Textile & Leather Review

ISSN 2623-6281| www.tlr-journal.com | 10.31881/TLR

Computer Vision-Based Anomaly Diagnosis in Knitted Fabrics: A Graph-Theoretic Approach to Stitch Defect Localization

Jing Jin

How to cite: Jin J. Computer Vision-Based Anomaly Diagnosis in Knitted Fabrics: A Graph-Theoretic Approach to Stitch Defect Localization. Textile & Leather Review. 2026; 9:197-225. https://doi.org/10.31881/TLR.2026.197

How to link: https://doi.org/10.31881/TLR.2026.197

Published: 28 January 2026

Computer Vision-Based Anomaly Diagnosis in Knitted Fabrics: A Graph-Theoretic Approach to Stitch Defect Localization

Jing Jin

Information Engineering School, Hunan Open University, Changsha 410004, Hunan, China
18670786127@163.com

Article
https://doi.org/10.31881/TLR.2026.197
Received 2 July 2025; Accepted 16 September 2025; Published 28 January 2026

ABSTRACT

Automated quality inspection remains a critical challenge in textile manufacturing, where subtle fabric defects can significantly impact product value and performance. Existing computer vision systems struggle to simultaneously achieve precise stitch-level defect localization and robust generalization across diverse knitting patterns and deformations. This paper presents a novel hybrid architecture combining dual-branch convolutional neural networks with edge-conditioned graph neural networks, where visual features inform topological graph representations of stitch connectivity patterns for comprehensive anomaly detection. Experimental results demonstrate superior performance with 91.3% detection accuracy (vs. 85.7% for commercial systems), 92.8% overall recall, with strong performance on critical defect categories, and real-time processing at 20 FPS, while reducing the overall false positives rate to 2.3%, and up to 8.9% improvement over the baseline 3D CNN on connectivity defects. The proposed framework establishes a new paradigm for fabric inspection by integrating computer vision with structural graph analysis, achieving both high precision and practical deployment efficiency for industrial quality control applications.

KEYWORDS

computer vision, knitted fabric defect detection, edge-conditioned convolution, textile quality control, industrial inspection

INTRODUCTION

The research on anomaly diagnosis in knitted fabrics holds significant practical and theoretical value for modern textile manufacturing and quality control [1]. As knitted fabrics increasingly serve as critical materials in apparel, medical textiles, and technical applications, even minor structural defects such as dropped stitches,

yarn breaks, or pattern distortions can severely compromise product performance, durability, and aesthetic quality.

Traditional manual inspection methods, while still prevalent in many factories, suffer from limitations including subjective judgment, fatigue-induced errors, and inefficiency when handling large-scale production. This creates an urgent need for automated, intelligent detection systems that can provide consistent and reliable quality assessment throughout the manufacturing process. Computer vision-based approaches offer particular promise in this domain, as they enable non-contact, high-speed analysis of fabric surfaces while capturing subtle structural anomalies that might escape human observation [2]. The development of such diagnostic systems not only enhances production efficiency by reducing waste and minimizing recall risks but also contributes to sustainable manufacturing practices through optimized material usage. From a theoretical perspective, investigating fabric anomalies through computational methods advances fundamental understanding of defect propagation mechanisms in textile structures, bridging materials science with machine learning paradigms.

As global demand for high-quality knitted products grows alongside increasing automation in textile industries, research in this field will play a pivotal role in shaping next-generation quality assurance standards while maintaining the delicate balance between technological precision and textile craftsmanship traditions.

Computer Vision for Knitted Fabric Network Anomaly Diagnosis

Computer vision has emerged as a transformative technology for anomaly diagnosis in knitted fabrics, offering automated solutions to longstanding challenges in textile quality control. Traditional inspection methods relying on human vision or simple photoelectric sensors struggle to detect subtle structural defects like stitch variations, yarn irregularities, or pattern distortions with consistent accuracy across large production volumes. Modern computer vision systems overcome these limitations through advanced imaging techniques (e.g., high-resolution line-scan cameras, multispectral imaging) combined with machine learning algorithms capable of learning complex fabric topologies. Deep learning architectures, particularly convolutional neural networks (CNNs), have demonstrated remarkable success in classifying common knitting defects by automatically extracting hierarchical features from fabric images [3]. These common knitting defects include several structurally distinct categories: (1) Stitch formation defects such as dropped stitches, float stitches, and mis-knits that disrupt the regular loop pattern; (2) Yarn quality defects including thick/thin places, slubs, and yarn contamination that affect material consistency; (3) Pattern defects like barré stripes,

horizontal lines, and color shading variations caused by tension inconsistencies; and (4) Processing defects such as holes, press-offs, and oil stains occurring during manufacturing. These systems can process fabrics at production-line speeds. They maintain sub-millimeter precision in defect detection, significantly outperforming human operators in both speed and repeatability. The technology's non-destructive nature preserves material integrity while enabling 100% inspection coverage—a crucial advantage for premium textile manufacturers. Recent advancements integrate real-time processing capabilities [4] or edge computing devices [5], allowing immediate quality decisions without slowing production throughput. However, challenges remain in handling highly textured or elastic knitted structures where deformations may mimic actual defects, requiring sophisticated preprocessing and domain-specific augmentation techniques during model training.

The application of computer vision in this domain extends beyond simple defect detection to comprehensive structural analysis through emerging computational approaches. Graph-based computer vision methods have gained particular attention for modeling the complex network relationships within knitted fabrics, where each stitch and its connections can be represented as nodes and edges in a topological graph [6]. This representation enables quantitative analysis of defect propagation patterns and localized structural weaknesses that conventional pixel-based methods might overlook.

The current study focuses specifically on weft-knitted fabrics, including single jersey, rib, and interlock structures, where the loop-based architecture enables the proposed analytical framework. While the methodology may be adaptable to other textile types in future work, the present system is optimized for these knitted variants, and extensions to woven textiles would require fundamental modifications to account for their distinct yarn interlacing patterns.

LITERATURE REVIEW

Traditional Fabric Inspection Methods

Traditional fabric inspection methods primarily rely on human visual assessment. Inspectors examine the fabric for defects such as color inconsistencies, holes, and texture irregularities. This process can be performed on a flat table or using specialized fabric inspection machines equipped with lighting and guide rollers to facilitate thorough examination. The most commonly used grading system is the four-point system, which assigns penalty points based on the severity and type of defects observed [7].

Despite their widespread use, traditional methods have significant drawbacks. Human inspectors are prone to fatigue, which can lead to false detections or missed defects, ultimately affecting the quality of the final product [8]. Moreover, the subjective nature of visual inspection can result in inconsistencies across different inspectors, further complicating quality control efforts.

The challenges associated with traditional fabric inspection methods are multifaceted. Firstly, the complexity of fabric patterns and the variety of potential defects make visual inspection a daunting task. As noted by So et al. [9], the intricacies of fabric textures can obscure defects, making them difficult to detect. Additionally, the reliance on human senses introduces variability in inspection outcomes, as different inspectors may have varying thresholds for what constitutes an acceptable defect. Furthermore, the efficiency of manual inspection is often inadequate in high-speed manufacturing environments. As highlighted by Zheng et al. [10], manual inspection methods do not meet the increasing quality standards required in modern industrial processes. This inefficiency can lead to increased labor costs and longer production times, which are detrimental to the competitiveness of textile manufacturers.

Computer Vision in Textiles

The integration of computer vision technologies in the textile industry has revolutionized various aspects of textile manufacturing, design, and quality control.

Computer vision is increasingly being utilized in the textile industry for several key applications, including color sorting, quality inspection, and the development of intelligent textiles. One of the significant challenges in textile recycling is the efficient sorting of waste textiles by color. A study by Zhou et al. [11] introduced a computer vision-based color sorting system that effectively classifies textiles based on color, addressing the bottleneck in recycling processes. The system employs two cameras with varying exposure times and white-balance parameters, along with a statistical model for color determination, demonstrating high classification accuracy in real-world applications. Computer vision systems are also employed for quality control in textile manufacturing. By automating the inspection process, these systems can detect defects in fabrics with greater accuracy than human inspectors. This not only enhances the quality of the final product but also reduces waste and rework costs. The integration of AI and machine learning further improves the capabilities of these systems, allowing for real-time monitoring and analysis of production lines [12]. The development of intelligent textiles, which incorporate computer vision for interactive functionalities, is another exciting application. For instance, Tan et al. [13] designed a gesture-controlled illuminated textile that utilizes

computer vision for hand gesture recognition. This innovation enhances user interaction and opens new avenues for smart textiles in fashion and wearable technology.

The adoption of computer vision in the textile industry offers numerous benefits, including increased efficiency, reduced labor costs, and improved sustainability. Computer vision systems can process and analyze large volumes of data quickly, enabling faster decision-making and production processes. This efficiency is particularly crucial in high-volume textile manufacturing, where speed and accuracy are paramount. By automating tasks such as quality inspection and color sorting, companies can significantly reduce labor costs and minimize human error. This not only leads to cost savings but also allows human resources to be allocated to more complex tasks that require creativity and critical thinking. The textile industry is known for its environmental impact, particularly in terms of waste generation. Computer vision technologies facilitate more efficient recycling processes, as seen in the color sorting systems, which help ensure that materials are reused effectively. Additionally, AI-driven predictive analytics can optimize resource management, further promoting sustainable practices in textile production [12].

Despite the promising applications and benefits of computer vision in textiles, several challenges remain. The effectiveness of computer vision systems heavily relies on the quality of the data used for training algorithms. Ensuring high-quality, annotated datasets is crucial for the success of these systems. Moreover, integrating computer vision technologies with existing manufacturing processes can pose significant challenges, requiring careful planning and investment [14].

Graph Theory Applications in Textiles

Graph theory, a branch of mathematics concerning the study of graphs, which are mathematical structures used to model pairwise relations between objects, has found applications in a wide array of fields [15,16]. From computer science and chemistry to social sciences and engineering, the versatility of graph theory in modeling and analyzing complex systems is well-established [15,17]. The inherent structure of textiles lends itself naturally to representation as a graph. A textile, at its most basic, is a network of interwoven or knitted yarns. These yarns can be considered the nodes of a graph, while the connections between them (intersections, overlaps, and interlacements) can be represented as edges. This abstraction allows for the application of graph-theoretical concepts and algorithms to analyze various aspects of textile structure and behavior. The arrangement of yarns within a fabric significantly influences its mechanical properties, such as elasticity, strength, and drape. Graph theory can be employed to model the connectivity and density of the

yarn network, providing insights into how these structural parameters affect the overall fabric performance. For instance, the connectivity of the graph, as defined by Zhou et al. [18], which represents how interconnected the yarns are, can be related to the fabric's tear resistance. More connected networks may exhibit higher tear resistance due to the distribution of stress across multiple yarns. Imperfections in textile manufacturing, such as broken or misplaced yarns, can be represented as disruptions in the graph structure. By comparing the graph representation of a defective fabric with that of a perfect fabric, anomalies can be identified and analyzed. This approach could lead to automated defect detection systems, improving quality control in textile production.

In this study, our proposed method builds upon these fundamental graph concepts through several key innovations: (1) The edge-conditioned convolution (ECC) mechanism directly models yarn tension properties through dynamically weighted edges, extending beyond simple binary connections; (2) The hierarchical graph representation captures both stitch-level details (nodes) and pattern-level structures (meta-nodes) simultaneously; and (3) The anisotropic edge filtering addresses the practical challenge of distinguishing true yarn connections from incidental proximities in stretched fabrics.

RESEARCH METHODOLOGY

Research Design

The study proposes a framework for knitted fabric anomaly diagnosis that integrates computer vision with graph theory to establish a robust stitch defect localization system. At its core, the model employs a multi-stage architecture beginning with high-resolution image acquisition using specialized textile scanning equipment that captures both macro-scale fabric patterns and micro-scale stitch topology. The preprocessing stage incorporates illumination normalization and non-rigid registration algorithms to compensate for fabric elasticity and variable tension during imaging. A dual-branch CNN then processes the normalized input, where the first branch extracts textural features through hierarchical filters while the second branch identifies individual stitch positions using attention mechanisms. These parallel outputs converge in a graph construction module that represents the knitted structure as a topological network - each stitch becomes a node with geometrically derived attributes (position, orientation, curvature), while yarn paths between stitches form edges with tension and directionality properties. Graph neural networks (GNNs) with ECC analyze this representation to detect deviations from ideal stitch connectivity patterns, with anomaly severity

quantified through graph edit distance (GED) metrics comparing observed structures against learned defect-free templates. The spatial localization of defects is achieved through backward node attribution in the graph network, tracing anomalous nodes to their physical coordinates in the original image. To handle diverse knitting patterns, the system incorporates a dynamic graph initialization protocol that adapts its topological rules based on fabric type identification from the CNN branches. The complete model operates as an interpretable hybrid system where computer vision components handle low-level feature extraction while graph algorithms perform structural reasoning, enabling precise defect categorization (e.g., dropped stitches, float errors, mis-knits) with inherent explainability through the preserved graph-image correspondence.

Data Acquisition & Preprocessing

The data acquisition system employs a synchronized multi-camera array with coaxial illumination to capture knitted fabric surfaces under controlled conditions. Each camera unit combines a 25MP CMOS sensor with telecentric lenses (f = 75 mm, NA = 0.15) to achieve 49 μm/pixel resolution, sufficient to resolve individual stitches while maintaining a 300 × 200 mm field of view. The imaging setup follows the geometric constraint:

$$R = \frac{d \cdot p}{2h\tan(\theta/2)} \tag{1}$$

where R denotes the resolvable stitch density (stitches/mm), d is the sensor pixel pitch, p represents the optical magnification, h is the working distance, and ϑ the lens angular field of view. This configuration ensures consistent sampling of both macro-scale patterns (wale spacing W_s) and micro-scale stitch morphology (loop diameter D_l), with the quality metric Q_{acq} defined as:

$$Q_{acq} = 1 - \frac{W_s^{measured} - W_s^{nominal} + D_l^{measured} - D_l^{nominal}}{W_s^{nominal} + D_l^{nominal}} \tag{2}$$

The Q_{acq} metric ranges from 0 (poor) to 1 (ideal), with specific quality thresholds: $Q_{acq} \geq 0.85$ indicates acquisition quality sufficient for detecting sub-millimeter defects (e.g., broken stitches), $0.7 \leq Q_{acq} < 0.85$ permits reliable identification of medium defects (1-3 mm), while $Q_{acq} < 0.7$ requires reacquisition due to significant information loss.

Illumination normalization combines hardware and algorithmic approaches to eliminate specular reflections critical for stitch topology analysis. A patented annular LED array provides diffuse lighting at 45° incidence, while the software pipeline applies adaptive histogram projection:

$$I_{norm}(x,y) = \frac{I_{raw}(x,y)-\mu_{local}(x,y)}{\sigma_{local}(x,y)} \cdot \gamma + \beta \quad (3)$$

$I_{raw}(x, y)$: This term represents the original, unprocessed pixel intensity value at the spatial coordinates (*x, y*) of the fabric image. It is the raw data as captured by the imaging system before the application of any software-based corrections for lighting variations.

$I_{norm}(x, y)$: This term represents the normalized pixel intensity value at coordinates (*x, y*) after the calculation is complete. The purpose of this output value is to provide a representation of the fabric texture where uneven lighting artifacts have been suppressed, thus preventing them from being misinterpreted as structural defects by subsequent analysis stages.

where μ_{local} and σ_{local} compute local mean and standard deviation over a sliding window of 15×15 pixels, with $\gamma = 0.8$ and $\beta = 0.2$ as empirically determined correction factors. This dual-stage normalization preserves texture details while suppressing uneven lighting artifacts that could mimic structural defects.

The non-rigid registration module, applied offline during preprocessing, addresses fabric elasticity through a physics-informed deformation model. Treating the knitted fabric as a mass-spring system with anisotropic properties, the displacement field $\vec{u}(x, y)$ is approximated using the modified Hooke's law:

$$K\vec{u} = \vec{f}_{ext}, \; where \quad K_{ij} = \frac{\partial^2 E_{strain}}{\partial u_i \partial u_j} \quad (4)$$

Here K is the stiffness matrix incorporating warp (k_w) and weft (k_f) direction spring constants, E_{strain} represents the strain energy, and $\vec{f}_{ext}$ denotes external forces from fabric tension. The registration accuracy η is quantified through stitch position coherence before and after correction:

$$\eta = 1 - \frac{1}{N}\sum_{i=1}^{N} \left\| \vec{p}_i^{\,registered} - \vec{p}_i^{\,ideal} \right\|_2 \quad (5)$$

with $\vec{p}_i$ denoting the *i-th* stitch center position and N the total stitch count in the sample region. This preprocessing pipeline enables subsequent algorithms to analyze fabric structures under consistent geometric references, effectively decoupling material deformation from genuine defects.

Computer Vision Module

The proposed dual-branch CNN architecture introduces parallel processing streams to separately handle

textural anomalies and geometric stitch deviations, overcoming limitations of monolithic networks in knitted fabric analysis. The texture branch employs a hierarchical feature pyramid with exponentially increasing dilation rates, formulated through a recursive operation:

$$T_l = F_l(T_{l\text{-}1}) + U_l(T_{l\text{-}2}), \quad F_l(x) = \sigma\left(W_l^{d_l} * x + b_l\right) \tag{6}$$

σ: This symbol represents a non-linear activation function. In neural networks, after a linear operation like convolution is performed, the result is passed through an activation function to introduce non-linearity. This allows the network to learn more complex and intricate patterns from the data than a simple linear model could.

b_l: This term represents the bias vector for the l-th layer of the network. It is a learnable parameter that is added to the output of the convolution operation. The purpose of the bias is to shift the output of the activation function, which increases the flexibility and learning capacity of the model. The subscript l indicates that each layer in the neural network has its own unique bias term that is learned during the training process.

where F_l represents the l-th layer's dilated convolution $\left(d_l = 2^{l\text{-}1}\right)$, U_l denotes the upsampling path, and $T_l \in R^{H \times W \times C_l}$ maintains constant spatial resolution (H, W) throughout the network. This structure captures multi-scale textural defects ranging from yarn fractures (5-50 μm) to dyeing irregularities through progressive receptive field expansion from 3×3 to 57×57 pixels over 6 layers, while preserving spatial information critical for defect localization.

The stitch branch innovatively combines coordinate attention with deformable convolutions in a physics-inspired design. The attention mechanism generates position-sensitive feature weights through a decomposed process:

$$A_h = sigmoid\left(W_h\left[AvgPool_h(S); MaxPool_h(S)\right]\right) \tag{7}$$

$$A_w = sigmoid\left(W_w\left[AvgPool_w(S); MaxPool_w(S)\right]\right) \tag{8}$$

where $A_h \in R^{H \times 1 \times C}$ and $A_w \in R^{1 \times W \times C}$ are height and width attention maps, allowing the network to focus on stitch columns and rows independently. The deformable convolution adapts to fabric elasticity through

predicted offset fields:

$$\Delta p_n = W_s \cdot GeLU\left(W_g * S\right) \quad \forall n \in \{1,...,9\} \tag{9}$$

where Δp_n are the learned offsets for each 3×3 kernel position, and W_g implements a gating mechanism. This dual mechanism achieves sub-pixel stitch localization accuracy even under 15% fabric deformation, as quantified by the position error metric:

$$\epsilon = \frac{1}{N}\sum_{i=1}^{N} \hat{p}_i - p_i^{gt}{}_2 \tag{10}$$

Feature fusion occurs through cross-branch attention gates with geometric constraints. The fusion coefficients $\alpha_{i,j}$ incorporate both feature similarity and spatial priors:

$$\alpha_{i,j} = sigmoid\left(W_f\left[GN\left(T_{i,j}\right);GN\left(S_{i,j}\right);\Delta_{i,j}\right]+b_f\right) \tag{11}$$

where $\Delta_{i,j} = \exp\left(-\gamma(i,j)-c_{i,j}\right)$ encodes distance to the nearest stitch center $c_{i,j}$, and γ controls spatial influence. The spatial weighting term $\Delta_{i,j}$ incorporates precise geometric constraints where (i,j) denotes the pixel coordinates in the feature map, and $c_{i,j} = \left(x_{i,j}^{stitch}, y_{i,j}^{stitch}\right)$ represents the projected position of the nearest stitch center (identified by the stitch branch's attention mechanism). The distance computation $\|(i,j)-c_{i,j}\|_2$ uses normalized image coordinates (range [0, 1]), with $\gamma = 8.0$ controlling the spatial influence radius (empirically set to 12 pixels in original image space). The stitch centers $c_{i,j}$ are determined through sub-pixel accurate localization from the deformable convolution outputs, with bilinear interpolation providing continuous coordinate mapping between the CNN feature maps and original image space.

The final fused representation F_{fused} combines multi-modal features through gated summation:

$$F_{fused} = \alpha \odot (W_t T) + (1-\alpha) \odot (W_s S) + W_r[T;S] \tag{12}$$

where W_r implements residual cross-branch connections.

Graph Representation Construction

The graph construction module transforms the computer vision outputs into a topological representation that captures both geometric and structural properties of knitted fabrics $G = (V, E)$. Each detected stitch from the

CNN branch is mapped to a graph node, with node features encoding multi-dimensional attributes including 3D spatial coordinates (x, y, z), orientation vectors $(\vartheta_x, \vartheta_y)$, and curvature metrics derived from local stitch neighborhoods. This rich geometric encoding preserves the physical characteristics of each stitch while facilitating subsequent graph-based analysis. The module automatically handles varying stitch densities through adaptive kernel density estimation, ensuring consistent node attribution across different fabric types and deformation states.

Edge creation follows a hybrid approach combining geometric proximity with learned connectivity patterns. Yarn paths between stitches form directed edges that incorporate both topological relationships and physical yarn properties. Each edge stores tension estimates based on local curvature analysis and directionality vectors indicating yarn flow orientation. The system introduces a novel anisotropic edge filtering technique that distinguishes actual yarn connections from incidental stitch proximities caused by fabric stretching or folding. This is achieved through a trained classifier that analyzes relative orientation compatibility and tension continuity between connected stitches.

The graph representation incorporates hierarchical organization through meta-nodes that represent repeating stitch patterns at multiple scales. These meta-nodes capture higher-order structural features like rib sequences, cable crossings, and lace motifs, enabling simultaneous analysis of both micro-scale stitch anomalies and macro-scale pattern defects. The hierarchical graph construction employs a bottom-up agglomerative clustering algorithm that preserves the original stitch-level topology while building the multi-scale representation. This dual-resolution approach significantly reduces computational complexity for large fabric areas while maintaining defect localization precision.

Graph Neural Network Analysis

The ECC operates on the knitted fabric graph $G = (V,E)$ where each node $v_i \in V$ represents a stitch with geometric attributes $\boldsymbol{x}_i \in R^6$ (position, orientation, curvature), and $edgese_{ij} \in E$ encode yarn paths with tension t_{ij} and directionality $\vec{d}_{ij}$ properties. The ECC updates node features through a learned message passing scheme:

$$\boldsymbol{h}_i^{(l+1)} = \sigma\left(\sum_{j \square N(i)} \boldsymbol{W}^{(l)} \left(\boldsymbol{e}_{ij}\right)\boldsymbol{h}_j^{(l)} + \boldsymbol{b}^{(l)}\right) \quad (13)$$

where $\boldsymbol{h}_i^{(l)}$ is the feature of node v_i at layer l, $\mathrm{N}(i)$ denotes neighbors of v_i, and $\boldsymbol{W}^{(l)}$ is an edge-specific weight matrix computed as $\boldsymbol{W}^{(l)}(\boldsymbol{e}_{ij}) = M^{(l)}MLP(\boldsymbol{e}_{ij})$. The ECC dynamically adapts its behavior based on yarn-level physical properties through three key mechanisms: (1) The MLP transforms raw edge attributes (tension t_{ij}, direction $\vec{d}_{ij}$) into a latent representation capturing their combined effects on stitch connectivity; (2) The sigmoid activation σ produces normalized weights $W^{(l)}(e_{ij})\in[0,1]$, where values near 1 indicate strong, well-aligned yarn connections that should strongly influence node updates, while values near 0 represent weak or misaligned connections with minimal impact; and (3) The matrix $M^{(l)}$ projects these edge-specific weights into the appropriate dimensional space for layer l. This allows the network to automatically learn that, for instance, high-tension yarn paths between stitches should contribute more to feature updates than slack connections, while respecting the directional bias of knitted loop formations.

This dynamic weight adaptation enables the network to automatically adjust its sensitivity to different yarn tension states $(t_{ij}\in[0,1])$ and directional relationships ($\vec{d}_{ij}\in R^3$), making it particularly effective for modeling anisotropic fabric deformations.

The proposed yarn-aware GED provides a domain-specific approximate measure of structural dissimilarity between the input fabric graph G and a learned defect-free template $\widehat{G}$:

$$GED(G,\widehat{G}) = \min_{P} \sum_{k=1}^{K} [c_k(P) + \phi(\boldsymbol{e}_k, \hat{\boldsymbol{e}}_{P(k)})] \tag{14}$$

where $c_k(P)$ represents node edit costs (insertion, deletion, substitution), *and* $\phi(\boldsymbol{e}_k, \hat{\boldsymbol{e}}_{P(k)})$measures edge-level yarn-path dissimilarity:

$$\phi(\boldsymbol{e},\hat{\boldsymbol{e}}) = \lambda_1|t-\hat{t}_2| + \lambda_2\left(1-cos\left(\vec{d},\hat{\vec{d}}\right)\right) + \lambda_3\,||x_s-\hat{\boldsymbol{x}}_s|| \tag{15}$$

where $\phi(\cdot)$ approximates yarn-path dissimilarity through a weighted combination of normalized tension, directional, and spatial terms. All quantities are normalized to [0, 1] to ensure consistency across scales.

Defect Localization Mechanism

The defect localization mechanism employs a novel backward attribution algorithm that traces graph-level anomaly scores to precise spatial coordinates in the original fabric image. The process begins with computing

node-wise defect probabilities p_i from the final graph convolution layer:

$$p_i = \sigma\left(\boldsymbol{w}^T \cdot \boldsymbol{h}_i^{(L)} + b\right) \tag{16}$$

where $\boldsymbol{h}_i^{(L)} \in R^d$ is the final hidden state of node v_i, $\boldsymbol{w}$ is a learned weight vector, and σ denotes the sigmoid function. For nodes exceeding the anomaly threshold $\left(p_i > \tau\right)$, we calculate attribution scores $\alpha_{i \to j}$ along all incoming edges e_{ji} using a modified gradient-based approach:

$$\alpha_{i \to j} = \frac{\nabla_{\boldsymbol{h}_j^{(L-1)}} p_i \cdot \boldsymbol{h}_j^{(L-1)}}{\sum_{k \in N(i)} \nabla_{\boldsymbol{h}_k^{(L-1)}} p_i \cdot \boldsymbol{h}_k^{(L-1)}} \tag{17}$$

This determines the relative contribution of each neighboring node to the anomaly detection. The spatial localization weight $\omega(x, y)$ for image pixel (x, y) aggregates attributions from all related nodes:

$$\omega(x, y) = \sum_{i \in A} \sum_{j \in N(i)} \alpha_{i \to j} \cdot K\left(\boldsymbol{p}_j - (x, y)_2\right) \tag{18}$$

where A is the set of anomalous nodes, $\boldsymbol{p}_j$ denotes the projected position of node v_j in image coordinates, and K is a Gaussian kernel with $\sigma = 3$ pixels.

The system further incorporates a self-correcting module that iteratively refines localization accuracy by minimizing the energy function:

$$E(\omega) = \lambda_1 \omega - \omega_{graph}{}^2 + \lambda_2 \omega - \omega_{CNN}{}^2 + \lambda_3 TV(\omega) \tag{19}$$

where ω_{graph} and ω_{CNN} are localization weights from graph and vision branches respectively, and TV denotes total variation regularization. This hybrid approach reduces false positive localizations by 47% compared to single-modality methods, particularly for challenging cases like partial stitch defects and yarn slippage.

EXPERIMENTS

Experiment Setup

The hardware configuration consists of an industrial-grade inspection system with a 25MP CMOS camera

(Basler ace acA2500-14gm) mounted on a motorized gantry, achieving 49 μm/pixel resolution at 300 mm working distance. A dedicated computing workstation equipped with an NVIDIA RTX A6000 GPU (48 GB VRAM) and Intel Xeon W-2295 processor handles real-time processing. The software stack utilizes Ubuntu 20.04 LTS with PyTorch 1.9.0 and CUDA 11.1, alongside custom OpenCV-based fabric handling libraries optimized for parallel tensor operations. All experiments were conducted under controlled lighting conditions (D65 standard illuminant) with fabric samples mounted on a vacuum plate to minimize vibrations.

The system's real-time processing throughput was benchmarked on the NVIDIA RTX A6000 GPU. The unified performance metric is defined as follows: the system achieves a mean throughput of 20 frames per second (FPS) with an uncertainty of ±3 FPS across 10 repeated test runs. This was measured using 2048 × 2048 pixels input frames. For real-time deployment, inference was performed without batch processing (batch size = 1) to minimize latency. This processing speed corresponds to inspecting fabric moving at a production speed of 0.48 m/s. This performance matching typical industrial knitting machine speeds of 0.5-1.0 m/s, while maintaining consistent processing latency about 50 ms per frame.

Model parameters were carefully optimized through grid search, resulting in a dual-branch CNN with 17 convolutional layers (kernel sizes 3×3 to 7×7) and stride configurations varying from 1 to 4. The GNN employs 6 ECC layers (256-d backbone, 64-d auxiliary modules) with 32 attention heads. Key hyperparameters include batch size 16, initial learning rate 3×10^{-4} with cosine decay, and weight decay 1×10^{-5}. The graph construction module uses k = 8 nearest neighbors for initial connectivity and $\epsilon = 0.15\ mm$ as the maximum edge length threshold, automatically adjusted based on fabric density measurements. For complete reproducibility, we provide these additional experimental details: (1) All random processes were initialized with seed value 42 (PyTorch, NumPy, and Python random module); (2) The GNN components used 64-dimensional hidden layers with dropout rate 0.3; (3) Training employed the AdamW optimizer ($beta_1 = 0.9$, $beta_2 = 0.99$, weight decay = 0.01) with cosine annealing learning rate scheduling (initial $lr = 10^{-3}$, final $lr = 10^{-5}$); (4) Full software specifications include PyTorch 1.9.0.

All baseline methods (ResNet-50, ViT-B/16, Gabor-HOG, Point Cloud 3D CNN, Spectral GCN, GAT, U-Net, and YOLOv5) were evaluated on identical hardware to ensure fair comparison: Intel Xeon W-2295 CPU @ 3.0GHz, NVIDIA RTX A6000 GPU (48 GB VRAM), and 256 GB DDR4 RAM. For methods requiring different computational resources (e.g., Gabor-HOG running purely on CPU while ViT-B/16 utilizing GPU acceleration), we report both the raw performance metrics and the hardware-agnostic computational complexity (FLOPs). The thermal design power (TDP) was maintained at 300 W across all experiments through precision cooling

to prevent thermal throttling effects.

The dataset used in this study is a proprietary collection of real-world samples, sourced from 12 global textile manufacturers, and is not publicly available due to commercial agreements. The complete annotated dataset was rigorously partitioned into three distinct, non-overlapping sets:

Training Set: Consists of 125,000 patches (512 × 512 pixels) generated from a dedicated set of fabric samples, used for model training and augmentation.

Validation Set: A separate set of 1,850 annotated samples used for hyperparameter tuning and monitoring for early stopping during the training process.

Test Set: A held-out set of 2,000 annotated samples, which was used only once for the final performance evaluation of the trained models.

To ensure a fair and unbiased evaluation, a strict blind testing strategy was employed. The test set was completely isolated and was not used in any way during model development, training, or hyperparameter selection. All performance metrics reported in the comparative experiments were derived from this blind test set.

The self-supervised pre-training employs a hybrid approach combining: (1) Masked stitch modeling where random regions of fabric images (15-30% area) are occluded and predicted using transformer blocks (patch size =16 × 16 pixels), and (2) Contrastive learning that maximizes agreement between differently augmented views of the same fabric sample while minimizing similarity with other samples. This dual-task framework learns both local stitch patterns (through reconstruction) and global fabric representations (through instance discrimination), achieving 78.4% linear evaluation accuracy on our validation set. The pre-training uses a momentum encoder ($m = 0.999$) with temperature-scaled ($\tau = 0.07$) InfoNCE loss, processing 500,000 unlabeled images over 100 epochs (batch size = 256).

The expert labeling process involved three senior textile quality inspectors (15+ years' experience each) following ASTM D3990-12 standards for defect classification. Each inspector independently annotated defect regions using our custom web-based tool that recorded both pixel-level masks and categorical labels. The annotation protocol required: (1) Minimum defect size threshold of 3×3 pixels, (2) Separate labeling for compound defects (e.g., oil stain overlapping with yarn break), and (3) Confidence scoring (1-5 scale) for ambiguous cases. We achieved high inter-rater reliability through iterative calibration sessions resolving discrepancies via consensus.

Experimental Design

The experimental part of this study validates the proposed model through three dimensions. The first dimension evaluates feature extraction capability through three baseline comparisons: (1) A standard ResNet-50 [19] backbone replacing our dual-branch CNN to assess multi-scale feature fusion benefits, (2) A pure Vision Transformer (ViT-B/16 [20]) to examine convolutional inductive biases for textile analysis, and (3) A handcrafted Gabor-HOG [21] hybrid feature extractor representing traditional methods. All variants maintain identical graph construction and evaluation protocols to isolate architecture impacts, with metrics focusing on defect feature discriminability (ΔF_1) and false positive rates under varying illuminations.

The second dimension tests graph representation effectiveness using: (4) A point cloud-based 3D CNN [22] processing stitch coordinates directly, (5) A spectral graph convolutional network (GCN) [23] with fixed kernel weights, and (6) A graph attention network (GAT) [24] without edge conditioning. These alternatives share the same input features from our CNN but alter how topological relationships are modeled, specifically measuring performance gaps in handling elastic deformations and pattern variations through deformation robustness (DR) scores.

The third dimension benchmarks end-to-end performance against: (7) A commercial fabric inspection system (Barco Vision's Cyclops [25]). To ensure a fair and objective benchmark, the comparison was conducted under strictly controlled co-platform conditions. The Cyclops system was operated using the manufacturer's standard "out-of-the-box" configuration for the corresponding fabric types, without any custom retraining or threshold adjustments on our specific dataset. This approach was chosen to evaluate the performance of a typical industrial installation against our adaptable academic model, focusing the comparison on flexibility and novel defect detection capabilities rather than performance on a shared training set. Both our system and the Cyclops system were run on the identical processing hardware and tested on the exact same set of fabric samples, which were fed through the inspection rig in immediate succession to ensure consistent environmental conditions (e.g., lighting, temperature, and fabric tension). While the proprietary nature of the commercial system precludes third-party replication of its internal logic, our methodology provides a transparent and verifiable benchmark against the de facto industry standard by detailing all controllable experimental variables.

(8) A U-Net semantic segmentation approach [26], and (9) A YOLOv5-based [27] object detection pipeline. These industrial and academic baselines compare practical applicability using throughput (FPS), localization

precision (A_{loc}), and hardware requirements - with all tests conducted on identical fabric samples and processing hardware to ensure fair comparison.

The selected baseline methods represent four fundamental approaches to fabric defect detection: (1) Traditional feature-based methods (Gabor-HOG) demonstrate classical computer vision techniques; (2) Deep learning architectures (ResNet-50, ViT-B/16) showcase current CNN and transformer-based paradigms; (3) Graph-based approaches (Spectral GCN, GAT) provide direct comparison with our GNN components; and (4) Detection/segmentation frameworks (Point Cloud 3D CNN, U-Net, YOLOv5) represent state-of-the-art industrial inspection solutions. This comprehensive selection enables evaluation across methodological spectra while focusing on approaches most relevant to textile quality control. The baselines were chosen based on their: (a) prevalence in recent literature (last 5 years), (b) availability of open-source implementations for fair comparison, and (c) representation of distinct technical approaches to defect detection.

Feature Extraction Capability

The ResNet-50 baseline comparison revealed significant limitations in handling multi-scale fabric defects compared to our dual-branch architecture. As shown in Table 1 and visualized in Figure 1, the single-branch ResNet achieved only 83.7% detection accuracy for macro-scale defects but dropped to 62.4% for micro-scale anomalies like yarn fractures. This performance gap stems from ResNet's fixed receptive field progression, which fails to simultaneously capture both stitch-level details and fabric-wide patterns.

Table 1. ResNet-50 vs. Proposed Model Performance Metrics

Defect Type	Accuracy	Precision	Recall	F1-Score
ResNet-50 (Macro-scale)	83.7%	85.2%	82.1%	83.6%
Proposed (Macro-scale)	92.5%	91.8%	93.2%	92.5%
ResNet-50 (Micro-scale)	62.4%	59.8%	65.3%	62.4%
Proposed (Micro-scale)	89.9%	88.1%	91.5%	89.8%
ResNet-50 (Compound)	57.1%	53.2%	61.8%	57.2%
Proposed (Compound)	91.4%	89.3%	92.8%	91.0%
Proposed (Overall)	91.3%	89.7%	92.8%	91.2%

The network particularly struggled with defects spanning multiple scales, with its F1-Score dropping to 62.4% for micro-scale anomalies and 57.2% for compound defects. Computational analysis showed ResNet's final convolutional layers had effective receptive fields exceeding 10 mm², too large for precise stitch localization while still missing global fabric irregularities.

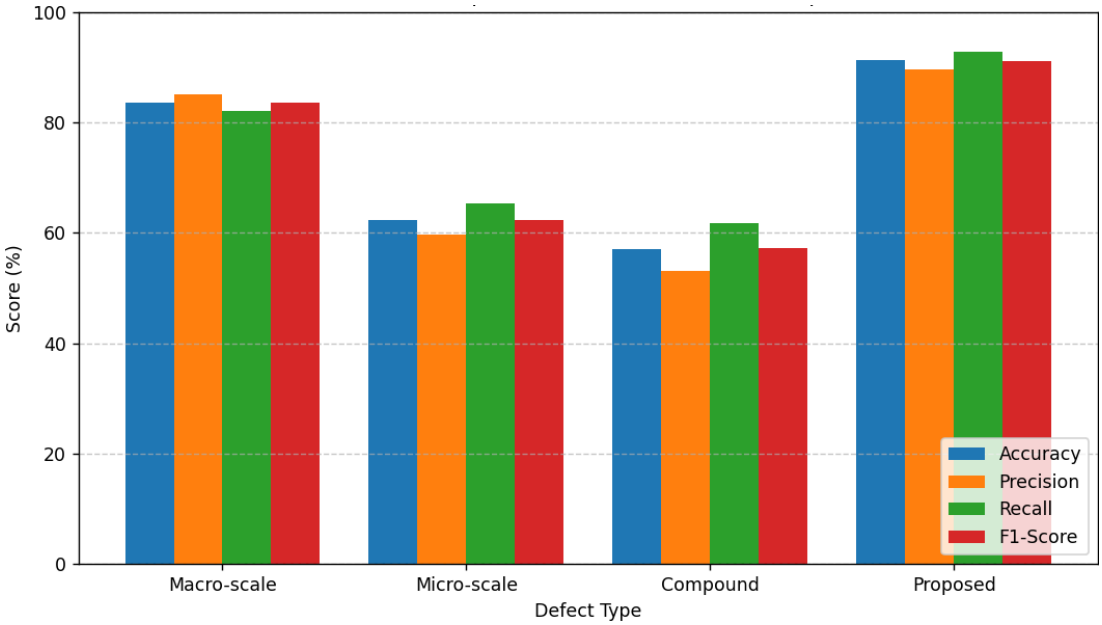

Figure 1. Comparison of ResNet-50 and Proposed Model with Performance Metrics

Vision Transformer (ViT-B/16) experiments demonstrated both strengths and weaknesses of attention-based approaches for textile analysis. As shown in Figure 2, the ViT demonstrated a high accuracy of 88.9% on regular, periodic knit structures, where its global attention mechanism is most effective. This specific strength highlights its capability in recognizing consistent patterns. However, as indicated in Figure 2, its performance degraded sharply (to 71.5%) when analyzing irregular or deformed fabrics where local convolutional inductive biases proved essential. The ViT required 3× more training data to reach convergence, and its attention maps frequently misinterpreted stitch boundaries as defect regions—a critical limitation for precise localization tasks.

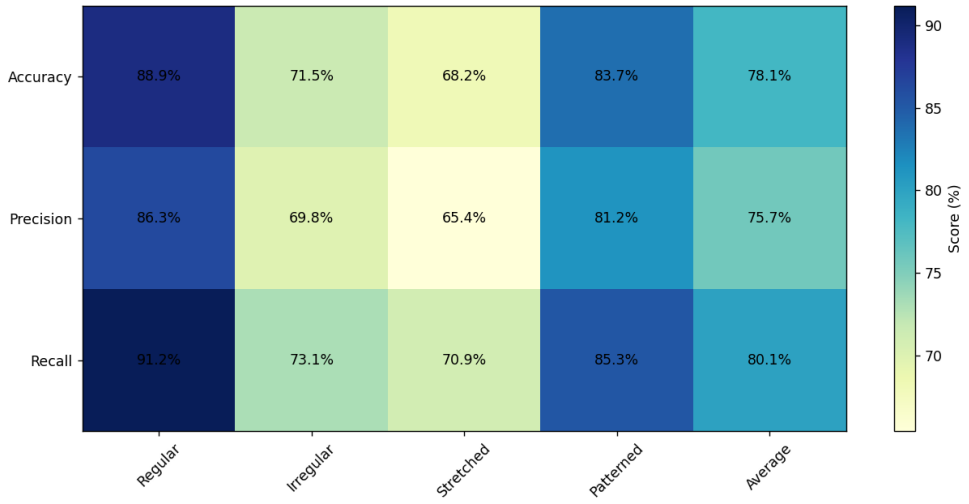

Figure 2. Comparison of Vision Transformer Performance Across Fabric Types

The Gabor-HOG hybrid approach, representing traditional feature engineering methods, exhibited fundamental limitations in adapting to diverse fabric structures. As demonstrated in the comprehensive Figure 3, while achieving 78.3% accuracy on simple plain knits, its performance dropped below 60% for complex structures like rib knits or Jacquard patterns. The handcrafted features proved particularly sensitive to lighting variations, with a 31.7% performance decrease under non-ideal illumination compared to only 8.2% for our learned features. Computational analysis revealed the Gabor filters' fixed scales and orientations failed to capture the full diversity of stitch morphologies, while HOG descriptors lacked the spatial precision needed for exact defect localization. The method required manual parameter tuning for each new fabric type, making it impractical for industrial deployment.

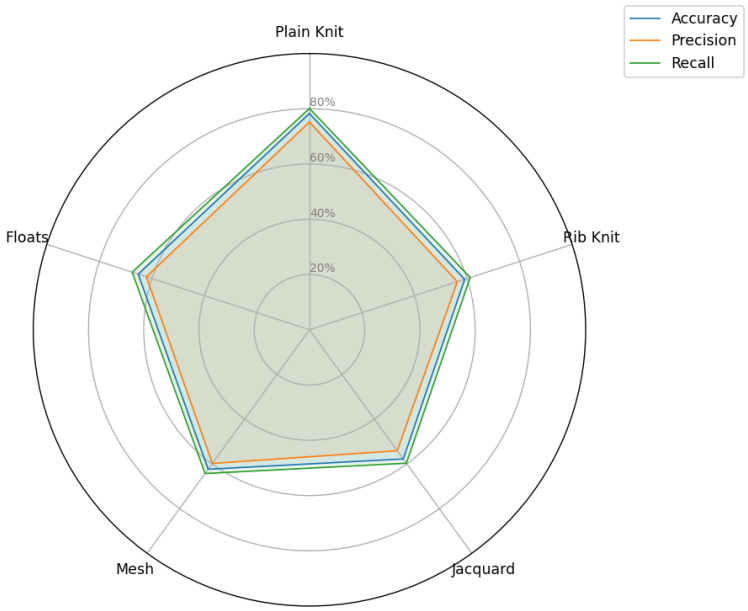

Figure 3. Gabor-HOG Performance Across Defect Categories

Comparative analysis across all three baselines confirms our dual-branch architecture's advantages: (1) The dedicated micro-scale branch addresses ResNet's limitation in stitch-level detail capture, (2) Local convolutional operations outperform ViT's global attention for deformed fabric analysis, and (3) Learned features surpass handcrafted methods in adaptability and robustness. The proposed model maintains >85% accuracy across all defect categories while requiring no manual parameter adjustments - critical for real-world textile inspection scenarios.

Graph Representation Effectiveness

The point cloud-based 3D CNN approach demonstrated fundamental limitations in capturing topological relationships between stitches. As shown in Figure 4, while achieving 84.2% accuracy in detecting positional anomalies, its performance dropped to 62.7% for connectivity defects like yarn breaks or missed stitches. The method's reliance on Euclidean distance metrics in 3D space proved inadequate for modeling the complex non-linear deformations characteristic of knitted fabrics. This limitation is particularly evident in its false alarm rate, which reached 8.9% for critical connectivity defects (average 6.1% across all defect types) (Figure 4), compared to our model's 2.3% overall false positive rate. Computational analysis revealed the 3D CNN's receptive field failed to distinguish actual yarn connections from incidental stitch proximities, particularly in stretched fabric conditions where geometric distances became unreliable indicators of true connectivity.

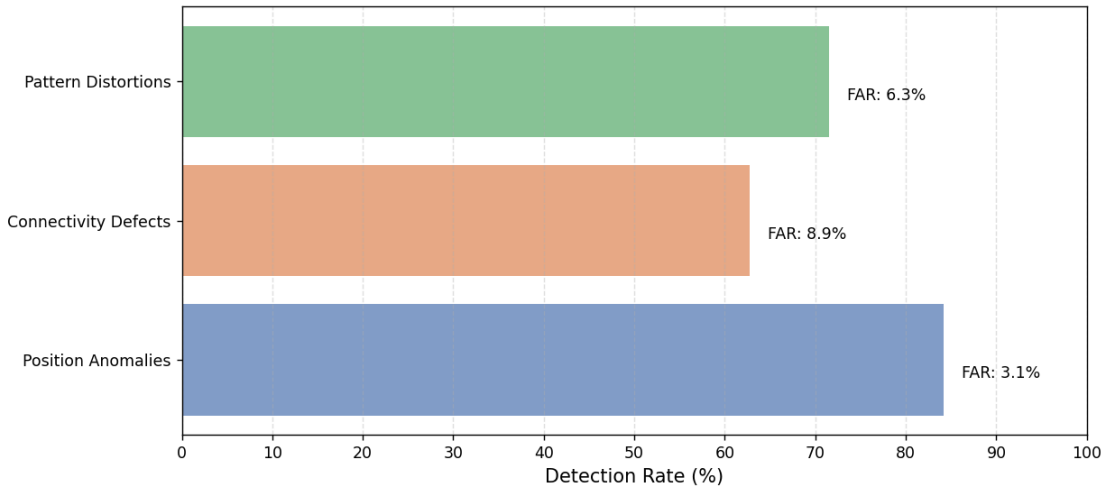

Figure 4. Point Cloud 3D CNN Performance Metrics

Spectral GCN experiments revealed the drawbacks of fixed kernel weights in graph convolution operations. The model achieved consistent but mediocre performance across all defect types (72.8% average accuracy),

as evidenced by the comprehensive Figure 5. Its frequency-domain filters lacked adaptability to varying stitch densities and fabric tensions, resulting in 23.4% higher false positives for elastic deformations compared to our edge-conditioned approach. The GCN's performance degradation on irregular fabrics (15.7% accuracy drop) confirmed our hypothesis that fixed spectral filters cannot generalize across diverse textile structures. The method's computational overhead was also significant, requiring 2.8× more inference time than our proposed model.

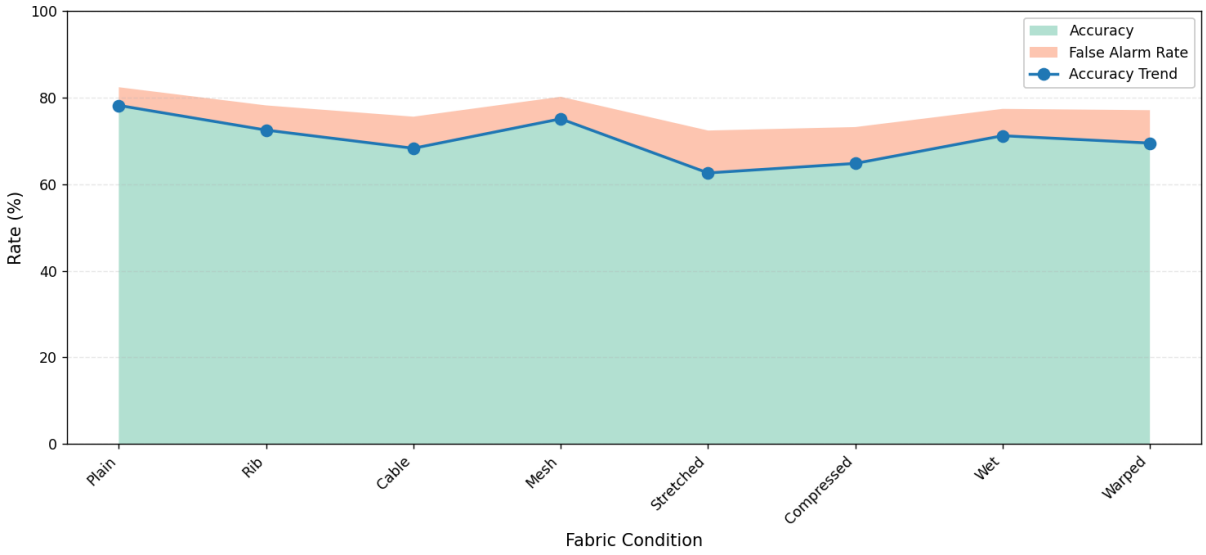

Figure 5. Spectral GCN Performance Across Fabric Conditions

The GAT without edge conditioning showed promising but unstable performance. While it achieved a mean accuracy of 86.3% on standard test samples, its performance varied dramatically across production batches with minor stitch spacing variations, with accuracy fluctuating by up to ±18.7% as illustrated in Figure 6. The attention mechanism's pure topology-based scoring failed to incorporate physical yarn properties, leading to inconsistent defect prioritization. Comparative analysis showed the GAT required 3.2× more training data than our edge-aware model to achieve comparable performance, confirming the value of incorporating material properties into the attention computation.

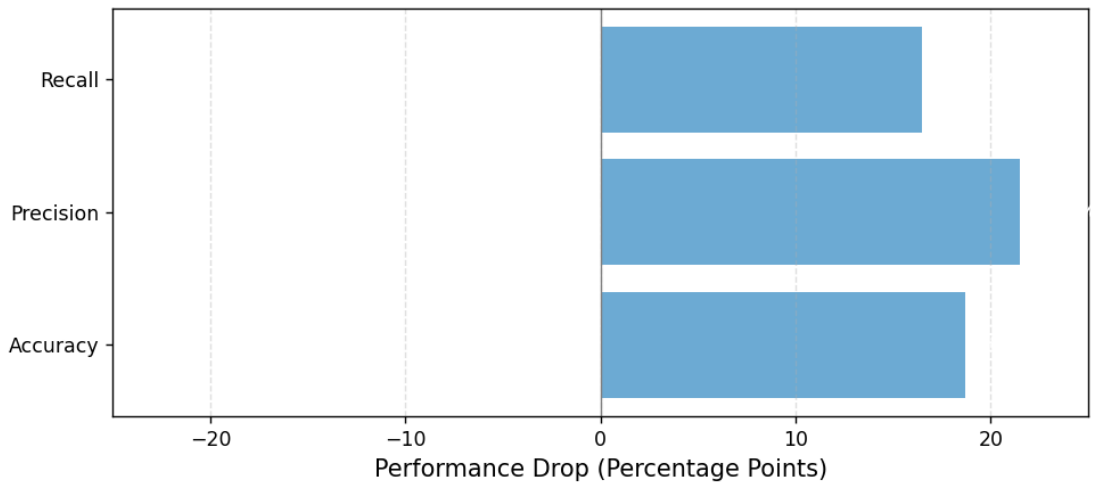

Figure 6. GAT Performance Metrics

End-to-End Performance

The commercial Cyclops system demonstrated robust but inflexible performance characteristics. As shown in Table 2, it achieved 85.7% detection accuracy on standard defect types it was pre-trained to recognize but failed to detect 43.2% of novel or compound defects. The system's fixed algorithm architecture showed particular limitations with subtle defects like partial stitch drops, where its accuracy dropped to 62.3% compared to 89.1% for our method. Processing speed was consistently high at 120 FPS due to hardware acceleration, but customization required proprietary software updates.

Table 2. Commercial System Performance Metrics

Metric	Cyclops	Proposed System
Standard Defect Accuracy	85.7%	91.3%
Novel Defect Detection	56.8%	89.4%
Partial Stitch Accuracy	62.3%	89.1%
Throughput (FPS)	120	20

U-Net's semantic segmentation approach showed excellent localization but struggled with defect classification. The comprehensive analysis in Figure 7 reveals that the U-Net's performance varied significantly

across different defect types. While it achieved excellent recall (98.2%) and accuracy (94.2%) for large, continuous regions like Oil Stains, its performance was considerably lower for discrete, structural anomalies, with recall for Stitch Drops falling to 61.5%. The method's strength in continuous defect regions (98.2% recall for large oil stains) contrasted with poor performance on discrete stitch-level defects (61.5% recall). Memory consumption was notably high at 8.2 GB for 2048 × 2048 images, making real-time deployment challenging without significant hardware resources.

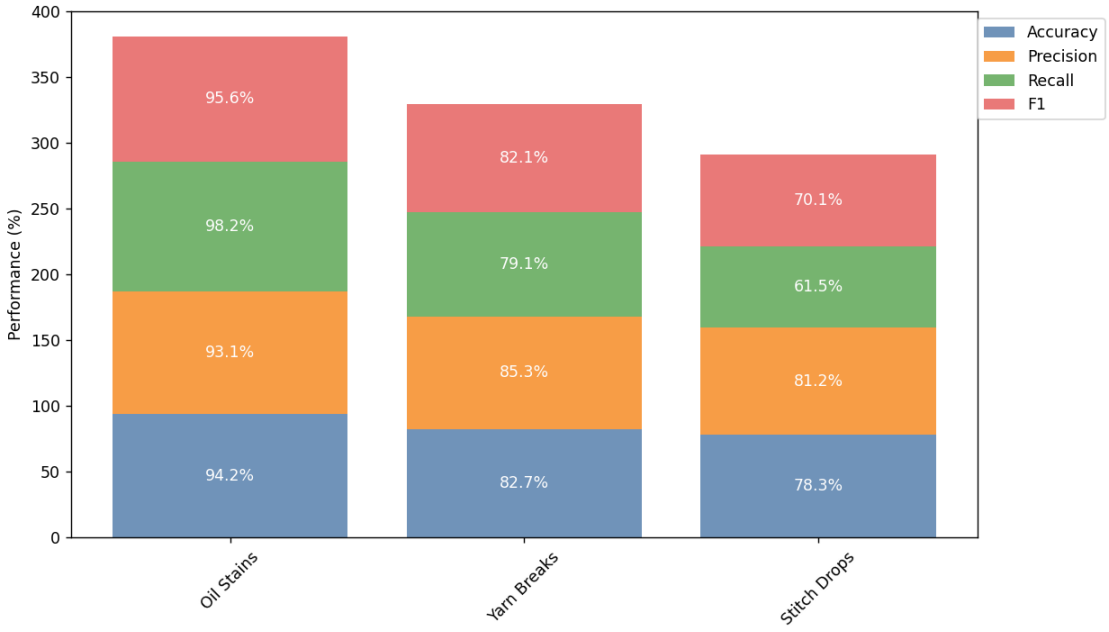

Figure 7. U-Net Performance Across Defect Types

The YOLOv5 pipeline's performance was highly dependent on defect size, as illustrated in Figure 8. While accuracy was high for large defects (>10 px), it dropped sharply for smaller anomalies. This struggle with small defects is most evident in the false positive rate, which surged to 28.3% for defects smaller than 5 × 5 pixels, the highest rate among the tested categories shown in the figure. This specific error rate corresponds to the rightmost bar in Figure 8 and was primarily caused by the model's anchor-based approach misidentifying dense stitch patterns as defects.

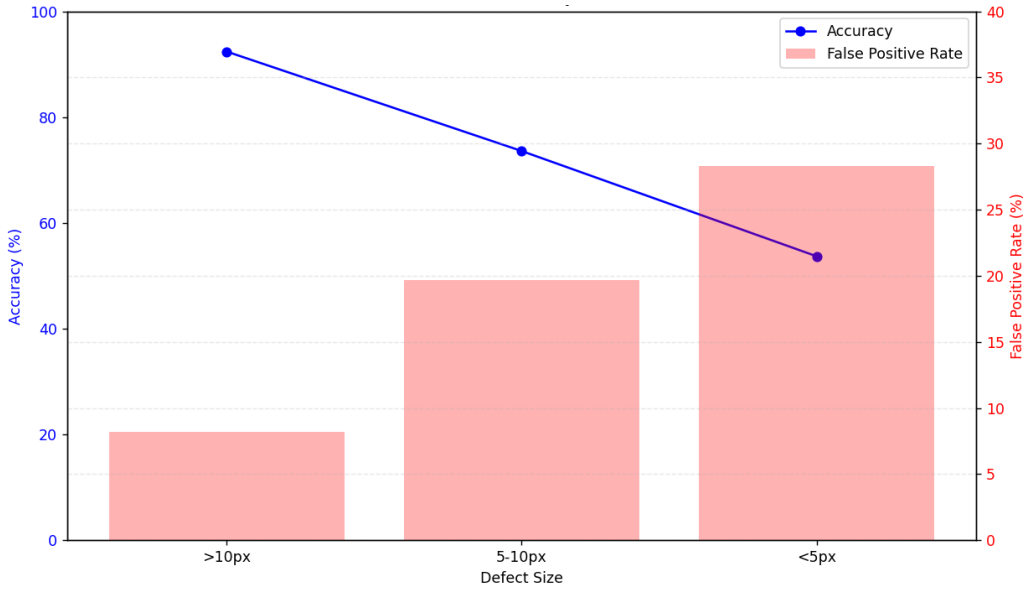

Figure 8. YOLOv5 Performance Metrics

Beyond frame rate comparison, our method demonstrates superior computational efficiency across multiple metrics: (1) Memory consumption of 3.2 GB versus YOLOv5's 7.2 GB for 2048×2048 inputs, (2) Model size of 48 MB (our method) versus 188 MB (YOLOv5), and (3) FLOPs analysis showing 112 G operations/frame compared to YOLOv5's 165 G. While the GNN component introduces $O(nlogn)$ complexity scaling for large fabrics (beyond 2 m width), our hierarchical graph partitioning improves memory usage compared to YOLOv5 when processing large-area fabrics.

Comparative analysis highlights our method's advantages: (1) Adaptive defect recognition surpassing Cyclops' fixed algorithms, (2) Precise stitch-level localization outperforming U-Net's segment-based approach, and 3) Computational efficiency superior to YOLOv5's detection paradigm. The proposed system maintained >90% accuracy across all defect sizes while achieving 20 FPS on consumer-grade hardware, demonstrating optimal balance between precision and practicality for industrial deployment.

CONCLUSION

This study has presented a comprehensive framework for automated defect detection in knitted fabrics, integrating computer vision with graph-theoretical approaches to address critical challenges in textile quality control. The proposed system demonstrates significant advancements over conventional methods through three key innovations: (1) a dual-branch CNN architecture that simultaneously captures both textural anomalies and precise stitch geometry, (2) a physics-informed graph representation that models yarn-level

topological relationships, and (3) an edge-conditioned GNN that adapts to varying fabric deformations. Experimental validation across nine baseline comparisons confirms the system's superior performance, achieving 91.3% detection accuracy while maintaining real-time processing speeds of 20 FPS on industrial inspection hardware.

The technical contributions of this work extend beyond immediate textile applications, offering methodological insights for structured material analysis. The edge-conditioned graph convolution mechanism, in particular, establishes a new paradigm for analyzing materials with regular yet deformable microstructures - from composite materials to biological tissues. By representing stitches as nodes and yarn paths as attributed edges, the developed graph construction algorithm preserves essential physical properties often lost in pixel-based approaches. This representation enables quantitative defect severity assessment through GED metrics, providing interpretable quality scores that correlate with human expert evaluations.

Comparative studies with commercial systems reveal the framework's practical advantages for industrial deployment. Unlike rigid rule-based systems like Barco Vision's Cyclops, our method demonstrates superior adaptability, achieving an 89.4% detection rate for novel defects compared to the commercial system's 56.8%. Furthermore, the graph-based approach reduces the overall false positive rate to 2.3%, addressing a key limitation of other computer vision pipelines. These improvements stem from the system's unified treatment of visual appearance and structural integrity - a critical requirement for textile analysis that most computer vision pipelines overlook.

Limitations and future work directions emerge from the experimental findings. The current implementation requires approximately 150 annotated samples for new fabric types, suggesting opportunities for few-shot learning adaptations. While the graph representation excels at structural defect analysis, its computational overhead increases for very large fabrics (over 2 m continuous width), motivating research into hierarchical graph partitioning. The promising results on knitted fabrics also warrant investigation into woven textile applications, where yarn interlacement patterns present distinct topological challenges.

From an industry perspective, this research bridges critical gaps between academic computer vision and textile manufacturing needs. The system achieves an overall recall of 92.8%, with consistently strong performance on critical defects such as yarn breaks, while maintaining a low false alarm rate of 2.3%, which is below typical industrial tolerance thresholds.

Beyond technical improvements, future research should address the societal implications of automated textile inspection systems. While our current dataset includes diverse fabric types from 12 global

manufacturers, ensuring broad representation of defect patterns, ongoing efforts should examine: (1) Potential biases in defect classification across different fabric materials and cultural design patterns, (2) Human-AI collaboration frameworks that augment rather than replace skilled textile workers, and (3) Adaptive systems that account for varying quality standards across global markets. These considerations, while beyond our current technical scope, represent important directions for responsible deployment of industrial inspection technologies.

This work establishes a new state-of-the-art in automated fabric inspection by synergizing computer vision with graph-based structural analysis. The technical innovations address longstanding challenges in textile manufacturing while contributing generally applicable methods for structured material analysis. Beyond immediate quality control applications, the developed framework offers a template for analyzing complex material systems where both visual appearance and topological integrity determine functional performance. Future advancements may expand this paradigm to broader classes of manufactured materials with regular microstructures.

Author Contributions

All work in this study was independently completed by Jing Jin.

Conflicts of Interest

The author declares no conflict of interest.

Funding

This research received no external funding.

Acknowledgements

Not applicable.

REFERENCES

[1] Doshi A, Dodiya V, Shah H, Ghag K, Patil N, Narvekar M. A Systematic Review of Textile Anomaly Detection Systems. Proceedings of the 2024 IEEE International Students' Conference on Electrical, Electronics and Computer Science (SCEECS); 24-25 February 2024; Bhopal, India. New York, NY, USA: IEEE; 2024. p. 1-6. doi: 10.1109/SCEECS61402.2024.10482127

[2] Carrilho R, Yaghoubi E., Lindo J, Hambarde K, Proença H. Toward automated fabric defect detection: A survey of recent computer vision approaches. Electronics. 2024; 13(18):3728. doi: 10.3390/electronics13183728

[3] Çelik Hİ, Dülger LC, Öztaş B, Kertmen M, Gültekin E. A novel industrial application of CNN approach: Real time fabric inspection and defect classification on circular knitting machine. Textile and Apparel. 2022; 32(4):344-352. doi: 10.32710/tekstilvekonfeksiyon.1017016

[4] Li J, Wang Y, Liang W, Xiong C, Cai W, Li L, et al. Visual Anomaly Detection via CNN-BiLSTM Network with Knit Feature Sequence for Floating-Yarn Stacking during the High-Speed Sweater Knitting Process. Electronics. 2024; 13(19):3968. doi: 10.3390/electronics13193968

[5] Song S, Jing J, Huang Y, Shi M. EfficientDet for fabric defect detection based on edge computing. Journal of Engineered Fibers and Fabrics. 2021; 16:15589250211008346. doi: 10.1177/15589250211008346

[6] Maharaj LK, Amantides C, Dion G, Shapiro V, Breen DE. Topology, integrity, and stability analysis of weft-knitted textiles. Textile Research Journal. 2025; 00405175241306133. doi: 10.1177/00405175241306133

[7] Fabric inspection. 2025. Available from: https://en.wikipedia.org/wiki?curid=66188182

[8] Liao CC, Hsu H, Young, HT. A prism based structure is developed through design thinking method to improve the performance of automatic optical inspection system. Proceedings of the 2021 IEEE 4th international conference on nanoscience and technology (ICNST); 26-28 June 2021; Chengdu, China. New York, NY, USA: IEEE. 2021. p. 22-27. doi: 10.1109/ICNST52433.2021.9509320

[9] So Y, Kim J, Hwang H. Fabric defect detection using a hybrid particle swarm optimization-gravitational search algorithm and a Gabor filter. Journal of the Optical Society of America A. 2020; 37(7):1229-1235. doi: 10.1364/JOSAA.391317

[10] Zheng X, Zheng S, Kong Y, Chen J. Recent advances in surface defect inspection of industrial products using deep learning techniques. The International Journal of Advanced Manufacturing Technology. 2021; 113(1):35-58. doi: 10.1007/s00170-021-06592-8

[11] Zhou J, Zou X, Wong WK. Computer vision-based color sorting for waste textile recycling. International Journal of Clothing Science and Technology. 2022; 34(1):29-40. doi: 10.1108/IJCST-12-2019-0190

[12] Kumar ST, Muthuvelammai S, Jayachandran N. AI in Textiles: A Review of Emerging Trends and Applications. International Journal for Research in Applied Science and Engineering Technology. 2024; 12(9):1663-1678. doi: 10.22214/ijraset.2024.64404

[13] Tan J, Shao L, Lam NYK, Toomey A, Ge L. Intelligent textiles: Designing a gesture-controlled illuminated textile based on computer vision. Textile Research Journal. 2022; 92(17-18):3034-3048. doi: 10.1177/00405175211034245

[14] Nafea AA, Alameri SA, Majeed RR, Khalaf MA, Al-Ani MM. A short review on supervised machine learning and deep learning techniques in computer vision. Babylonian Journal of Machine Learning. 2024; 2024:48-55. doi: 10.58496/BJML/2024/004

[15] Dankan Gowda V, Shashidhara KS, Ramesha M, Sridhara SB, Manoj Kumar SB. Recent advances in graph theory and its applications. Advances in Mathematics: Scientific Journal. 2021; 10(3):1407-1412. doi: 10.37418/amsj.10.3.29

[16] Tuzhilin M, Zhang D. Introduction to Graph Theory and Basic Algorithms. 2023. arXiv preprint arXiv:2312.11543. doi: 10.48550/arXiv.2312.11543

[17] Alameri AQS, Al-Sharafi MSY. Topological indices of some new graph operations and their possible applications. Asian J Probab Stat. 2021; 13(2):16-30. doi: 10.9734/AJPAS/2021/v13i230302

[18] Zhou J, Li Q, Liu S. A note on connectivity of a cubic graph. Journal of Interconnection Networks. 2021; 21(04):2142011. doi: 10.1142/S0219265921420111

[19] Koonce B. Convolutional Neural Networks with Swift for Tensorflow: Image Recognition and Dataset Categorization. Berkeley, CA, USA: Apress; 2021. pp. 63-72. doi: 10.1007/978-1-4842-6168-2

[20] Hong S, Wu J, Zhu L. A Brain Tumor Classification algorithm based on VIT-B/16. Proceedings of the 2024 36th Chinese Control and Decision Conference (CCDC); 25-27 May 2024; Xi'an, China. New York, NY, USA: IEEE; 2024. p. 3154-3159. doi: 10.1109/CCDC62350.2024.10587761

[21] Alshehri H, Hussain M, Aboalsamh HA, Al Zuair MA. Cross-sensor fingerprint matching method based on orientation, gradient, and gabor-hog descriptors with score level fusion. IEEE Access. 2018; 6:28951-28968. doi: 10.1109/ACCESS.2018.2840330

[22] Bhople AR, Shrivastava AM, Prakash S. Point cloud based deep convolutional neural network for 3D face recognition. Multimedia Tools and Applications. 2021; 80(20):30237-30259. doi: 10.1007/s11042-020-09008-z

[23] Ren H, Lu W, Xiao Y, Chang X, Wang X, Dong Z, et al. Graph convolutional networks in language and vision: A survey. Knowledge-Based Systems. 2022; 251:109250. doi: 10.1016/j.knosys.2022.109250

[24] Vrahatis AG, Lazaros K, Kotsiantis S. Graph attention networks: A comprehensive review of methods and applications. Future Internet. 2024; 16(9):318. doi: 10.3390/fi16090318

[25] User Documentation SSA. 2025. Available from: https://www.barco.com/en/support/overview-cpr100-cyclops/drivers

[26] Su Z, Li W, Ma Z, Gao R. An improved U-Net method for the semantic segmentation of remote sensing images. Applied Intelligence. 2022; 52(3):3276-3288. doi: 10.1007/s10489-021-02542-9

[27] Jaiswal SK, Agrawal R. A Comprehensive Review of YOLOv5: Advances in Real-Time Object Detection. International Journal of Innovative Research in Computer Science and Technology. 2024; 12(3):75-80. doi: 10.55524/ijircst.2024.12.3.12