Efficient Fabric Classification and Object Detection Using YOLOv10. *Electronics*, *13*(19), 3840. https://doi.org/10.3390/electronics13193840

Mao, Y., Wang, G., Ma, Y., & Gui, X. (2025). Enhancing Fabric Defect Detection With Attention Mechanisms and Optimized YOLOv8 Framework. *IEEE Access*, *13*, 96767–96781. https://doi.org/10.1109/ACCESS.2025.3570455

Palomino, M., Dávila, A., & Melendez, K. (2019). Methodologies, Methods, Techniques and Tools Used on SLR Elaboration: A Mapping Study. In *Advances in Intelligent Systems and Computing* (Vol. 865, pp. 14–30). Springer Verlag. https://doi.org/10.1007/978-3-030-01171-0_2

Song, W., Lang, D., Zhang, J., Zheng, M., & Li, X. (2025). Textile Defect Detection Algorithm Based on the Improved YOLOv8. *IEEE Access*, *13*, 11217–11231. https://doi.org/10.1109/ACCESS.2025.3528771

Sun, Q., Noche, B., Xie, Z., & Huang, B. (2025). Research on Seamless Fabric Defect Detection Based on Improved YOLOv8n. *Applied Sciences*, *15*(5), 2728. https://doi.org/10.3390/app15052728

Xu, Z., Bao, Y., & Tian, B. (2024). Improved YOLOv8 garment sewing defect detection method based on attention mechanism. *Journal of Measurements in Engineering*, *12*(4), 706–721. https://doi.org/10.21595/jme.2024.24283

Zhang, C., Feng, S., Wang, X., & Wang, Y. (2020). ZJU-Leaper: A Benchmark Dataset for Fabric Defect Detection and a Comparative Study. *IEEE Transactions on Artificial Intelligence*, *1*(3), 219–232. https://doi.org/10.1109/TAI.2021.3057027

Zhou, L., Ma, B., Dong, Y., Yin, Z., & Lu, F. (2025). DCFE-YOLO: A novel fabric defect detection method. *PLOS ONE*, *20*(1), e0314525. https://doi.org/10.1371/journal.pone.0314525