



A Keen Digital Eye: Using Computer Vision to Develop Low-cost, High Accuracy Defect Detection Systems for Small Garment Manufacturers

Sparsh Goyal

Student

Delhi Public School Agra

Abstract

Small and medium enterprise (SME) garment manufacturers in developing nations operate under severe structural and financial constraints. Raw fabric accounts for 60% to 70% of total production costs. With razor-thin profit margins, material efficiency and quality assurance are vital to commercial survival. However, SMEs predominantly rely on manual visual inspection, an error-prone process where inspection accuracy peaks at only 60% to 75% under ideal conditions. Unflagged defects lead to severe material waste, lost labour, and costly penalties imposed by buyers.

Computer vision (CV) and deep learning offer a scalable framework to automate quality control through early defect interception, defect-aware CAD marker nesting, and in-line seam inspection. While traditional automated systems historically required large capital outlays, recent advancements in low-cost edge computing hardware—such as NVIDIA’s Jetson series—and open-source machine learning models have significantly lowered deployment barriers. In addition, modern CV systems support simple, “bolt-on” brownfield integration avoiding costly production overhauls. This technological shift allows garment SMEs to protect operating margins, satisfy international compliance standards, and maintain global competitiveness.

Introduction

The ready-made garment sector serves as a vital catalyst for economic development, industrialization, and employment generation in the Global South. However, while multinational fashion conglomerates capitalize on agile production networks, small and medium enterprise (SME) garment manufacturers in developing nations operate under severe structural constraints. These include capital constraints, unfavourable supply chains, and international compliance requirements. All of these constraints collude to make profit margins razor-thin.

These thin margins make material efficiency and quality assurance primary determinants of commercial survival. Raw fabric constitutes between 60% and 70% of the total production cost of a finished garment (Yeasar, 2026). Measures to reduce avoidable material waste and post-assembly defects directly enhances a garment SME’s operating capital.

Despite operating in an increasingly automated global market, garment SMEs in developing nations largely rely on manual visual inspection for quality control. However, human eyes and brains are fundamentally constrained by physiological thresholds. Even veteran workers struggle to monitor defects and wastage accurately beyond 60-75 minutes at a stretch. As small factories face rising buyer standards, compressed lead times, and escalating material costs, computer vision (CV) and deep learning technologies offer a scalable framework to eliminate structural waste, safeguard product quality, and restore competitiveness against larger industrial conglomerates.

The Dual Crisis of Waste and Quality Control in Small Apparel Factories

The operational toll of manual inspection manifests in two interconnected areas: raw material waste and downstream quality failure. Traditionally, fabric rolls arriving at a factory are inspected on light tables. Trained human operators visually scan moving textiles under intense illumination. Empirical studies demonstrate that human inspection accuracy peaks at approximately 60% to 75% under ideal conditions, deteriorating rapidly after an hour of continuous labour due to eye strain and cognitive fatigue (WarpDriven, 2025; Yeasar, 2026). Subtle surface anomalies, such as tension slubs, needle lines, minor oil stains, broken pick threads, and subtle shade variations, frequently pass undetected into the cutting room.

When defective fabric moves unflagged into automated or manual cutting operations, the financial loss escalates exponentially. A defect cut into individual garment panels, such as sleeves, fronts, or collars, compromises the entire piece. If the flaw is discovered during assembly, the cut panel must be scrapped, resulting in wasted fabric, lost operator labour, and line disruption.

If the flaw remains undetected throughout assembly, finishing, and packaging, the finished garment is exported to international brand buyers. When buyers discover defects at the destination port, they can trigger severe commercial penalties. These include full-batch chargebacks, mandatory price discounts, or even cancelled supplier contracts. Data indicates that uncorrected fabric defects account for up to 85% of all garment defects identified during final audit, reducing the market value of compromised textile lots by 45% to 65% (Carrilho et al., 2024; Yeasar, 2026). Therefore, automated systems that eliminate human variability in quality control could greatly reduce waste and increase profitability for garment SME.

Mechanics of Computer Vision-Based Defect Detection

Automated computer vision systems replace human observation with continuous optical inspection, removing subjectivity in favour of clearly defined defect metrics. These systems combine high-resolution line-scan industrial cameras, specialized LED backlighting, optical filters, and deep learning neural networks.

Line-scan cameras capture the unrolling fabric line by line across its full width, compiling a high-resolution, distortion-free digital image of moving textile rolls operating at production speeds exceeding 60 meters per minute (Brightpoint AI, 2026). The captured visual feed is analyzed in real time using Convolutional Neural Networks (CNNs) (WarpDriven, 2025; Yeasar, 2026). These deep learning models are trained on extensive industrial datasets containing thousands of labeled textile anomalies across woven, knitted, and denim substrates.

Computer vision resolves quality and waste challenges through three distinct mechanisms:

- 1. Early Defect Interception at the Edge:** When retrofitted onto circular knitting machines, dyeing lines, or fabric-spreading tables, edge-based CV hardware processes image frames locally without cloud latency. Upon detecting a structural anomaly, such as a broken yarn or a hole, the edge controller signals the machine to halt instantly (Carrilho et al., 2024). Catching a defect at meter five of a 100-meter roll prevents the machine from producing 95 meters of ruined fabric, containing raw material losses at the point of origin.
- 2. Defect-Aware Marker Nesting:** For raw fabric containing scattered, unavoidable flaws, the computer vision system generates a digital coordinate map ("defect map") of the entire roll. This spatial data is exported directly into Computer-Aided Design (CAD) pattern-nesting software. The nesting algorithm automatically arranges garment pattern pieces around the mapped coordinates, ensuring pattern blades cut exclusively from pristine fabric zones (BrightPoint AI, 2026). This optimization increases material yield and reduces total fabric waste by 10% to 15%.
- 3. In-Line Seam and Assembly Integrity:** Beyond raw textiles, specialized CV camera arrays mounted at sewing workstations inspect completed seams and stitching. Deep learning models evaluate stitch density, seam puckering, thread tension, and pattern alignment across joined panels frame by frame (Richardson, 2026). Identifying a skipped stitch or misaligned seam directly at the sewing line allows operators to rework the garment immediately, preventing defective items from progressing through finishing and packaging.

Put together, computer vision-based defect detection systems can identify textile defects faster and with a lower likelihood of error than textile workers. Unfortunately, the costs involved in purchasing such systems and retrofitting existing production lines to use them has kept automated defect detection out of reach for many SMEs. However, advancements in hardware and machine learning are changing this.

A Comparison of Defect Detection Systems

Historically, automated visual inspection required capital outlays exceeding USD 100,000, placing it out of reach for smaller manufacturers. However, advances in edge computing, lower camera costs, and open-source machine learning models have lowered deployment expenses (Udayamurthy, 2025).

Low-cost edge hardware platforms such as NVIDIA’s Jetson series can process continuous image streams to identify defects, using appropriate image processing algorithms. Two categories of CNNs have emerged in recent years as leading contenders in textile defect detection. Single-stage anchor-free object detectors are one, while deep residual backbones within multi-stage detection frameworks are the other.

You Only Look Once, or YOLO is a prime example of the former, designed for direct end-to-end detection. Its backbone relies on a Cross-Stage Partial structure with C2f (Cross-Stage Partial with 2 Convolutions) modules, concatenating bottleneck outputs to enrich feature representation while reducing parameters (Johnson, 2026). The neck incorporates a Path Aggregation Network combined with a Feature Pyramid Network (PAN-FPN) to fuse low-level spatial detail (crucial for micro-pinholes) with high-level semantic context. Unlike earlier anchor-based models, YOLOv8 decouples classification and bounding box regression into separate parallel streams, removing anchor box tuning and reducing post-processing overhead during Non-Maximum Suppression (NMS) (Ma et al., 2024).

Residual Networks (ResNet) are a good example of the latter. These algorithms were originally introduced to solve the vanishing gradient problem in deep networks using identity shortcut connections. However, ResNet is strictly an *image classification backbone*. To detect and locate fabric defects, it must be paired with external region-proposal or detection heads. For instance, a Region Proposal Network (RPN) extracts candidate defect bounding boxes, followed by RoI Align and final classification layers (He et al., 2016). This is a sequential process with high accuracy but it carries a significant time penalty.

The table below compares some of the performance metrics of YOLOv8 and ResNet-50 based systems running on NVIDIA Jetson Orin Nano hardware with a 640x640 pixel input stream:

Table 1: YOLOv8 vs ResNet-50 benchmarks on NVIDIA Jetson Orin Nano

Architectural Metric	YOLOv8s (Small)	ResNet-50 + Faster R-CNN
Parameter Count	11.2 M	~34.2 M
FLOPs (at 640x640)	28.6 GFLOPs	~130.0+ GFLOPs
Inference Latency (TensorRT FP16)	~18.2 ms	~68.0 ms
Throughput (Jetson Orin Nano)	~55 FPS	~14 FPS
Peak VRAM / RAM Footprint	~850 MB	~2.6 GB
Detection Accuracy (mAP@50)	94.8%	95.1%

(He et al., 2016; Ma et al., 2024; NVIDIA Corporation, 2026)

On a lower cost Jetson Orin Nano board (less than USD 299), YOLOv8 runs nearly four times as fast as ResNet-50 + Faster R-CNN. To achieve comparable throughput, ResNet-50 + Faster R-CNN requires more

advanced and expensive hardware, such as NVIDIA's Jetson AGX Orin, which can cost as much as USD 2,000. That is nearly 7 times more expensive.

While ResNet-50 paired with multi-stage detection heads yields exceptional spatial feature extraction capable of identifying minute, low-contrast fabric defects achieving up to 95.1% defect detection accuracy, its computational complexity and high memory usage make it ill-suited for low-cost edge platforms.

YOLOv8 represents the superior architecture for industrial edge deployment. By combining an anchor-free decoupled head, efficient feature extraction, and PAN-FPN multi-scale feature fusion, YOLOv8 delivers comparable accuracy while achieving 4x to 8x higher frame rates and using less than half the memory footprint. Consequently, YOLOv8 enables textile manufacturers to deploy real-time fabric defect inspection on low-cost hardware like the NVIDIA Jetson Orin Nano, delivering a cost-effective and scalable solution for quality control.

The Economics of Retrofitting Automated Defect Detection Systems

A major misperception among small garment producers is that adopting computer vision requires replacing legacy machinery with expensive automated hardware. In practice, modern industrial computer vision systems are designed specifically for *brownfield integration*, i.e. retrofitting existing factory floors with minimal operational disruption.

Features that make brownfield integration simple and cost-effective (Nguyen, 2026)

- **Non-Invasive Hardware Mounting:** Cameras, optical sensors, and LED light bars are attached to existing machine frames using modular bolt-on brackets. These devices can be easily attached to specific machines or areas of the production line, such as fabric spreaders, circular knitting rigs, or manual inspection tables, while keeping the core of the production line unchanged.
- **Simple Signal Interfacing:** To halt a machine upon defect detection, technicians connect an optocoupled relay switch from the edge computer directly to the machine's existing emergency stop (E-stop) circuit. This modification takes hours rather than days and requires no structural alteration to legacy equipment
- **Wireless Data Pipelines:** Edge hardware processes image feeds locally and transmits lightweight metadata (e.g., defect counts and location coordinates) over local Wi-Fi networks to factory dashboards. This eliminates the need for expensive structural network cabling across the shop floor.

As described in the previous section, single-stage CNNs hold clear advantages for SMEs. A complete defect detection system that uses YOLOv8 could cost between USD 3,000 and 10,000 per production line, including hardware costs and retrofitting expenses. These systems use pre-trained transfer learning models accessed via low monthly Software-as-a-Service (SaaS) subscriptions, costing an additional USD 100-300 per month (Nguyen, 2026; Udayamurthy, 2025).

Larger textile units with both the resources and need for more accurate and capable systems can consider using ResNet based CNNs. Comprehensive multi-camera deployments featuring customized optical enclosures, high-resolution line-scan hardware, and direct integration with factory-wide Enterprise Resource Planning (ERP) systems can cost anywhere from USD 50,000 to 150,000 (Nguyen, 2026). While such systems may be out of reach for the typical textile manufacturing SME in the developing world, they represent the next stage of industrial hardware as SMEs grow and their quality control requirements change.

Conclusion

For small garment manufacturers in developing countries, material waste and quality inconsistency represent severe threats to long-term commercial viability. Relying on manual inspection leaves small enterprises vulnerable to human error, material loss, and buyer penalties. Deploying computer vision and deep learning inspection systems provides an accessible, high-precision alternative. By identifying fabric flaws at the source,

optimizing cutting room yields, and verifying seam integrity in real time, computer vision technology allows small manufacturers to protect their operating margins, meet strict international compliance standards, and remain competitive within global supply chains.

As computer vision algorithms continue to grow in capability while the hardware needed to run them becomes smaller and less expensive, we may see ever greater deployment of these systems among garment manufacturers across the world. These automated systems will also enable textile workers to focus on the creative and detail-oriented aspects of garment manufacturing, and can usher a new phase of accessible and unique clothing.

1. References

- a. Brightpoint AI. (2026, June 12). *AI-Based Fabric Defect Detection: A Complete Guide for Textile Manufacturers in 2026*. <https://www.brightpoint.ai/post/ai-based-fabric-defect-detection-a-complete-guide-for-textile-manufacturers-in-2026>
- b. Carrilho, R., Yaghoubi, E., Lindo, J., Hambarde, K., & Proença, H. (2024). Toward Automated Fabric Defect Detection: A Survey of Recent Computer Vision Approaches. *Electronics*, 13(18), 3728. <https://doi.org/10.3390/electronics13183728>
- c. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- d. Johnson. (2026, July 14). *AI Vision vs Manual Inspection: Total Cost and Accuracy Comparison*. iFactory AI. <https://ifactoryapp.com/ai-vision-camera/ai-vision-vs-manual-inspection-total-cost-accuracy-comparison>
- e. Ma, S., Liu, Y., & Zhang, Y. (2024). Fabric Defect Detection Based on Improved Lightweight YOLOv8n. *Applied Sciences*, 14(17), 8000. <https://doi.org/10.3390/app14178000>
- f. Nguyen, D. A. (2026, January 15). *How Much Does It Cost to Build Computer Vision Systems for Manufacturing: Complete Price Breakdown*. SmartDev. <https://smartdev.com/computer-vision-cost-manufacturing-complete-cv-system-pricing-breakdown/>
- g. NVIDIA Corporation. (2026). *Jetson Benchmarks*. NVIDIA Developer. <https://developer.nvidia.com/embedded/jetson-benchmarks>
- h. Richardson, C. (2026, June 5). *Garment Defect Detection AI Stitching Seam Trim Inspection*. iFactory AI. <https://ifactoryapp.com/industries/textile-manufacturing/garment-defect-detection-stitching-seam-ai>
- i. Udayamurthy, S. (2025, September 25). *AI Visual Inspection System Pricing: What Manufacturers Should Expect*. Jidoka Tech. <https://www.jidoka-tech.ai/blogs/ai-visual-inspection-pricing-guide-manufacturers>
- j. WarpDriven. (2025, November 1). *The Future of Computer Vision in Fabric Inspection and Detection*. WarpDriven AI. <https://warpdriven.ai/blog/industry-1/computer-vision-in-fabric-inspection-and-detection-254>
- k. Yeasar, S. (2026, February 2). *Fabric Defect Detection and Reduction Using Vision-Based Deep Learning Techniques*. Textile Learner. <https://textilelearner.net/fabric-defect-detection-and-reduction/>