Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X Volume 4, Issue 3, 2026

INTELLIGENT IMAGE PROCESSING AND MACHINE LEARNING FOR AUTOMATED FABRIC IDENTIFICATION AND QUALITY INSPECTION: A CRITICAL REVIEW

Urooj Aftab Khilji[*1], Syeda Iqra Hassan[2]

[*1]*Mir Chakar Rind University Sibi Balochistan*
[2]*Software Engineering Department, Sir Syed University of Engineering and Technology PK*
[2]*Multimedia University MY, Universiti Telekom Sdn. Bhd. Research Management, Multimedia University Persiaran Multimedia Cyberjaya, Selangor 63100 Malaysia*

[*1]arooj.aftab24@gmail.com, [2]iqra.hassan@ssuet.edu.pk/syeda.iqra@mmu.edu.my

DOI: https://doi.org/10.5281/zenodo.24126084

Keywords
textile inspection; fabric identification; weave recognition; texture analysis; defect detection; computer vision; CNN; multimodal learning; industrial quality control.

Article History
Received: 20 January 2026
Accepted: 05 March 2026
Published: 19 March 2026

Copyright @Author
Corresponding Author: *
Urooj Aftab Khilji

Abstract

Automated visual inspection has become a critical part of intelligent textile manufacturing, yet existing literature is quite scattered in terms of fabric material recognition, classification of weaves/texture, defect detection, localization and real-time implementation. The earliest methods used texture and color descriptors constructed by hand, while more recent research moved toward the use of convolutional neural networks, transfer learning, multiscale architecture, object detector networks, segmentation, attention-based networks and multimodal representations. The current review focuses on this evolution and not the sameness of all computer vision approaches. In particular, it addresses the problem of separating two related problems: recognition of what a piece of fabric is and recognition of if its surface is of acceptable quality. Woven fabrics' recognition research demonstrates the benefits of learned visual representations, while material identification literature shows the relevance of small-sample and lightweight models. Surface defect detection literature evolves from unsupervised image reconstruction and compact CNNs to increasingly advanced real-time detectors and segmentation algorithms. Multimodal literature suggests that texture, local microstructure and complementing surface descriptors may help in material identification. However, despite significant advances in the field, there are still several limitations: lack of datasets, unbalanced classes, poor control of fabric roll leakage, restricted laboratory conditions, poor consistency in performance measures, lack of cross-dataset evaluation and ignoring of uncertainty/explainability issues. An open direction for future research suggested by this review is a multimodal and robust framework of inspection of fabric identity, texture/weave, surface defects and quality.

1. INTRODUCTION

Fabric inspection in industrial textile manufacturing is not a single, uniform computer-vision task. In practice, a production-line visual inspection system must perform several interconnected analyses before reaching a quality-control decision. It must identify the fabric or material class, analyze weave patterns and surface textures, assess color consistency and surface structure, detect abnormal regions, accurately localize defect boundaries, and ultimately determine whether the material should be accepted, graded, or rejected. Treating all these requirements as one classification problem overlooks important differences in data-acquisition procedures, annotation methods, neural-network architectures, and task-specific evaluation metrics. Recent surveys of

automated fabric-defect recognition commonly divide the overall process into four main stages: image acquisition, preprocessing, feature extraction or selection, and final classification or detection. These studies also emphasize that the effectiveness of such methods is strongly influenced by real-world factory conditions and the availability and quality of datasets [1].

Traditionally, automated textile inspection relied either on manual inspection or on computer-vision pipelines based on handcrafted spatial and spectral features. Over the past decade, however, the field has shifted substantially toward deep learning and automatically learned visual representations. An important early contribution by Mei et al. introduced an unsupervised framework based on a multi-scale convolutional denoising autoencoder. The model learned to reconstruct defect-free fabric patches, allowing defects to be detected and localized as reconstruction anomalies without requiring large collections of pixel-level defect annotations [2]. Building on this direction, subsequent research focused on improving computational efficiency for production-line and edge-device deployment. This resulted in wide-and-compact convolutional neural network (CNN) architectures that reduced model size and inference latency while maintaining effective defect-detection sensitivity [3].

At the same time, the related problem of recognizing and classifying woven-fabric patterns increasingly adopted transfer learning and deep CNN backbones. These approaches demonstrated the ability of learned feature representations to capture the complex and repetitive characteristics of textile textures [4]. More recent developments have extended computational textile analysis beyond basic defect screening to include multi-category material identification [5], multimodal sensing that combines visual information with tactile or spectral data for more accurate textile classification [6], and advanced real-time defect-detection systems designed for high-speed industrial production lines [7].

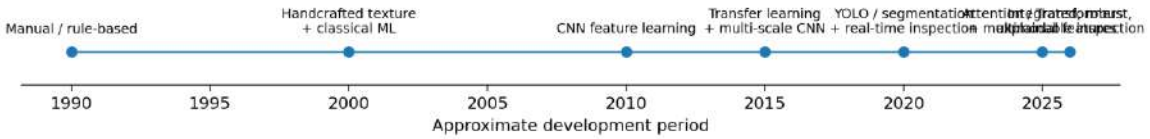


Figure 1. Evolution of automated fabric inspection: from manual/rule-based analysis to integrated multimodal and robust intelligent inspection.

2. Review methodology and analytical framework

The review is organized thematically rather than as a chronological list of studies. Its analytical framework is divided into six main evidence streams: (i) image acquisition and preprocessing, (ii) handcrafted representations of texture and color, (iii) fabric material and weave identification, (iv) defect classification and detection, (v) defect localization, segmentation, and real-time inspection, and (vi) multimodal, robust, and industrially deployable learning approaches. This structure is broadly consistent with recent review literature, while extending the scope beyond defect detection to address the wider challenges of fabric identity and quality assessment.

Table 1. Analytical decomposition of the automated fabric-inspection literature.

Evidence stream	Primary question	Typical methods	Main limitation
Acquisition	Can the surface be captured consistently?	Industrial/controlled imaging, ROI extraction	Lighting, background, scale and roll variation
Texture/weave	What surface structure is present?	GLCM, LBP, CNN, transfer learning	Confusion among visually similar structures
Material ID	What textile/material category is represented?	DenseNet variants, CNNs	Small datasets and limited material diversity
Defect detection	Is there an abnormality and what type?	CNNs, YOLO-family detectors	Small/subtle defects and domain shift
Localization	Where is the defect and what area does it occupy?	Segmentation, pixel masks, repeated-pattern analysis	Annotation cost and boundary ambiguity
Integrated inspection	Can all observations support one reliable quality decision?	Feature fusion, multi-task learning	Fragmented datasets and lack of end-to-end validation

3. Image acquisition and preprocessing

Image acquisition is often regarded as a preliminary engineering consideration, but it directly influences the statistical problem that the model must solve. Fabric images can vary substantially because of illumination, camera distance, viewing angle, lens characteristics, shadows, background conditions, fabric tension, orientation, and surface reflectance. As a result, a model trained exclusively on highly controlled images may learn acquisition-specific patterns rather than features that genuinely characterize the fabric. Appropriate preprocessing methods therefore include image resizing, color-space conversion, noise reduction, contrast enhancement, illumination normalization, segmentation, background removal, and feature normalization.

An important methodological consideration is that data augmentation cannot replace genuine diversity in image acquisition. For example, rotating an image or increasing its brightness does not necessarily reproduce the physical effects of changes in lighting geometry, fabric tension, folds, yarn orientation, or camera optics. Consequently, industrial validation should be treated as an independent experimental condition rather than simply as another form of augmented training data.



Figure 2. Image-acquisition and preprocessing chain. The objective is not merely to make images visually attractive, but to reduce nuisance variation while preserving fabric-specific evidence.

4. Handcrafted texture, color and structural descriptors

Before deep learning became dominant, fabric analysis relied heavily on engineered texture descriptors. Gray-Level Co-occurrence Matrix (GLCM) statistics characterize spatial relationships between gray-level values and can capture properties such as contrast, correlation, energy, and homogeneity. Local Binary Patterns (LBP), in contrast, describe local intensity patterns and are particularly useful for analyzing fine-scale surface texture. Color descriptors can further complement these features when differences in material or surface appearance are associated with chromatic information.

These traditional methods remain scientifically relevant even when they do not outperform modern deep-learning models. Their advantages include interpretability, low computational cost, effectiveness with small datasets, and the ability to provide explicit hypotheses about the physical surface characteristics that distinguish different fabrics. Recent multimodal textile-identification research likewise uses complementary representations, including co-occurrence histograms, LBP features, and albedo-dependent feature maps, rather than relying exclusively on raw-image inputs [6]. The key research question, therefore, is not simply whether handcrafted features or deep learning is preferable. Instead, it is how complementary representations can be effectively combined and evaluated under controlled and reproducible conditions.

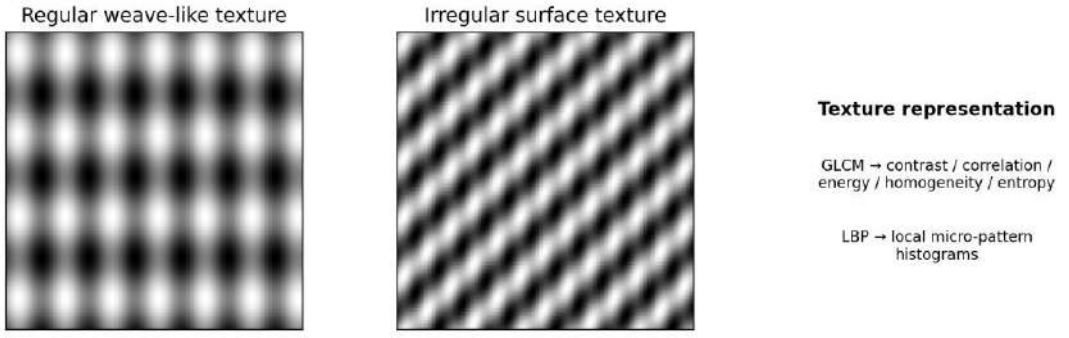


Figure 3. Representative texture-analysis concept. Regular and irregular surface structures generate different statistical and local-pattern representations, motivating multimodal texture encoding.

Table 2. Complementarity of common fabric-image representations.

Representation	Captures	Strength	Weakness
GLCM	Spatial gray-level relationships	Interpretable; useful for regular texture	Sensitive to quantization and imaging conditions
LBP	Local micro-patterns	Simple; efficient; useful for fine texture	Limited global structural context
Color descriptors	Chromatic distribution	Useful for color/material discrimination	Affected by illumination and white balance
Structural/weave descriptors	Repeated geometry and yarn arrangement	Closer to physical weave organization	Can be difficult under deformation
Deep features	Hierarchical visual patterns	Learns complex representations	Data-hungry; less transparent

5. Fabric material identification

Material identification is conceptually distinct from defect detection. A fabric may appear completely normal yet still be difficult to classify because different materials can share similar colors, weave structures, or surface appearances. Cheng et al. investigated textile material identification using a lightweight DenseNet-based model and introduced a five-material image dataset. Their AONet contained only 0.64 million parameters and achieved a reported recognition accuracy of 95.61% [5]. The study is

significant because it demonstrates the potential of lightweight deep-learning models for material identification when training data are limited. However, the relatively small dataset also raises concerns about the external validity and generalizability of the reported results.

Material-identification models should therefore be evaluated using genuinely unseen samples and, where possible, unseen fabric sources or production rolls. Randomly dividing highly correlated image patches between training and test sets can lead to overly optimistic performance estimates because visually similar or near-duplicate regions from the same fabric may appear in both subsets. A rigorous review should therefore distinguish between image-level random splitting and source-level or roll-level separation, with the latter providing a more realistic assessment of a model's ability to generalize to new fabric sources.

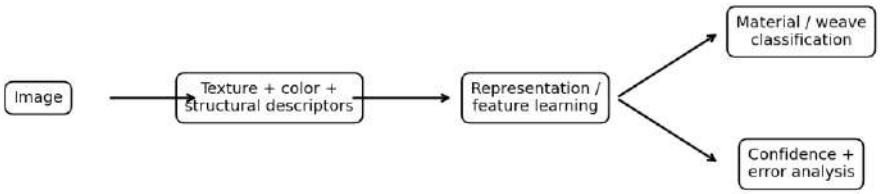


Figure 4. Conceptual pathway for fabric-material and weave identification, showing complementary surface representations and confidence-aware classification.

Table 3. Representative evidence for material and weave identification.

Study / stream	Task	Representative approach	Critical observation
Cheng et al. (2023) [5]	Material identification	DenseNet variant / AONet	Lightweight architecture is promising, but small-sample generalization remains a concern
Hussain et al. (2020) [4]	Woven-pattern recognition	Deep CNN + augmentation/transfer learning	Shows value of learned weave representations
Gnoupa et al. (2025) [6]	Textile identification	Multimodal texture inputs + deep learning	Supports combining macro/micro texture cues

6. CNNs and deep representation learning for textile surfaces

Convolutional neural networks (CNNs) transformed fabric analysis by enabling models to learn hierarchical visual features directly from image data. Hussain et al. applied deep CNNs to the recognition and classification of woven-fabric patterns, addressing the limitations of manual inspection and traditional pipelines based on handcrafted features [4]. For fabric-defect inspection, Mei et al. introduced a multi-scale convolutional denoising autoencoder, while Li et al. developed a wide-and-compact CNN designed to balance detection accuracy with model size and computational efficiency [2,3].

The significance of CNNs extends beyond their potential to achieve higher classification or detection accuracy. Their hierarchical representations can capture features ranging from edges and fine-scale textures to repetitive patterns, defect boundaries, and larger structural

configurations. Nevertheless, strong performance on a single dataset does not by itself demonstrate industrial superiority. Meaningful comparisons should also consider data efficiency, parameter count, inference speed, robustness to variations in imaging conditions, calibration, and performance on genuinely independent fabric sources.

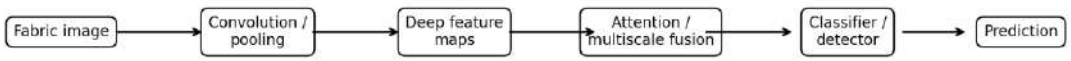


Figure 5. Generic deep-learning pathway for fabric inspection. Modern architectures may add attention, multi-scale feature fusion or transformer modules, but the central issue remains the quality and independence of the training evidence.

7. Fabric defect classification and detection

Defect detection is the most extensively studied area within the literature. Mei et al. addressed defect detection and localization using an unsupervised reconstruction approach, thereby reducing the need for large numbers of manually labeled defective samples [2]. Subsequent studies increasingly adopted supervised CNNs and compact architectures, while YOLO-based detectors have been widely explored for real-time defect localization. More recent research continues this development, including the use of fine-tuned YOLO-12 models for fabric-defect detection [7].

A clear distinction should be maintained between classification, object detection, and segmentation. Image-level classification determines whether a particular defect category is present in an image or image patch. Object detection goes further by identifying the defect and estimating its location, typically through bounding boxes. Segmentation provides more detailed information by identifying the pixels or boundaries associated with the defect. These tasks differ not only in their annotation requirements and computational demands but also in their industrial applications. In a production environment, knowing the defect category alone may be insufficient; its location and spatial extent may also be required because defect severity and subsequent quality-control decisions can depend on its size and position.

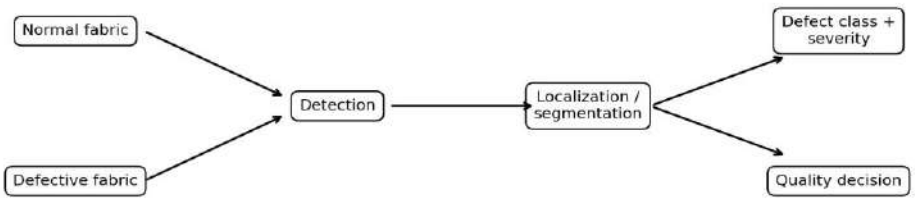


Figure 6. Defect-inspection pathway from normal/defective recognition to localization, severity estimation and quality decision.

Table 4. Distinction between fabric-inspection tasks and their outputs.

Task	Output	Typical metrics	Industrial meaning
Binary classification	Defective / non-defective	Accuracy, precision, recall, F1	Basic screening
Multi-class classification	Defect category	Macro-F1, per-class recall	Defect diagnosis
Object detection	Boxes + classes	mAP, precision, recall	Locate actionable defects
Segmentation	Pixel mask	IoU, Dice, pixel F1	Estimate defect extent
Quality assessment	Grade / accept-reject	Class metrics + calibration	Production decision

8. Localization, segmentation and repeated-pattern analysis

Localization is an important step beyond image-level recognition. Repetitive textile patterns present a particular challenge because periodic structure is an inherent characteristic of normal fabric. An inspection model must therefore distinguish legitimate pattern repetition from genuine structural anomalies. Pixel-level segmentation is especially useful when the size, shape, or affected area of a defect contributes to quality assessment and grading.

Recent studies have investigated repeated-pattern analysis and segmentation, reflecting a broader shift from the question of "Is a defect present?" to the more precise question of "Where exactly is the defect?" However, the literature remains inconsistent in its annotation practices. Bounding boxes, coarse region annotations, patch-level labels, and pixel-level masks represent different levels of ground-truth information and should not be treated as interchangeable. Future benchmark datasets should therefore clearly specify the annotation level used and, where possible, report inter-annotator agreement to establish the reliability of the ground truth.

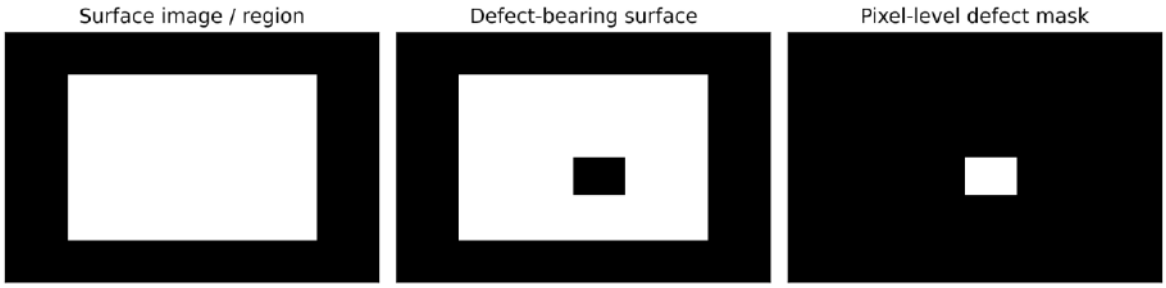


Figure 7. Conceptual transition from surface recognition to pixel-level defect localization. The same fabric image can support different levels of supervision depending on the annotation protocol.

9. Attention, transformers and multimodal textile representations

Attention mechanisms and transformer-based architectures have expanded the ability of deep-learning models to integrate local and global visual information. This is particularly relevant for fabric images, where local yarn irregularities may provide important diagnostic cues, while the overall weave structure offers essential contextual information. However, architectural novelty alone is becoming a relatively weak basis for a PhD contribution, as many recent studies already propose modified attention mechanisms, multi-scale architectures, and new variants of YOLO-based detectors.

A more strategically important direction is multimodal representation. Gnoupa et al. combined macro- and micro-texture information, co-occurrence histograms, LBP features, and albedo-dependent feature maps within deep-learning architectures for automated textile identification [6]. Their work demonstrates that complementary surface representations can be integrated rather than

requiring a single model to infer every relevant property directly from the same raw image. This principle provides a promising foundation for extending textile analysis beyond material identification toward an integrated framework that jointly represents fabric identity, texture, defects, and overall quality.

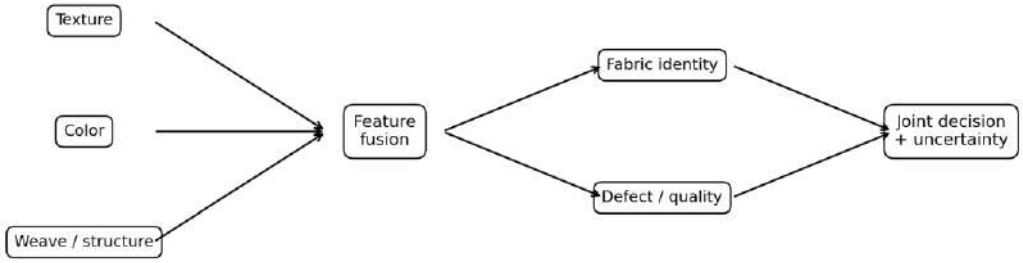


Figure 8. Multimodal feature-fusion concept for textile inspection. Texture, color and structural information are represented separately and fused before joint prediction and uncertainty estimation.

10. Dataset quality, leakage and evaluation validity

Dataset design may be more important than making another marginal modification to the model architecture. Recent review studies explicitly identify the quality and structure of fabric-defect datasets as a major challenge in the field [1]. Publicly available datasets are useful for benchmarking, but many contain only a limited range of defect categories, highly controlled backgrounds, repeated samples, or data distributions that differ substantially from those encountered in production environments. Consequently, results obtained from different datasets should not be compared as though they represent the same level of difficulty or generalization.

Patch-based datasets create an additional and potentially serious risk of data leakage. When multiple patches are extracted from the same fabric roll or continuous image, randomly splitting those patches at the image level can place highly similar or near-identical texture patterns in both the training and test sets. This can produce artificially high performance estimates. A rigorous PhD-level study should therefore use group-aware data splitting, with groups defined by fabric roll, physical sample, production batch, or acquisition session. Even stronger evidence of generalization can be obtained through external validation using a second dataset or images collected under a different acquisition environment.

Table 5. Experimental threats that should be controlled in a PhD-level fabric-inspection study.

Validity threat	How it arises	Recommended control
Patch leakage	Train/test patches come from the same fabric source	Group split by roll/sample/session
Class imbalance	Common defects dominate rare defects	Macro-F1, balanced sampling, per-class recall
Domain shift	Different cameras/lighting/backgrounds	Cross-session or cross-device testing
Annotation uncertainty	Ambiguous defect boundaries/classes	Multiple annotators and agreement analysis
Metric inconsistency	Different papers report different measures	Predefine a common metric suite
Overfitting to benchmark	Repeated tuning on one public dataset	Hold-out external dataset

11. Robustness, explainability and industrial deployment

Industrial inspection requires more than achieving high test accuracy on a single benchmark dataset. A practical model must remain reliable under realistic variations in lighting, fabric orientation, scale, focus, motion blur, camera position, and background conditions. Robustness should therefore be evaluated explicitly through controlled stress tests and, where possible, cross-domain validation. The objective is not to demonstrate that performance remains completely unchanged under all conditions, but to determine how rapidly, consistently, and predictably performance deteriorates as conditions become more challenging.

Explainability is equally important because an industrial quality-control system may need to provide evidence for why a fabric was classified as defective or rejected. Saliency maps, defect masks, feature-attribution methods, and confidence estimates can help determine whether model decisions are based on physically meaningful fabric characteristics or irrelevant background and acquisition artifacts. For deployment, evaluation should also extend beyond predictive accuracy to include inference latency, model size, memory consumption, and hardware requirements.

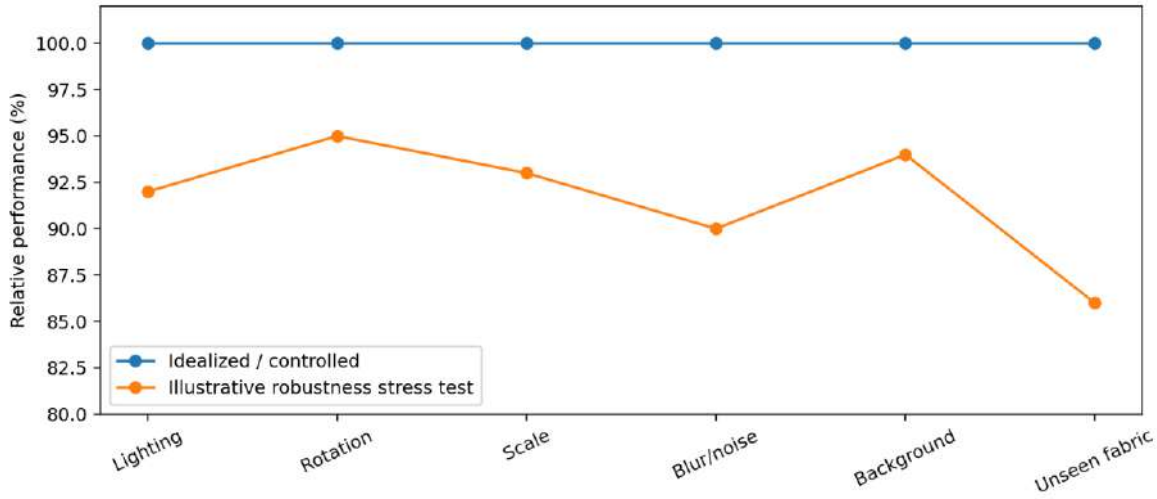


Figure 9. Example robustness-evaluation concept. The plotted values are illustrative, not experimental results; the figure demonstrates how performance degradation can be quantified under controlled perturbations.

12. Critical synthesis: what the literature has established—and what it has not

The existing literature provides several findings with reasonable consistency. First, computer vision can automate important aspects of textile inspection. Second, learned visual representations generally achieve strong performance in fabric recognition and defect analysis. Third, lightweight network architectures are particularly important when inspection must be performed close to the production line. Fourth, multimodal texture representations are promising because textile surfaces contain complementary information across different spatial scales [3–6]. Fifth, defect inspection is increasingly moving beyond simple classification toward more precise localization and segmentation [1,2,7].

Despite these advances, the literature does not yet provide a widely accepted integrated framework that jointly addresses material identity, weave and texture, surface condition, defect localization, quality grading, uncertainty, and realistic variations in the operating environment. Most existing studies are optimized for a specific task, dataset, or model architecture. Consequently, the evidence is relatively strong for individual inspection components but remains much weaker when these components must be combined into a unified system for industrial decision-making.



Figure 10. Research agenda emerging from the critical synthesis: data provenance → multimodal representation → multi-task prediction → robustness/explainability → industrial validation.

13. Research gap and defensible PhD direction

The most defensible research gap is not the development of yet another CNN, YOLO variant, or attention mechanism, as these areas are already extensively studied. A stronger PhD contribution would be an integrated, multimodal, and experimentally rigorous framework that treats fabric inspection as a hierarchy of related decisions: identifying the fabric type, characterizing its texture and weave, determining whether its surface is normal or defective, identifying the defect category, locating the defect, estimating its severity, and quantifying the confidence of each decision.

Such a study should begin with a carefully designed dataset that records fabric identity, weave and texture characteristics, color, acquisition conditions, normal or defective status, defect category, and defect localization where feasible. Classical descriptors such as GLCM and LBP should not simply be discarded; instead, they can provide interpretable baselines and potentially serve as complementary inputs to learned representations. Deep-learning features should then be systematically compared with, and where appropriate fused with, these handcrafted representations. The resulting system should be evaluated using group-aware data splits, cross-condition stress testing, external validation, uncertainty analysis, and practical deployment benchmarks. This evaluation strategy would provide stronger evidence of generalization and industrial usefulness than performance on a single controlled test set.

Table 6. A defensible PhD research agenda derived from the reviewed evidence.

Proposed contribution	Why it is stronger than an architecture-only contribution	How it can be tested
Curated multimodal dataset	Addresses evidence quality and task fragmentation	Independent roll/session/source splits
Feature fusion	Tests whether complementary textile cues add information	Ablation of each modality
Multi-task inspection	Links identity, texture, defects and quality	Joint vs single-task comparison
Robustness protocol	Measures industrial generalization	Lighting/orientation/device/domain stress tests
Explainability + uncertainty	Supports trustworthy quality decisions	Calibration, attribution and error analysis
Deployment benchmark	Connects accuracy to factory constraints	Latency, memory, model size and throughput

14. Conclusions

Automated fabric inspection has evolved from handcrafted texture analysis and conventional classification methods to deep CNNs, lightweight networks, object detectors, segmentation techniques, attention mechanisms, and multimodal feature learning. Although the literature demonstrates considerable technical capability, it remains fragmented: fabric material identification, weave recognition, defect detection, and quality assessment are generally investigated as separate problems. Recent research on multimodal representations offers a particularly promising direction because textile surfaces contain complementary information at the macro-, micro-, color, and structural levels. The next meaningful step is therefore not simply the development of another, more sophisticated detector, but the design of a rigorous integrated system that combines these complementary sources of information while explicitly addressing dataset leakage, domain shift, and variability in industrial imaging conditions. A PhD study built around this gap could make a stronger scientific contribution by improving the evidence base, experimental methodology, and industrial validity of intelligent textile inspection rather than relying primarily on incremental architectural modifications.

REFERENCES

He, F., Xin, B., Zhan, Z., Yang, W., & Liu, J. (2026). A review of fabric defect detection in textile manufacturing. The Journal of The Textile Institute, 117(2), 327–345. DOI: 10.1080/00405000.2025.2502188.

Mei, S., Wang, Y., & Wen, G. (2018). Automatic fabric defect detection with a multi-scale convolutional denoising autoencoder network model. Sensors, 18(4), 1064. DOI: 10.3390/s18041064.

Li, Y., et al. (2019). Automatic fabric defect detection with a wide-and-compact network. Neurocomputing, 329, 329–338. DOI: 10.1016/j.neucom.2018.10.070.

Iqbal Hussain, M. A., Khan, B., Wang, Z., & Ding, S. (2020). Woven fabric pattern recognition and classification based on deep convolutional neural networks. Electronics, 9(6), 1048. DOI: 10.3390/electronics9061048.

Cheng, L., Yi, J., Chen, A., Zhang, Y., & Hou, C. (2023). Fabric material identification based on Densenet variant networks. The Journal of The Textile Institute, 114(10), 1527–1538. DOI: 10.1080/00405000.2022.2144600.

Enow Gnoupa, M. G., Augousti, A. T., Duran, O., Lanets, O., & Liaskovska, S. (2025). Multimodal feature inputs enable improved automated textile identification. Textiles, 5(3), 31. DOI: 10.3390/textiles5030031.

Ahmad, W., Ashraf, R., Mahmood, T., Hussain, T., & Abid, M. H. (2026). Fabric defect detection using fine-tuned Yolo-12. PLOS ONE. DOI: 10.1371/journal.pone.0353550.

Mao, J., & Hong, Z. (2025). Review of YOLO-series methods and their development for visual detection. Sensors.

Guo, et al. (2024). Intelligent quality control of surface defects in fabrics: review and research directions. IEEE Access.

Carrilho, et al. (2024). Recent computer-vision approaches for textile/fabric inspection: a survey. Electronics.