PROJECT ABSTRACT

Fabric Defect Detection using Deep Learning

A CNN-based textile inspection system that detects weaving defects (holes, slubs, misweaves, broken ends, stains, snags) on fabric rolls. A transfer-learned ResNet-50 classifier trained on the published TILDA textile dataset is paired with Grad-CAM localization and a web inspection console for garment-manufacturing quality control.

Project type: Software

Availability: Built to order

Suitable branches: AI & Machine Learning, Computer Science, Textile Engineering

Technologies: Python, PyTorch, ResNet-50, TILDA dataset, OpenCV, Flask, ONNX

Price: Request a quotation

Prepared by Projectech · projectech.in — for academic project reference.

1. Abstract

This project presents a deep-learning system for automated fabric defect detection in garment manufacturing. A ResNet-50 convolutional neural network, pre-trained on ImageNet and fine-tuned on the published TILDA textile texture dataset, classifies fabric frames as defect-free or as one of seven weaving defect types (hole, slub, misweave, broken end, oil spot, snag, thread error). Grad-CAM heatmaps are thresholded into bounding boxes that localize each defect with a confidence score, and a web-based inspection console provides single-frame inspection, roll batch review and a defect-rate dashboard. The system is built to order: the training notebook, trained model, inspection application, report, presentation and viva preparation are prepared fresh for each order, and all performance figures are documented honestly as design targets until measured on the order's own training run.

2. Problem Statement

Fabric inspection in garment manufacturing is still largely manual: an operator watches the roll unspool and marks defects by eye. The process is tiring and inconsistent, subtle defects slip through, and no quantitative defect-rate log is kept per roll — so quality problems surface only as rejected garments or customer complaints downstream.

This project provides an automated alternative: a camera plus a convolutional neural network that classifies and localizes weaving defects on fabric frames, with a console that logs defect rates per roll against a configurable rework threshold.

3. Objectives

- Classify fabric frames into defect-free or one of seven weaving defect types using a transfer-learned CNN.
- Localize each defect with a bounding box, defect-type label and confidence score via Grad-CAM.
- Provide a web inspection console: single-frame inspection, roll batch queue and defect-rate dashboard.
- Train and evaluate reproducibly on the published TILDA textile dataset with imbalance-aware training.
- Document every performance figure honestly — design targets up front, measured metrics in the report.
- Deliver a complete viva kit: report PDF, presentation and Q&A; preparation.

4. System Architecture

The system is organized in three layers. The **data layer** holds the TILDA textile texture dataset (approximately 3,200 fabric images, 768×512, 8 fabric types) plus a supplementary mill collection for slub, misweave, snag and broken-end examples, with a roll-stratified 70/15/15 train/validation/test split. The **model layer** is a ResNet-50 backbone (ImageNet-pretrained, fine-tuned) with a global-average-pooling plus dense classification head and a Grad-CAM localization stage, trained with focal loss and class weights. The **application layer** is a Flask-backed web console with an HTML/CSS/JS front end for inspection, batch review and dashboarding.

5. Technologies Used

- Python 3
- PyTorch
- ResNet-50 (transfer learning)
- TILDA textile texture dataset
- OpenCV and Albumentations
- Focal loss with class-weighted training
- Flask inspection application
- HTML5, CSS and JavaScript console
- ONNX model export

6. Working Principle

- A fabric roll frame is captured (or a swatch photo uploaded) and resized to 512×512 with the same normalization used in training.
- The ResNet-50 backbone extracts texture features; the dense head outputs probabilities over 8 classes (defect-free + 7 defect types).
- For each predicted defect class, a Grad-CAM heatmap is computed from the final convolutional layer.
- The heatmap is thresholded and connected-component labeled into bounding boxes; boxes below 0.50 confidence are discarded.
- Each surviving box is drawn on the frame with its defect-type label and confidence; the frame receives a PASS or DEFECTIVE verdict.
- Frame results aggregate per roll into a defect-rate log (defects per 100 m); rolls crossing the rework threshold are flagged.
- The dashboard renders trends across rolls, and the training notebook regenerates per-class precision/recall and the confusion matrix.

7. Key Features

- ResNet-50 defect classifier fine-tuned to an 8-way softmax (defect-free + 7 weaving defect types).
- Grad-CAM defect localization: bounding boxes with defect-type labels and confidence scores.
- Inspection web console with bounding-box overlays, ranked detection list and per-frame statistics.
- Roll batch queue with PASS/DEFECTIVE verdicts and defect counts for quick triage.
- Defect-rate dashboard tracking defects per 100 m across rolls with a configurable rework threshold.
- Imbalance-aware training: focal loss with class weights and a roll-stratified split.
- Complete training notebook with per-class metrics, confusion matrix and training curves.

8. Applications

- In-process quality control for woven fabric rolls in garment manufacturing.
- Supplier quality auditing: quantitative defect-rate evidence per roll.
- Textile and apparel engineering laboratories and coursework.
- Demonstration of transfer learning and model interpretability for AI/ML curricula.

9. Prerequisites

- Basic knowledge of Python programming.
- Understanding of convolutional neural networks and image classification.
- Familiarity with transfer learning concepts (helpful, covered in the report).

- A computer capable of running the training notebook (GPU recommended; CPU workable for inference).
- Even, diffuse lighting and a fixed camera distance when capturing real fabric frames.

10. Limitations and Future Scope

- TILDA covers mostly patternless woven fabrics — bold printed or highly patterned textiles are a known hard case.
- Detection quality depends on capture conditions; harsh shadows or motion blur degrade results.
- Subtle defects (snags, faint oil spots on dark fabric) have the lowest recall and are called out honestly.
- The ~92% recall and <5% false-alarm figures are design targets; the report documents measured values after training.
- CPU inference (~120–180 ms/frame) suits roll-sample inspection, not full-speed line scanning without a GPU.
- Future scope: printed-fabric training data, real-time line integration, mobile capture app, multi-camera roll coverage.

11. Conclusion

Fabric Defect Detection using Deep Learning offers a practical, student-buildable approach to textile quality control: a transfer-learned ResNet-50 classifier trained on the published TILDA dataset, Grad-CAM localization for interpretable bounding boxes, and a working inspection console with batch review and defect-rate dashboarding. The project is honest about its limits — patterned fabrics, capture conditions and the design-target status of its metrics are stated up front — which makes it both a credible academic submission and a realistic starting point for real mill deployment.